UNIVERSITÀ COMMERCIALE "LUIGI BOCCONI"
PHD SCHOOL

PhD Program in Economics and Finance
Cycle: 36
Disciplinary field: SECS/P06

Essays on Information, Politics, and Digital Media

Advisor: **Prof. Carlo Schwarz**
Co-advisor: **Prof. Guido Tabellini**

PhD Thesis by
Julian Streyczek
ID number: 3125418

Year 2026

To all those
who aim to inspire
the best in others

Acknowledgements

I am deeply grateful for the time, energy, and trust invested by my advisors, which allowed me to learn so much from each of you. To Carlo Schwarz, thank you for your constant availability and unmatched enthusiasm, which helped me keep my spirits up. To Guido Tabellini, thank you for your thoughtful guidance, which helped me stay on the right path. To Luca Braghieri, thank you for sharing profound insights about work and life, both inside and outside academia, which helped me reflect on this experience. I could not have asked for a better team.

I owe special thanks to Jim Snyder for sponsoring my research visit to Harvard and for making me feel welcome from day one. I am also grateful to Elliott Ash for helping examine my thesis, providing valuable feedback, and generously sharing data. I would like to thank Bocconi University, Fondazione Invernizzi, and the German Academic Scholarship Foundation for the generous funding of my research activities.

I am grateful to the Economics faculty at Bocconi, as well as to my fellow PhD students for the countless corridor conversations and coffee chats that sparked inspiring ideas and brightened so many days. Your availability, kindness, and continuous feedback were crucial throughout this journey. It truly takes a village to raise a PhD student.

Finally, my heartfelt thanks go out to my friends and loved ones for providing a listening ear and thoughtful judgment over all these years, and for reminding me time and again that the world has many more colors than it may sometimes seem.

Introduction

Digital technologies are transforming how individuals access and engage with information, with consequences that can be beneficial or detrimental to society. While broader access to information offers clear advantages, substantial evidence also documents negative side effects, including the spread of misinformation, increased distraction, and the decline of traditional information media. These developments have wide-ranging implications for individuals' opinions and behavior, from voters selecting competent representatives to researchers identifying relevant scientific questions. Given the pace and scope of these transformations, many important questions about the societal effects of digital media remain unanswered.

This dissertation contains three chapters that empirically examine distinct settings in which contemporary digital media shape individuals' opinions and behavior. In the first chapter, I show that the introduction of paywalls on US newspaper websites reduced political knowledge and electoral participation. By restricting access to information, paywalls pushed individuals who were unwilling or unable to pay for subscriptions away from traditional news, lowering their engagement with politics. The second chapter develops a novel instrumental variable to demonstrate that economists' usage of Twitter (now X) increased publications and citations. The platform facilitated professional networking and encouraged the selection of research topics with greater relevance, both within academia and among the broader public. In the third chapter, I show that disruptive protests by a German environmental activist group failed to shift public opinion on climate policy despite generating substantial media coverage, suggesting limits to the persuasive power of media attention alone.

Contents

1 Political Effects of Newspaper Paywalls **4**

1 Introduction 5
2 Data 9
3 Effect on news consumption 15
4 Effect on political knowledge and electoral participation 21
5 Conclusion 46
A Appendix 50

2 #Science: How Twitter Shapes Economic Research **60**

1 Introduction 61
2 Data 65
3 Empirical Strategy 74
4 Effect of Twitter on Scientific Output 83
5 Effect of Twitter on Scientific Network 90
6 Effect of Twitter on Research Topics 94
7 Conclusion 99
A Appendix 103

3 Disruptive Activism and Environmental Policy Support **114**

1 Introduction 115
2 Data 117
3 Empirical strategy 122
4 Results 124
5 Conclusion 129
A Appendix 132

Chapter 1

Political Effects of Newspaper Paywalls

Abstract

I study how the introduction of paywalls on newspaper websites in the early 2010s affected political knowledge and electoral participation in the United States. Exploiting the staggered adoption of paywalls across newspapers, I first show that paywalls reduce readership of online news content, decreasing page views by a 25–30 percent on average. Next, I analyze the impact of county-level exposure to paywalls, measured as the share of online news consumption affected by paywalls, on political knowledge and electoral participation. I show that survey respondents in counties most exposed to paywalls become 1.6–2.4 percent less likely to correctly answer factual knowledge questions on party affiliations of political representatives and majorities in legislative bodies. This effect is driven by declines in knowledge about both national and regional politics, and is most pronounced among lower-income, less-educated individuals, and those least likely to subscribe to paid news. Additionally, higher paywall exposure reduces electoral participation among these non-subscribers by approximately 2 percent, particularly in federal and state congressional elections. These findings suggests that paywalls restrict access to politically relevant information, prompting substitution toward free alternatives with less political content. My results underscore the importance of easy access to high-quality news for democratic processes.

Keywords: News, Information, Politics, Paywalls
JEL Codes: D72, L82, Z13

1 Introduction

Over the past 15 years, paywalls have transformed much of newspapers' online content from a public good (back) into a club good: By the 2000s, most newspapers offered free online access to their articles, relying on advertising to generate revenue. However, as the shift to online readership undermined traditional print revenue streams, many publishers implemented paywalls on their websites – subscription models that require a monthly fee for full access – ending more than a decade of free news. By 2020, more than 80 percent of the most popular US newspapers had implemented such paywalls. Although these models help publishers remain profitable, they also restrict the availability of information, especially for those unwilling or unable to pay.

Access to information can shape voter knowledge and participation, which in turn can influence the accountability of elected representatives and the allocation of public funds (Besley and Burgess, 2002; Strömberg, 2004; Ferraz and Finan, 2008; Snyder Jr and Strömberg, 2010). Past shifts in digital media availability have prompted consumers to move away from news-intensive outlets, with significant consequences for voting patterns (Gentzkow, 2006; Drago et al., 2014; Falck et al., 2014; Campante et al., 2018; Gavazza et al., 2019).

In this paper, I study how the emergence of newspaper paywalls affected political knowledge and electoral participation, exploiting the staggered introduction of paywalls across time and geographic regions for causal identification. My analysis proceeds in two parts: First, I show that paywalls reduce consumption of newspapers' online content. Second, I examine the effect of paywalls on political knowledge and electoral participation by creating a measure of county-level exposure to paywalls and linking it to large-scale survey data.

In my analysis, I combine data on newspaper consumption, paywall introductions, and survey data on political knowledge and voting. I focus on the 100 largest daily newspapers in the United States based on print circulation in 2010, the year before the introduction of most paywalls, provided by the *Alliance for Audited Media*. For the largest 80 of these, I also measure online consumption using the daily number of page views between 2010 and 2017, collected by *Alexa Internet*. I assemble paywall introduction dates by checking mentions in news articles or historical website snapshots in the *Internet Archive*. Further, I measure individual-level knowledge and voting using the *Cooperative Election Study* [1], a yearly, nationally representative survey of the adult US population. I construct a political

[1] Formerly called *Cooperative Congressional Election Study (CCES)*.

knowledge index using questions that ask participants to specify the party affiliation of politicians and to name the majority party in US legislative bodies. I also use self-reported as well as validated data on participation in elections.

In the first part of my analysis, I identify the causal effect of paywalls on website page views using the sequential adoption of paywalls across websites. I compare newspapers that introduced a paywall between 2010 and 2017 to those that introduced paywalls later or never. I control for newspapers' market characteristics to account for trends among different audiences, and use estimators that are robust to heterogeneous treatment effects across newspapers or time. To corroborate the identifying assumption that trends in page views would be parallel in the absence of paywalls, I conduct two sets of robustness checks. First, I verify that selection of newspapers into paywalls is unlikely to drive the results, as pre-paywall trends in page views are parallel, inclusion of controls has minimal impact, and the result is robust to varying the composition of the control group. Second, I show that potential spillover biases, arising from readers' substitution of paywalled with non-paywalled outlets, are negligible, by conducting a series of robustness checks excluding newspapers that are expected to generate the highest spillovers.

I find that paywalls reduce the number of page views for paywalled newspapers by 25–30 percent, on average. These effects materialize in the first two months after the paywall and are persistent over the following years, suggesting that paywalls considerably displaced readers from mainstream US newspapers.

In the second part of my analysis, I start by quantifying the relevance of paywalls for survey respondents based on their county of residence. For this purpose, I construct a measure for counties' exposure to paywalls over time. This measure combines the timing of paywall introductions with newspapers' initial subscription rates in counties per household, which I validate as a strong predictor for regional online readership. Specifically, I define paywall exposure as the share of households subscribed to newspapers that have a paywall in place, which proxies the number of online readers of newspapers affected by paywalls. Then, I exploit variation in this measure using a staggered difference-in-differences design, classifying counties as treated once their paywall exposure exceeds the 75th percentile.[2]

Identification relies on the assumption that no other contemporaneous shock differentially influences political knowledge in treated counties at the time of paywall rollout, conditional on fixed effects and controls. I provide extensive evidence in support of this assumption. First, all specifications include state-by-year fixed effects that capture any

[2]The results hold when varying this threshold.

time-varying shocks to the regional political landscape. Second, I include survey respondent characteristics interacted with year fixed effects that capture changes in outcomes of certain demographics unrelated to paywalls, such as preferences for digital versus analog platforms. These variables also capture potential measurement error of paywall exposure correlated with regional characteristics. Third, the staggered nature of the treatment helps separate the effect of paywalls from aggregate trends.

I find that in counties with high paywall exposure, individuals are on average 1.2–1.8 percentage points (1.6–2.4 percent) *less* likely to answer political knowledge questions correctly, compared to a baseline accuracy of 74.5 percent correct answers. The event study rejects the hypothesis that prior trends in outcomes are driving the results. Interestingly, the declines in knowledge pertain to politicians and majorities on both state and federal level. Furthermore, heterogeneity analyses suggest that the effects are driven by the by groups most likely to be sensitive to price increases, namely lower-income and less-educated individuals.

As a placebo check, I show that these reductions in political knowledge can be explained by differences in the likelihood of purchasing news subscriptions. I use an additional survey to I predict subscription propensities based on individual characteristics. The effects of paywall exposure are entirely concentrated among those unlikely to pay for news, with no effect among individuals predicted to subscribe.

Finally, I investigate the effects of paywall exposure on participation in elections. Effects in the full sample are consistently negative but statistically insignificant. However, restricting the analysis to individuals predicted not to pay for news, and excluding highly partisan counties, yields larger and statistically significant reductions in participation. Both self-reported and validated turnout data indicate approximately a 2 percent decline in participation within this subgroup, driven primarily by lower turnout in federal and state congressional elections.

My results are most consistent with the interpretation that paywalls reduce news consumption from outlets providing extensive coverage of contemporary politics, particularly among individuals unwilling or unable to pay. These readers likely substitute toward free alternatives with less political focus, such as free news websites and social media, or disengage from news more broadly. As a result, they may pay less attention to politics, reducing political knowledge and ultimately electoral participation. While the absence of detailed individual-level consumption data prevents me from directly observing who reduces consumption of paywalled news and which substitutes they choose, several observations support this interpretation: First, substitution to newspapers' print versions is unlikely to explain the decline in page views, as aggregate print readership continued to decrease over

the sample period, and case studies find little evidence for substitution to print (e.g., Pew Research Center, 2023; Pattabhiramaiah et al., 2019). Second, the absence of paywall effects among likely newspaper subscribers contradicts the hypothesis that the effects are driven by changes to newspapers' content alongside paywall introductions. Instead, the result that declines in knowledge are concentrated entirely among likely non-subscribers supports the conclusion that switching away from traditional news sources drives the effect. Third, the finding that individuals with lower education and lower income are most affected by paywalls may reflect the monetary nature of the paywall barrier. Such individuals might have a lower willingness to purchase subscriptions due to relatively higher monetary costs.

My findings confirm that easy access to news is important for political knowledge and participation. Moreover, they underscore that in a predominantly profit-driven news market, economic shocks can ultimately affect democratic outcomes. To counteract these effects, modern democracies may benefit from guaranteeing broad and equitable access to relevant information.

This paper contributes to several streams of literature in economics and political science. First, it relates to existing work on news, knowledge, and accountability. Notable papers include the diffusion of television (Gentzkow, 2006; DellaVigna and Kaplan, 2007; Durante et al., 2019) and high-speed internet (Falck et al., 2014; Campante et al., 2018; Gavazza et al., 2019), which demonstrate that biased media can persuade voters, while substitution away from media with high news content can reduce political engagement. Studies on the entry and exit of newspapers underline that access to news can improve political participation (Gentzkow et al., 2011; Drago et al., 2014; Cagé, 2020; Gao et al., 2020; Djourelova et al., 2024), which may affect the performance of political representatives and the allocation of public resources (Besley and Burgess, 2002; Strömberg, 2004; Ferraz and Finan, 2008; Snyder Jr and Strömberg, 2010; Ash and Galletta, 2023).

To the best of my knowledge, I provide the first evidence on political effects of newspaper paywalls. Unlike most historical trends in media innovation and digitization, paywalls *increased* barriers to information. Moreover, I show that the *monetary* nature of these barriers disproportionally affected lower-income and less-educated individuals.

Second, my paper relates to the literature studying the news industry. Prior studies have shown that entry and exit of newspapers affect existing newspapers' revenues, prompting adjustments in topics, slant, or journalistic intensiveness (George and Waldfogel, 2006; Gentzkow et al., 2014; Angelucci and Cagé, 2019). Digitization has accelerated these pressures, as newspapers faced declining print advertising revenues and increased online competition (Seamans and Zhu, 2014; Bhuller et al., 2024). In response, many newspapers

have introduced paywalls as part of their subscription models. Existing case studies indicate that paywalls may cause newspapers to lose readers (Cook and Attari, 2012; Chiou and Tucker, 2013; Pattabhiramaiah et al., 2019), although some outlets have succeeded in offsetting these losses through increased subscriptions (Chung et al., 2019; Aral and Dhillon, 2021). In a broader study, Kim et al. (2020) find that paywalls on 42 large US-based newspapers reduced page views by 30 percent on average.

I add to this literature in three ways. First, I provide the most comprehensive empirical study to date on how paywalls affect news consumption, studying the 100 largest US-based newspaper websites at the time. Second, I identify the causal effect of paywalls using state-of-the-art econometric methods. Third, I demonstrate that the effects on the news industry are both economically significant and persistent.

The remainder of this paper is organized as follows. In Section 2, I introduce my data and provide background on paywalls. In Section 3, I show how paywalls affected news consumption, and in Section 4, I analyze downstream effects on political knowledge and electoral participation.

2 Data

2.1 Newspapers

My analysis focuses on the 100 largest daily US newspapers by print circulation in 2010, based on data from the *Alliance for Audited Media (AAM)*. The dataset provides the number of physical copies sold for each newspaper per county in 2010, covering all counties where circulation exceeded 25 copies.[3] Exceptions are the three newspapers classified as national by *AAM*, namely the *New York Times*, *USA Today*, and *Wall Street Journal*, whose circulation is available at the Designated Market Area (DMA) level.[4,5]

To supplement information on online readership, I use website traffic data from *Alexa Internet*, a former web analytics service that was discontinued in 2021. It offered a free browser toolbar that provided website statistics to its millions of users, while collecting anonymized data on website visits. That scale makes *Alexa Internet* a uniquely valuable source of high-frequency readership data, providing reliable figures even for smaller local news websites. The dataset covers the daily number of page views from 2010 to 2017 for

[3] Specifically, circulation data for each newspaper correspond to a four-quarter period ending between Q1 and Q4 of 2011, representing the earliest year with complete coverage across newspapers in the sample.

[4] The *Washington Post* is sometimes regarded as a national newspaper as well. In this paper, I stick to the definition by *AAM*.

[5] DMAs are 210 geographical regions in the United States, defined by the media analytics firm Nielsen.

the 80 largest of the 100 newspapers in my sample. Page views reflect the total number of daily visits per million toolbar users, including all clicks to each newspaper's main website and sub-domains.[6] Therefore, the variable should be interpreted as a proxy for total web traffic, which I verify using survey data. The data is limited to desktop users, thereby excluding traffic from mobile and tablet and does not provide information at the individual user level.

The newspapers in my sample represent a large share of total news consumption in the United States. They account for 76 percent of US print circulation recorded by *AAM* in 2010. According to *Alexa Internet*, their combined page views in 2010 were more than 1.5 times the combined page views of the websites of ABC, CBS, NBC, Fox, AP, and Reuters. Furthermore, none of the newspapers have exited the market as of the writing of this paper.

2.2 Paywalls

I obtain information on paywall launches through news coverage by the paywall-introducing newspapers themselves or competing outlets. Where ambiguous, I check historical website snapshots for the appearance of a digital subscription button using the Wayback Machine. Figure 1.1 shows the yearly share of newspapers in my sample that have a paywall on their website. Among the 100 newspapers, only five had introduced and upheld a paywall before 2010.[7] Between 2010 and 2016, 72 newspapers adopted paywalls, with an additional 14 doing so by 2021.

These paywalls were designed in a similar way: First, they usually adopted a "metered" model, allowing free access to the full website up to a certain number of clicks per month (usually 10-20), after which a subscription was necessary to view additional content.[8] Among four newspapers, the paywall had the form of new, dedicated "premium" website.[9] Second, paywalls were generally "leaky", meaning that they could be circumvented with sufficient effort (often intentionally), for example by clearing one's browser cache or finding links to specific articles on search engines or social media. Therefore, paywalls increased either the monetary or non-monetary costs of access. For the purpose of this paper, I focus only on the existence of a paywall, abstracting from design features

[6] *Alexa Internet* reported having several million toolbar users but never disclosed exact figures.

[7] Most notably, the *Wall Street Journal* introduced its paywall in 1997 and has maintained it continuously thereafter.

[8] The prominence of metered paywalls was partly due to Google's "First Click Free" policy, which until fall 2017 required news websites to grant each visitor at least three free articles per month, otherwise the website would be penalized in Google's search ranking.

[9] The four newspapers are the *Boston Globe*, *Houston Chronicle*, *Philadelphia Inquirer*, and *San Francisco Chronicle*.

like the number of free articles, leakiness, and subscription pricing.

Figure 1.1: Share of newspapers with paywalls by year

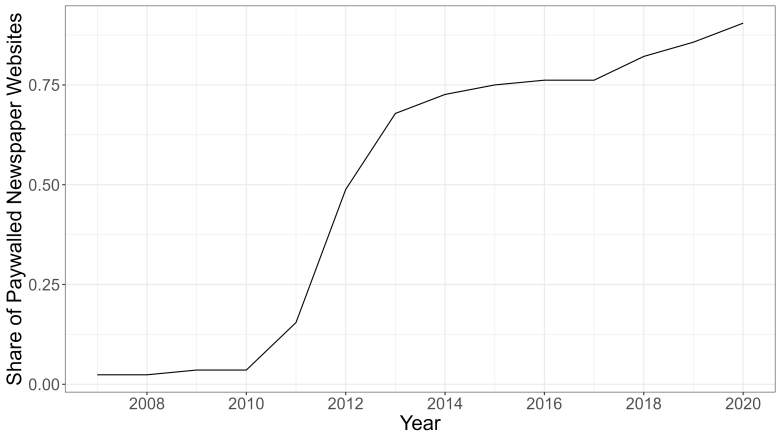

Notes: Share of US newspapers that have ever implemented a paywall on their website, among largest 80 newspapers by number of page views in 2010.

2.3 Survey outcomes

I measure knowledge and voting in the US population using the Cooperative Election Study (CES) for the years 2006–2021 (Ansolabehere and Schaffner, 2022). The *CES* is an annual, nationally representative survey of the US population administered by *YouGov* that uses repeated cross sections of 10,000 to 60,000 respondents.

Political knowledge

I measure political knowledge using the following two types of questions:

1. *Which party has the majority in [legislative body]?*
2. *Which party does [name of political representative] represent?*

These questions are asked separately for four legislative bodies (US Senate, US House, State Senate, State House) and four political representatives (governor, House representative, two senators), based on the respondent's place of residence.

To avoid biases from multiple hypothesis testing, I compile responses to these eight questions into a political knowledge index by calculating the share of correct answers for each respondent in a given year. To account for possible correlations in the signal of the answers, I use as an alternative measure the first principal component of the eight answers, ensuring that a higher value corresponds to a larger share of correct answers.

Figure 1.2 presents summary statistics. On average, respondents answer 63 percent of questions correctly, where correct answers are most frequent for the governor, and least frequent for their state's Senate and House majorities. Across all questions, respondents may choose "I don't know", which I classify as an incorrect answer. A third of the respondents answers all questions correctly, indicating that my measure discriminates well between low and high, but less so between high and very high knowledge. The average knowledge index is relatively stable over time, with a slight upward trend. Note that the two more difficult questions on state legislative majorities were omitted in 2006 and 2009, producing higher values in these years.

Electoral participation

The *CES* includes self-reported data on voting in presidential, gubernatorial, and federal and state legislative elections, covering both House and Senate. These are collected in the weeks after elections, but have been shown to suffer from misreporting (Ansolabehere and Hersh, 2012). Therefore, I also use validated data that *YouGov* matches to official voting records based on survey respondents' identifiable information. Indeed, validated turnout is consistently 10–15 percentage points lower than self-reported participation. For comparison, I present both self-reported and validated measures throughout the analysis.

News consumption

Data on news consumption is available only for selected years and measures whether individuals report consuming certain types of media in the 24 hours before answering the survey. In 2010, 56 percent reported reading a newspaper, of which 35 percent read exclusively online, 29 percent used both online and print sources, and 36 percent relied solely on print. Moreover, around 45 percent of respondents report that they follow the news "most of the time", the highest among four possible categories.

Paying for News

I use data from a nationally representative survey of US. adults conducted by the *Pew Research Center* and *Ipsos* (Pew Research Center, 2018), which includes 34,897 individuals surveyed between October 15 and November 8, 2018. It records whether individuals have paid for local news in the past year, along with detailed demographic information. According to the survey, approximately 14 percent of adults paid for local news sources in 2018. Although the data does not specify which local outlets respondents support and technically exclude national news subscriptions, they provide a rare and informative signal of the types of individuals generally willing to pay for newspaper content.

Figure 1.2: Descriptive statistics for political knowledge

(a) Distribution of correctly answered political knowledge questions per respondent

(b) Share of correct answers by question

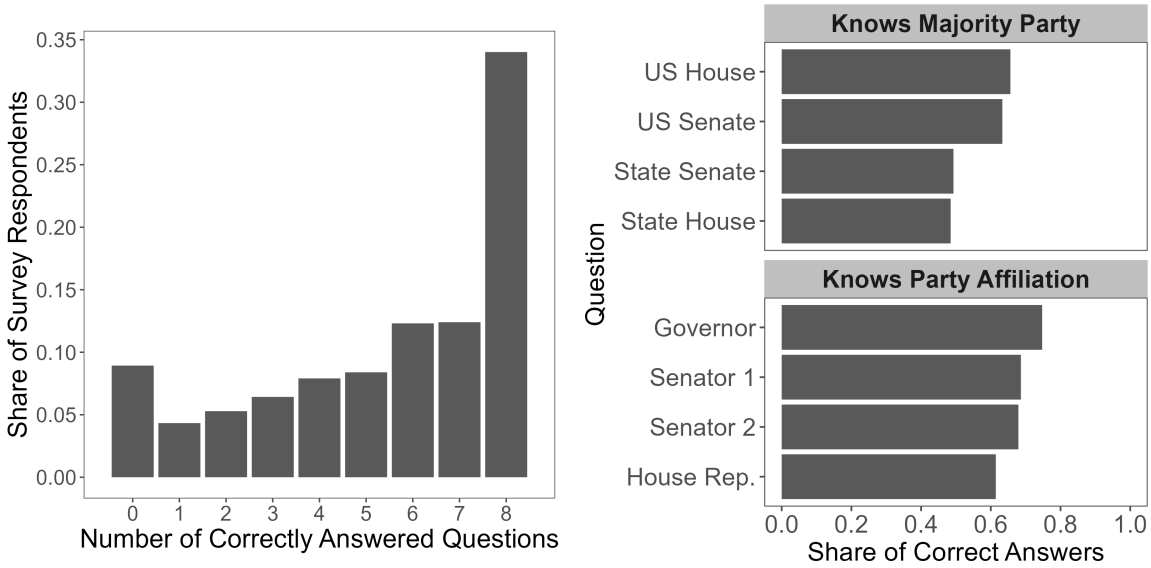

(c) Share of correct answers by year

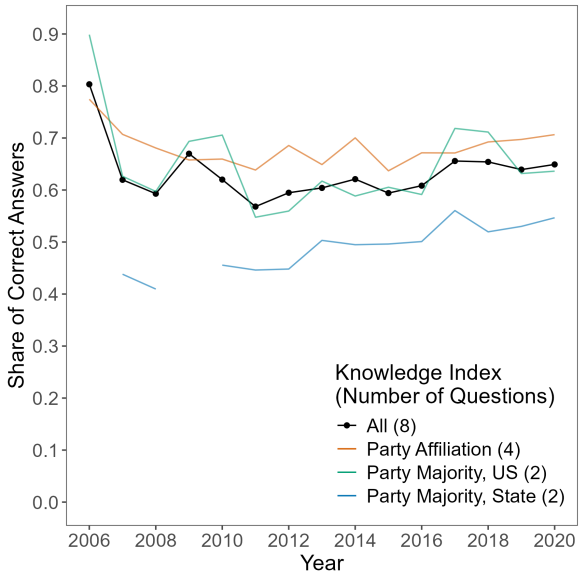

Notes: The response "I don't know" is coded as an incorrect answer. The questions on State legislative majorities were not asked in years 2006 and 2009. Shares are computed among all questions that were answered by an individual (around 1 percent of individuals skip at least one answer). Survey weights are included. Panel (a) shows the distribution of the number of correct answers over survey respondents. Where less than eight questions were asked, I round the share of correctly answered questions, multiply by eight, and round to the closest integer. Panel (b) shows the number of correct answers by question. Panel (c) shows the share of correct answers among all questions asked in that year (black), as well as separately for the 4 questions regarding representatives' party affiliation, 2 questions regarding majorities in the US chambers, and 2 questions regarding majorities in the State chambers.

Table 1.1: Summary statistics for selected variables derived from CES

	Mean	SD	Min	Max	N
Political knowledge					
Knows majority: US Senate	0.63	0.48	0	1	520,495
Knows majority: US House	0.66	0.47	0	1	523,703
Knows majority: State Senate	0.49	0.50	0	1	475,198
Knows majority: State House	0.48	0.50	0	1	474,936
Knows party: Governor	0.75	0.43	0	1	528,258
Knows party: Representative	0.61	0.49	0	1	524,663
Knows party: Senator 1	0.69	0.46	0	1	527,706
Knows party: Senator 2	0.68	0.47	0	1	527,682
Knowledge index: Share correct answers	0.63	0.35	0	1	530,966
Vote participation (self-reported)					
Voted in general election	0.85	0.36	0	1	343,614
Voted in presidential election	0.71	0.45	0	1	212,796
Voted in gubernatorial election	0.66	0.47	0	1	196,897
Voted in House election	0.65	0.48	0	1	420,133
Voted in Senate election	0.68	0.47	0	1	282,406
Voted in State House election	0.64	0.48	0	1	379,062
Voted in State Senate election	0.63	0.48	0	1	373,816
Vote participation (validated)					
Voted in general election	0.54	0.50	0	1	420,712
Voted in presidential election	0.56	0.50	0	1	212,796
Voted in gubernatorial election	0.52	0.50	0	1	168,069
Voted in House election	0.51	0.50	0	1	383,760
Voted in Senate election	0.53	0.50	0	1	256,062
Individual Characteristics					
Age	46.95	17.00	18	109	556,746
Female	0.52	0.50	0	1	531,087
White	0.72	0.45	0	1	556,746
Has child	0.27	0.44	0	1	509,736
College degree	0.37	0.48	0	1	556,746
Post-graduate degree	0.10	0.30	0	1	556,746
Employed	0.46	0.50	0	1	556,746
Unemployed	0.22	0.41	0	1	556,746
Retired	0.19	0.39	0	1	546,796
Family Income $>$ 60k	0.35	0.48	0	1	556,746
Family Income $>$ 100k	0.15	0.36	0	1	556,746
Democrat	0.38	0.49	0	1	506,342
Republican	0.30	0.46	0	1	506,342
Media consumption					
Read newspaper in past 24h	0.48	0.50	0	1	384,282
Read newspaper online in past 24h	0.32	0.47	0	1	384,282
Follow news "most of the time"	0.45	0.50	0	1	556,746

Notes: Selected summary statistics of survey respondent–level variables derived from CES for years 2006–2021, averaged across all years and including survey weights.

2.4 Other data sources

I incorporate data on county demographic and economic characteristics for 2010 from the 5-Year *American Community Survey* (U.S. Census Bureau, 2010). Moreover, I use county-level estimates for visits to selected newspaper websites from the *MRI-Simmons Local* dataset, which assigns responses to its *National Consumer Survey* (MRI-Simmons, 2014) to geographies based on respondent and county characteristics. I measure county partisanship using the partisanship index in the *National Neighborhood Data Archive (NaNDA)* (Chenoweth et al., 2020), which is based on Democratic and Republican vote shares in past elections. To link counties with media markets, I use the crosswalk from Gentzkow and Shapiro (2008).

3 Effect on news consumption

In this section, I analyze how the launch of paywalls on major US newspaper websites between 2010 and 2017 affected page views. I employ a staggered difference-in-differences design that exploits the sequential timing of paywall introductions across newspapers.

3.1 Empirical strategy

My analysis is based on the following two-way fixed effects (TWFE) regression model:

$$\log(Pageviews_{n,t}) = \sum_{\tau \neq -1} \beta_\tau D^\tau_{n,t} + \alpha_n + \gamma_t + \delta' \mathbf{X_{n,t}} + \varepsilon_{n,t} \tag{1.1}$$

$Pageviews_{n,t}$ denotes the number of page views for newspaper n in year-month t. I apply a log transformation to account for skewness, enabling interpretation of the coefficients as (approximate) percentage effects on page views. $D^\tau_{n,t}$ is an indicator equal to one if and only if newspaper n introduces a paywall τ months after t. α_n and γ_t denote newspaper and year-month fixed effects, respectively. I cluster standard errors at the newspaper level.

To account for differential trends in newspapers' page views as well as shocks to their respective audiences over time, I include in $\mathbf{X_{n,t}}$ characteristics of the markets in which the newspapers operated as of 2010, interacted with year-month fixed effects. Specifically, using the counties in which the newspaper is circulated, I compute the average population density, the shares of population with yearly income below $\$50k$ and above $\$100k$, the share of college-educated individuals, and the *NaNDA* partisanship index, weighted by the number of print copies sold in each county.

For this analysis, I restrict my sample as follows: First, I exclude the four newspapers that introduced new premium websites as part of their paywalls. Second, I exclude the newspaper *Newsday*, which introduced a paywall in October 2009, just before the sample period. Third, I exclude three small newspapers with significant data gaps, likely due to data collection errors by *Alexa Internet*. The final sample comprises the monthly page views for 72 newspapers from 2010 to 2017. Among these newspapers, two introduced paywalls before the sample period, 55 during the sample period, eleven afterward, and four never introduced a paywall (as of 2022). In my preferred specification, I assign the two always-treated and eleven later-treated newspapers as untreated, in addition to the four never-treated. These newspapers arguably act as valid control units because they implemented their paywalls well outside the sample period. I demonstrate that my results are robust to varying definitions of the control group.

3.2 Identification

Identification of the monthly treatment effects β_τ relies on the assumption that page views of paywalled newspaper websites would have evolved parallel to those of non-paywalled websites in the absence of paywalls. In this sub-section, I discuss several challenges to this assumption.

Heterogeneous effects

I account for potential biases from heterogeneous treatment effects across units or time, which may arise when estimating Equation (1.1) via Ordinary Least Squares (for example, see Baker et al., 2022). I use the robust estimator proposed by Callaway and Sant'Anna (2021), which aggregates doubly-robust estimates for treatment group-time effects using only "clean" control units (Sant'Anna and Zhao, 2020).[10] Control variables are included in each of those group-time-specific estimates both as a regression adjustment and for the construction of probability weights. Additionally, I employ the robust estimator by Sun and Abraham (2021), which estimates treatment effects by including only treated and never-treated units in the outcome regression and re-weighting the estimates accordingly.

Selection

Since the decision to introduce a paywall is endogenous, paywalls may correlate with unobserved factors not accounted for in the regression. To mitigate concerns about selec-

[10]Throughout the analysis, I use a common ("universal") base period, consistent with Equation (1.1), estimating newspaper-specific treatment effects relative to the month preceding paywall introduction.

tion, I do the following. First, I show that pre-treatment trends in page views are parallel, ensuring that such trends are not biasing the result. Second, I control for market-based factors interacted with month-year fixed effects, which accounts both for different incentives for newspapers to introduce a paywall, and for time-varying shocks to different types of newspaper audiences. Third, to address concerns about selection in the control group, I show that results are statistically and economically significant for different compositions of the control group. Fourth, I explore heterogeneity in treatment effects to ensure that the results are not just driven by a small subset of newspapers.

Spillovers

Another potential concern involves biases from spillover effects: Readers displaced by a paywall may switch to a newspaper in the control group, which leads to an upward bias (in absolute terms) in the estimate for the paywall effect. Formally, such spillovers violate the stable unit treatment value assumption (SUTVA), which requires that units' potential outcomes do not depend on the treatment status of other units. However, in this setting, spillovers are unlikely to be a first-order concern due to the segmented nature of the newspaper market: All but three newspapers are classified as regional by *AAM*, catering primarily to their respective regional markets. As a result, these newspapers likely act as poor substitutes for one another. For example, it seems unlikely that a reader of the *Chicago Tribune* would respond to a paywall by switching to the *Los Angeles Times*.

To address any remaining concerns, I verify that my results hold when removing the largest paywalled newspapers and the smallest control newspapers. Intuitively, paywalls on larger newspapers displace more readers, whose substitution has a larger relative effect on the control newspapers the fewer readers these newspapers have.

3.3 Results

Main result

Figure 1.3 presents estimates for the dynamic effect of paywalls on website page views, using the robust Callaway and Sant'Anna (2021) estimator. On average, this "first wave" of paywalls led to a sharp and persistent reduction in the number of page views for paywalled newspapers. Table 1.2 reports estimates for the static effect. On average, page views decreased by 25–30 percent, depending on the choice of control variables. Figure A1.2 and Table A1.1 demonstrate that these results are robust to using alternative estimators.

Figure 1.3: Effect of paywalls on website page views

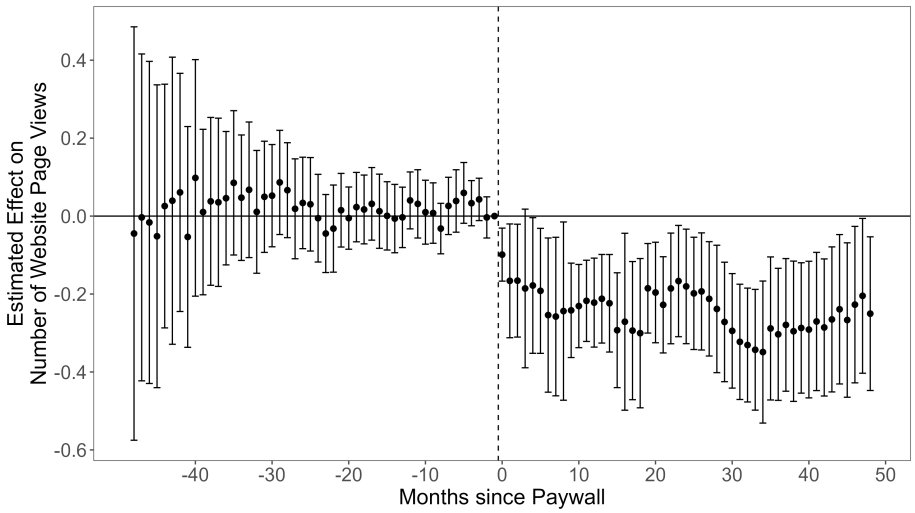

Notes: Monthly coefficients for β_τ from Equation (1.1): Regression of monthly number of newspaper website views on indicators denoting the number of months since paywall implementation. The omitted category is the month before the paywall. Includes newspaper fixed effects, year-month fixed effects, and log population density, shares of three income buckets, college-educated share, and partisanship index for newspapers' audiences, interacted with year-month fixed effects. Estimated using Callaway-Sant'Anna estimator. Vertical bars denote 95 percent pointwise confidence intervals. Standard errors are clustered by newspaper.

Table 1.2: Effect of paywalls on pageviews

Dependent Variable:	Log(Pageviews)				
	(1)	(2)	(3)	(4)	(5)
Paywall	-0.257***	-0.291***	-0.302***	-0.280***	-0.310***
	(0.070)	(0.062)	(0.064)	(0.063)	(0.076)
Newspaper FE	✓	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓	✓
Controls x Month-Year FE:					
Log(Population Density)		✓	✓	✓	✓
Sh. HH income $\leq 50k$			✓	✓	✓
Sh. HH income $> 100k$			✓	✓	✓
Partisanship index				✓	✓
Sh. college-educated					✓
Observations	6,912	6,912	6,912	6,912	6,912

Notes: Coefficients for static (pre-post) version of Equation (1.1): Regressions of monthly number of newspaper website views on indicator denoting active paywall. All specifications include newspaper fixed effects and year-month fixed effects. Controls represent average county characteristics at 2010-levels where newspaper is active, weighted by number of print copies sold, interacted with year-month fixed effects: Log population density, share of households with yearly income below $50,000 and above $100,000, share adults with college degree, and NaNDA partisanship index. Estimated using Callaway-Sant'Anna estimator. Standard errors in parentheses are clustered by newspaper. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Interpretation

The dynamic estimates reject the presence of pre-treatment trends, indicating that newspapers did not select into introducing paywalls based on prior trends in page views. The immediate decline after introducing a paywall supports the interpretation that paywalls raise the opportunity cost of website access, prompting some users to reduce their visits. Conversely, paywalls may also encourage new subscribers to visit more frequently. Thus, the estimates represent the net outcome of these two opposing effects. However, without individual-level user data, it is not possible to determine whether the observed decline reflects the displacement of a few highly active readers or many rather casual users.

The timing of the effects is inconsistent with alternative explanations that rely on gradual changes accompanying paywalls, such as shifts in topic composition or ideological slant, as these would likely materialize more slowly. Nonetheless, I cannot entirely rule out that these factors contributed to the modest downward trend observed in the treatment effects. However, other long-term factors may also explain this trend. For instance, the diffusion of broadband and 3G internet during the sample period likely increased overall interest in online compared to offline news. If new users entering the online news market disproportionately favored non-paywalled over paywalled websites, the initial effect of paywalls would become more pronounced over time.

How many individuals respond to paywalls by switching to the print version of the paywalled newspaper? While I cannot provide a definite answer based on the available data, this channel is unlikely to explain the observed effects. During the sample period, digital news consumption grew rapidly while print news consumption declined: The growth in the number of newspaper website visits often exceeded 10 percent *per year*, while the number of print copies sold declined by around 30 percent between 2009 and 2017 (Pew Research Center, 2023). Moreover, studying the *New York Times* paywall, Pattabhiramaiah et al. (2019) conclude that spillovers from digital to print were negligible, with a 17 percent decline in website visits accompanied by only a 1-4 percent increase in print circulation. This increase was primarily driven by bundling online subscriptions with the print version, rather than by substitution to standalone print subscriptions.

Further robustness

In this sub-section, I confirm that my results are robust to changes in the choice of which newspapers constitute the control group, and are not driven by users' substitution between paywalled and non-paywalled newspapers.

Table 1.3 shows results under different control group compositions. Column (2) shows that excluding the pre-treatment period observations for the 55 not-yet-treated newspapers from the control group has minimal impact on the results. Columns (3) and (4) adopt the most restrictive definition for the control group, comparing newspapers that implemented a paywall during the sample period to those that introduced one just afterwards. If anything, the effects are larger under these specifications, with a 34–36 percent decline in page views compared to the baseline estimate of a 28 percent decrease.

Table A1.2 shows results when excluding different sets of newspapers from the sample. Columns (1) and (2) indicate that excluding large newspapers, in particular the three national newspapers and the fifteen largest newspapers by page views, increases the estimate slightly. Columns (3) and (4) exclude the fifteen largest paywalled newspapers and the five smallest non-paywalled newspapers. If spillovers were significant, one would expect the absolute magnitude of the estimates to decrease under these restrictions. However the opposite is true: the estimates are either unchanged or increase slightly. Intuitively, while theoretically an issue, spillovers may not play a significant role empirically because of the regional segmentation of the news markets, making most newspapers poor substitutes for one another. Moreover, the control group is sufficiently large to dilute potential biases.

Table 1.3: Effect of paywalls on pageviews - Control groups

Dependent Variable:	Log(Pageviews)			
Control group:	All	w/o Treated	w/o Earlier/Later	Only Later
	(1)	(2)	(3)	(4)
Paywall	-0.280***	-0.277***	-0.342***	-0.355***
	(0.067)	(0.070)	(0.068)	(0.068)
Newspaper FE	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓
Market Controls	✓	✓	✓	✓
Control group:				
Earlier Treated	✓	✓		
Treated (pre-treatment)	✓		✓	
Later Treated	✓	✓	✓	✓
Never Treated	✓	✓		
Observations	6,912	6,912	6,336	6,336

Notes: Coefficients for static (pre-post) version of Equation (1.1), for different compositions of the control group: "Earlier treated" are 2 newspapers that implemented paywall before 2010. "Treated (pre-treatment)" are 55 newspapers that implemented paywall between 2010 and 2017. "Later Treated" are 11 newspapers that implemented paywall after 2017. "Never Treated" are 4 newspapers that have not implemented a paywall by 2022. Regressions of monthly number of newspaper website views on indicator denoting active paywall. All specifications include newspaper fixed effects and year-month fixed effects. Controls represent average county characteristics at 2010-levels where newspaper is active, weighted by number of print copies sold, interacted with year-month fixed effects: Log population density, share of households with yearly income below $50,000 and above $100,000, share adults with college degree, and NaNDA partisanship index. Estimated using Callaway-Sant'Anna estimator. Standard errors in parentheses are clustered by newspaper. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

4 Effect on political knowledge and electoral participation

In this section I explore how paywalls affected factual knowledge about contemporary US politics as well as participation in elections by leveraging temporal and regional variation in exposure to paywalls.

4.1 Measuring regional exposure to paywalls

A key challenge is the lack of granular data that directly links web traffic from specific counties to individual websites.[11] As a practical alternative, I proxy regional online readership using newspapers' regional print circulation, which provides high-quality information on county level. For regional newspapers, which constitute the majority of my sample, print circulation reveals in which localities the regional focus of the newspaper is relevant. For the *New York Times*, instead, print circulation rather captures where its predominantly left-leaning and highly educated audience resides.

To validate this approach, I present two pieces of evidence. First, Figure 1.4 presents a binned scatter plot for the shares of respondents who report reading a newspaper offline versus online in the past 24 hours, averaged by county and year. It reveals that consumption of newspapers offline is strongly positively correlated with consumption online, suggesting that print circulation successfully captures regional variation in the overall intensity of engagement with newspaper content, regardless of format. Second, to verify that this relationship also holds for individual newspapers, I use *MRI-Simmons* data on estimated website visits by county for the *Washington Post* and *Los Angeles Times* – the only two non-national newspapers for which this data is available. Figures 1.5a and 1.5b show scatter plots for the number of print copies sold per household in a given county, versus the (estimated) share of the population visiting the respective website in the past 30 days. In both cases, higher print circulation is positively and approximately linearly associated with more frequent website visits for these websites. Together, these results suggest that newspapers' print circulation captures relevant variation in the regional engagement with their respective websites.

[11] For example, at the beginning of my sample, *comScore*'s browser panel of 50,000 individuals records fewer than 100 unique weekly visitors to most newspaper websites outside the top three.

Figure 1.4: County-level correlation of newspaper readership online vs. offline

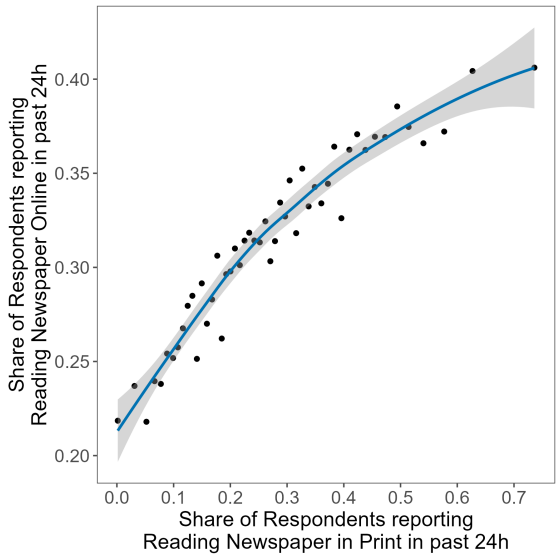

Notes: Binned scatterplot using self-reported newspaper readership data from the CES 2006–2021. The x- and y-axes show the share of individuals who report reading a newspaper in print and online, respectively, in the past 24 hours. Observations are aggregated at the county-year level prior to binning. Sample weights are included. The blue line displays the LOESS-smoothed conditional mean with 95 percent confidence intervals.

Figure 1.5: County-level correlation of print circulation and website visits for selected newspapers

(a) Washington Post

(b) Los Angeles Times

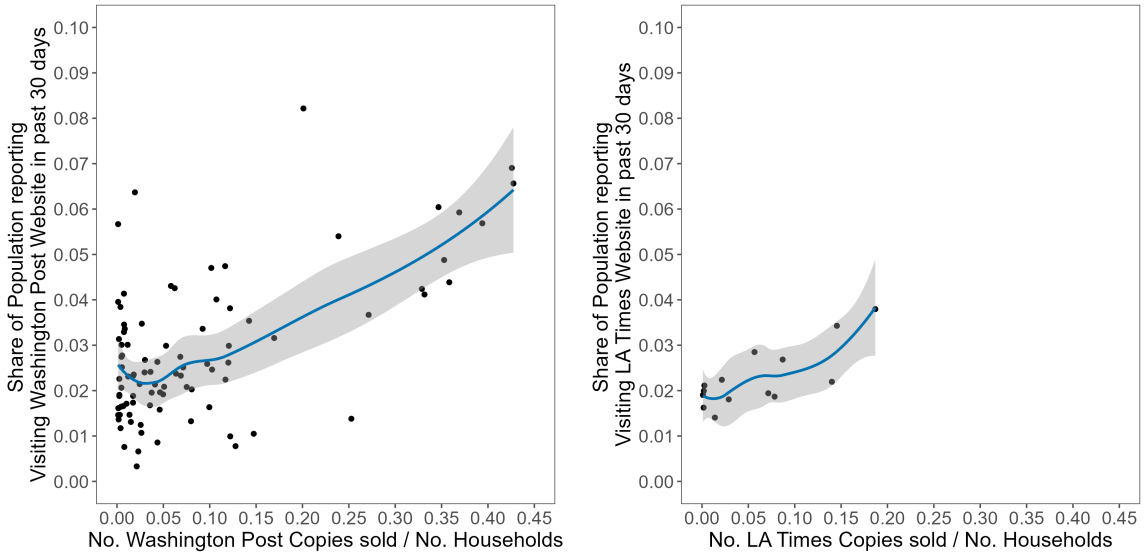

Notes: Each dot represents a county. Panel (a) shows data for the Washington Post; Panel (b) for the Los Angeles Times. The y-axis reports the percent of respondents who visited the newspaper's website in the past 30 days in 2013, based on county-level MRI-Simmons estimates constructed via multilevel regression and post-stratification (MRP). The x-axis shows the number of print copies sold per household in 2011, including counties with print circulation greater than 100 and excluding outliers at the top and bottom of both axes. The blue line displays the LOESS-smoothed conditional mean with 95 percent confidence intervals.

Motivated by this evidence, I define paywall exposure for each county and year as the share of households subscribed to newspapers that have a paywall in place. Formally:

$$Paywall\ Exposure\ _{c,t} = \sum_{n \in Newspapers} \frac{Circulation\ ^{n}_{c,\,2010} \times I(Paywall\ ^{n}_{t})}{Households\ _{c,\,2010}} \tag{1.2}$$

where $Circulation\ ^{n}_{c,\,2010}$ is the number of print copies sold of newspaper n in county c in 2010. $I(Paywall\ ^{n}_{t})$ is a binary variable that indicates whether newspaper n has implemented a paywall by year t. I fix circulation at their 2010 levels, prior to the widespread introduction of paywalls, to avoid endogeneity arising from paywalls themselves potentially affecting circulation. This measure of paywall exposure resembles shift-share variables (see Bartik, 1991; Blanchard et al., 1992), where initial market shares serve as shares and paywall introductions act as shifts.

To provide insights into the variation captured by paywall exposure, Figure 1.6 presents key descriptive statistics. Panel (a) shows that paywall exposure increased significantly between 2010 and 2014, and plateauing afterward. Panel (b) highlights substantial variation across counties, with only few experiencing the highest levels of exposure. To account for differences in county populations, Panel (c) displays the distribution of paywall exposure among respondents in the *CES* survey. Accounting for county sizes reveals that several large counties experience intermediate to high paywall exposure, producing a less skewed distribution. Note that since paywall exposure is intended as a proxy variable, its relative differences across counties and years are more informative than exact numerical values.[12]

Figure 1.7 shows geographic variation, revealing that paywalls affected all major regions of the US, while concentrating around urban areas. To gain a deeper understanding of which regions are most likely to be affected by paywalls, Figure A1.4 presents coefficients from a regression of paywall exposure in 2020 on (standardized) county characteristics in 2010. Paywalls disproportionately impacted high-income, Democratic-leaning counties. However, paywalls do not appear to be correlated with political knowledge, conditional on controls.

[12] In particular, paywall exposure should not be interpreted as the number of households affected by paywalls, for the following reasons: First, a household may subscribe to multiple newspapers. Second, each household member may have access to multiple distinct subscriptions for the same newspaper (e.g., through their employer or university). Third, because print and digital newspaper access are often bundled, subscribing households are typically *not* restricted by paywalls.

Figure 1.6: Paywall Exposure

(a) By year

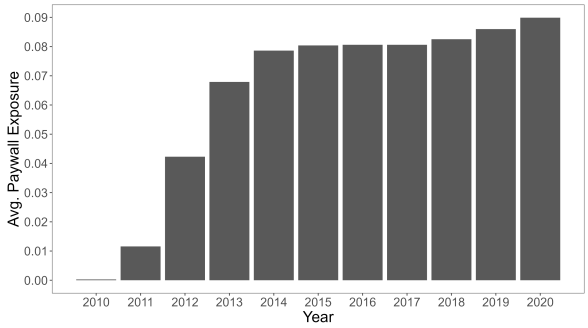

(b) Distribution over counties

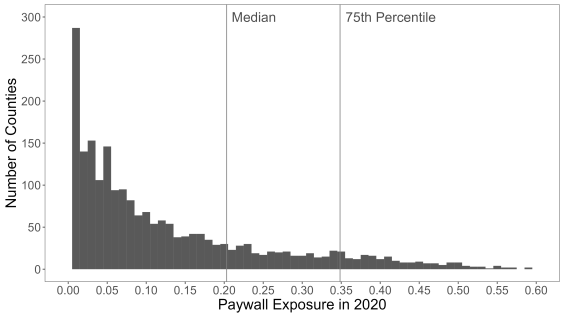

(c) Distribution over survey respondents

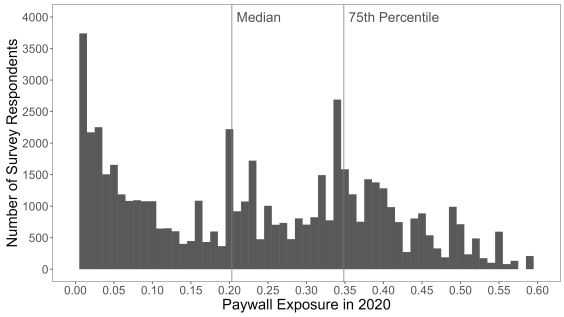

Notes: Panels (a) and (b): Average paywall exposure across counties by year. Panels (c) and (d): Distribution of paywall exposure over counties and survey respondents (weighted), respectively, in 2020.

Figure 1.7: Geographical variation of paywall exposure

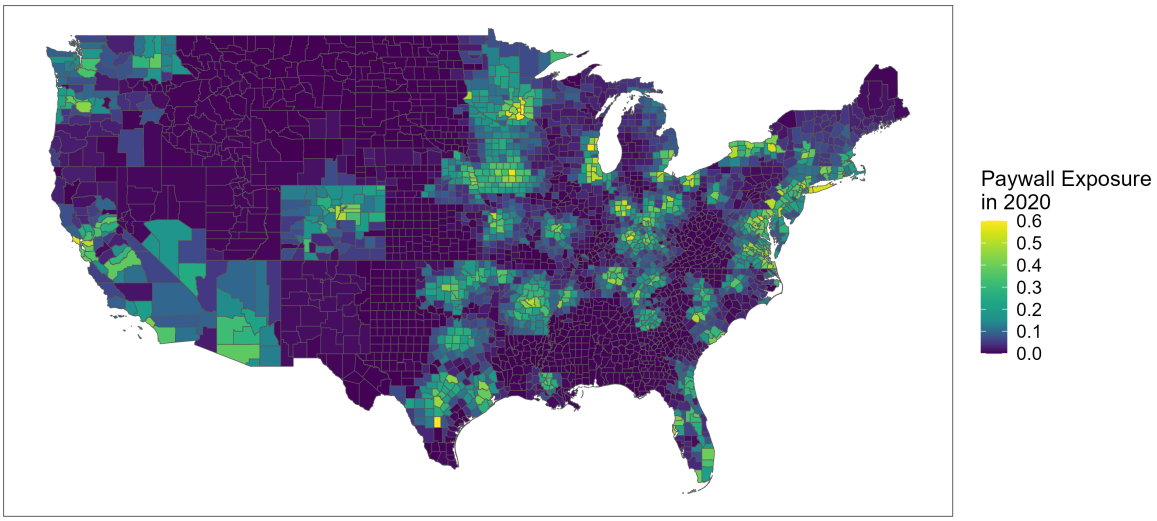

Notes: Paywall exposure by county in 2020.

4.2 Empirical strategy

I exploit the staggered exposure of counties to paywalls using a difference-in-differences design. For this purpose, I divide counties into a low- and a high-paywall exposure group, with counties entering treatment once paywall exposure exceeds a fixed threshold (Guriev et al., 2021). The baseline specification is based on the 75$^{\text{th}}$ percentile of paywall exposure at the end of the sample period, which ensures high power by splitting observations roughly equally into treated and untreated observations, across all years. In robustness checks, I confirm that the results are robust to varying this threshold.

The causal analysis is motivated by the following specification:

$$y_{i,c,t} = \sum_{\tau \neq -1} \beta_\tau D^\tau_{c,t} + \delta' \mathbf{X}_{i,c,t} + CountyFE_c + YearFE_t + \varepsilon_{i,c,t} \tag{1.3}$$

The dependent variable $y_{i,c,t}$ is an outcome variable for survey respondent i residing in county c in year t, such as the share of correctly answered political knowledge questions. County fixed effects account for pre-existing regional differences in knowledge, for example due to education or political composition, and year fixed effects capture yearly shocks to knowledge. Alternative specifications include state-by-year fixed effects instead, which capture changes in knowledge due to elections.[13]

$\mathbf{X}_{i,c,t}$ denotes time-varying control variables. In the baseline specification, these include individual demographic variables in the form of categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2). In robustness checks, I interact these with year fixed effects. I cluster standard errors at the county level.

The treatment variables $D^\tau_{c,t}$ are indicators equal to one if and only if paywall exposure in county c first exceeds the sample-weighed 75$^{\text{th}}$ percentile τ years after t. Splitting at the 75$^{\text{th}}$ percentile (as opposed to the median, for example), ensures a suitable number of treated and control observations across the panel. In particular, roughly equal parts of observations are never-treated, treated early (by 2017), and treated late (by 2021). Therefore, this approach ensures high power.

I estimate Equation (1.3) using the robust estimator proposed by Callaway and Sant'Anna (2021).[14] In each year, the control group consists of counties that are either never treated

[13] Note that the specification does not include individual fixed effects because the survey is a repeated cross section in which each individual appears only once.

[14] Throughout the analysis, I use a common ("universal") base period, consistent with Equation (1.3), estimating county-specific treatment effects relative to the year preceding treatment.

or not yet treated. While county fixed effects are implemented natively, this is not the case for state-by-year fixed effects and individual-level controls.[15] Therefore, I proceed in two steps. First, I residualize the outcome variable on the individual controls, county fixed effects, and state-by-year fixed effects. Second, I apply the Callaway and Sant'Anna (2021) estimator to these residualized outcomes. I verify empirically that the inclusion of control variables in this way improves precision of the estimator without considerably changing the point estimates.

4.3 Identification

I identify the average effect of paywall exposure on political knowledge across the adult US population, comparing regions with high versus low exposure. Identification requires that, in the absence of paywalls, trends in political knowledge would have evolved similarly in high- and low-exposure regions. In the following, I discuss four key challenges to identification and strategies to address them: Omitted county characteristics, correlations with regional politics, heterogeneous effects, and selection of newspapers.

Omitted county characteristics

A primary concern is that paywall exposure may be correlated with county-level characteristics that independently affect political knowledge. For example, in Section 4.1 I show that paywall exposure is higher in counties with higher initial income. If such counties experienced shocks to political knowledge that are unrelated to paywalls, estimates may be biased. For instance, high-income individuals may have a higher propensity to respond to the rise of social media by substituting away from newspapers toward softer news and entertainment content, even in the absence of paywalls. If not accounted for, this could lead to a downward bias in the estimated effect of paywalls.

More subtly, this bias may also arise from non-random measurement error in paywall exposure. If high-income counties have higher subscription shares for a given level of online engagement, then paywall exposure is overstated in these areas. Conversely, if paywalls are on average more restrictive in high-income counties, actual exposure may be understated. In both cases, failing to account for potential shocks to high-income individuals may bias estimates.

[15] Although the estimator allows for the inclusion of control variables, their implementation is considerably more demanding than specified in Equation (1.3): Controls enter the estimation of each group-time average treatment effect, effectively interacting them with treatment-group-by-year fixed effects. This approach requires sufficient support in each cell, and estimation often fails when group sizes are small.

I address these concerns in three main ways. First, I verify that results are robust to including interactions of individual-level controls with year fixed effects. These coefficients flexibly absorb year-specific shocks to specific subpopulations, such as changes in media consumption habits. I also check robustness to including state-by-year fixed effects, which account for unobserved state-level shocks. Second, the staggered implementation of paywalls helps distinguish their effects from contemporaneous trends. Even if some regions exhibit changes in political knowledge due to unobserved factors, it is unlikely that the timing of these trends would systematically coincide with the timing of paywall adoption across counties. Third, I show that results are robust to alternative threshold definitions for discretizing counties into high- and low-exposure groups. This reduces sensitivity to measurement error near the cutoffs and ensures that results are not driven by misclassification of a small number of units.

Correlation with regional politics

The identifying assumption may also be violated if regions with high paywall exposure experienced divergent trends in political knowledge for reasons unrelated to paywalls. For instance, if states with more paywalls happen to elect many newcomers shortly before paywall adoption, residents might find political survey questions easier or harder to answer, leading to lower measured knowledge.

To address this concern, I estimate two sets of specifications, one with only year fixed effects, and one that instead includes state-by-year fixed effects. The former leverages variation in paywall exposure across all counties and years, and should therefore expected to have more power. The latter relies on comparisons of counties within state and year, effectively holding constant statewide political shifts such as election outcomes, political events, or policy changes.

Heterogeneous Effects

Survey respondents may respond differently to paywalls depending on individual characteristics. For example, initial news consumption patterns may play a role. Holding paywall exposure fixed, high-income counties may receive a smaller share of news from television, making them more affected by paywalls. In this case, one would expect larger effects in high-income areas. Alternatively, high-income individuals may be more likely to purchase subscriptions in response to paywalls, while low-income individuals may switch to free news sources with less news content, implying smaller effects in high-income areas. Moreover, the effect of paywalls may also depend on their business and editorial decisions as well as the presence of other local newspapers, and may unfold or vary over time.

Such heterogeneous effects do not pose a threat to identification but are important for interpretation. The Callaway-Sant'Anna estimator accounts for heterogeneity by ensuring that comparisons are made between appropriate units and that group-time specific treatment effects are aggregated correctly. Therefore, it explicitly allows that different groups experience effects with varying intensity and and different patterns over time. However, when interpreting the results, it is important to keep in mind that the reported average effect may mask substantial heterogeneity across individuals and counties.

Selection of Newspapers

Finally, one may be concerned that omitting newspapers outside the largest 100 in the analysis could bias the constructed measure of paywall exposure. In fact, this course of action effectively ignores any paywalls of these smaller newspapers, leading to underestimation of paywall exposure in these regions.

To address this concern, I verify that the results are robust to excluding counties where omitted newspapers account for a large share of total circulation. I also check robustness to dropping small counties, as these are more likely to rely on smaller, local newspapers outside the top 100.[16]

4.4 Effects on political knowledge

In this subsection, I present results for the effects of paywall exposure on political knowledge.

Motivating evidence

I start by providing motivating evidence on the negative correlation between paywall exposure and political knowledge. Figure 1.8 plots 25 quantiles of paywall exposure against the share of correctly answered political knowledge questions (residualized on county, year, and individual controls) using data from 2011 to 2021 (all years affected by paywalls). While the relationship is somewhat noisy, respondents exposed to higher levels of paywalls clearly answer fewer knowledge questions correctly.

Table A1.3 confirms this pattern using regression analysis based on Equation (1.3), replacing treatment indicators with the continuous Paywall Exposure variable. Columns (1)–(3) include county and year fixed effects, while columns (4)–(6) add state-by-year fixed

[16]This check also addresses the fact that *AAM* data are only available for counties with at least 25 sold copies.

Figure 1.8: Correlation of paywall exposure and political knowledge

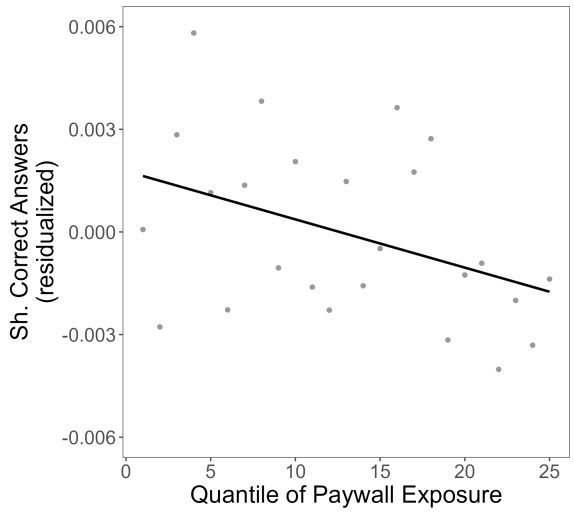

Notes: Dots represent the average share of correct answers to eight political knowledge questions, residualized on individual characteristics as well as county- and year fixed effects, within each of 25 quantiles of paywall exposure. Sample weights are included. The black line shows the corresponding linear trend.

effects. Comparing Columns (1)–(2) and (4)–(5) reveals that the inclusion of individual demographic controls reduces the correlation between paywall exposure and knowledge while improving explanatory power. This result reflects both the correlation of paywall exposure with individual characteristics, and confirms that different types of individuals differ in their level of informedness about politics. However, comparing Columns (2)–(3) and (5)–(6) shows that interacting demographics with year fixed effects yields little additional explanatory power. This supports the central assumption that the negative correlation of paywalls and political knowledge is not systematically driven by unrelated shocks to demographic subgroups.

The results persist when relying only on within state-year variation, though slightly attenuated and only significant at the 10 percent level. This pattern may reflect that this specification is more restrictive, as paywall exposure varies more across than within states. Alternatively, higher paywall exposure may simply coincide with greater political change across states, making it more difficult to correctly guess majorities or party affiliations.

Table A1.4 replicates the previous analysis using bins corresponding to terciles of paywall exposure. Although estimates are mostly insignificant at conventional levels due to limited statistical power, the coefficient for the second tercile is consistently negative across specifications, and the coefficient for the third tercile is more negative and marginally significant in some cases. This pattern suggests that the results are not driven by outliers, and instead supports a monotonic negative relationship between paywall exposure and political knowledge.

Main result

I now present the main results, motivated by the causal specification in Equation (1.3) and estimated via Callaway-Sant'Anna. Figure 1.9 displays the dynamic effect of paywalls on the share of correctly answered political knowledge questions, comparing survey respondents in counties with high versus low paywall exposure. Panel (a) includes county and year fixed effects, while Panel (b) adds state-by-year fixed effects. Both specifications confirm that the negative effect of paywall exposure emerges in the year of treatment and remains persistent after one to two years. Importantly, the estimates show no evidence of pre-treatment trends in political knowledge. Moreover, the similar-sized effect obtained using only within state-year variation helps rule out confounding from cross-state differences in the political environment.

Table 1.4 demonstrates that the result is robust to the choice of control variables and fixed effects. Column (1) reports results from a specification with county and year fixed effects and no individual controls, which corresponds to the most basic application of the Callaway-Sant'Anna estimator. Columns (2) and (3) first residualize the outcome on demographic controls, in levels and interacted with year fixed effects, respectively. These adjustments increase precision considerably, without affecting the magnitude of the estimate. Columns (4) and (5) repeat the estimation using state-by-year fixed effects, and yield similar conclusions. Together, the results suggest that, conditional on fixed effects and demographic controls, the bias from omitted variables is negligible. In particular, neither unrelated shocks to specific subpopulations nor potential measurement error in paywall exposure appear to drive the results.

The most conservative estimate implies that high paywall exposure reduces the probability of correctly answering a political knowledge question by 1.2 percentage points (1.6 percent) relative to low exposure. For reference, this estimate corresponds to roughly 1 additional incorrect answer per 8 respondents, or 1 out of every 46 respondents now answering all questions incorrectly. Although modest in absolute terms, the effect is sizable given that only about half the sample regularly consumes news, and only a subset is expected to reduce consumption in response to paywalls. I return to the discussion of effect magnitudes at the end of this Section.

Figure 1.9: Effect of paywalls on political knowledge

(a) County and year fixed effects

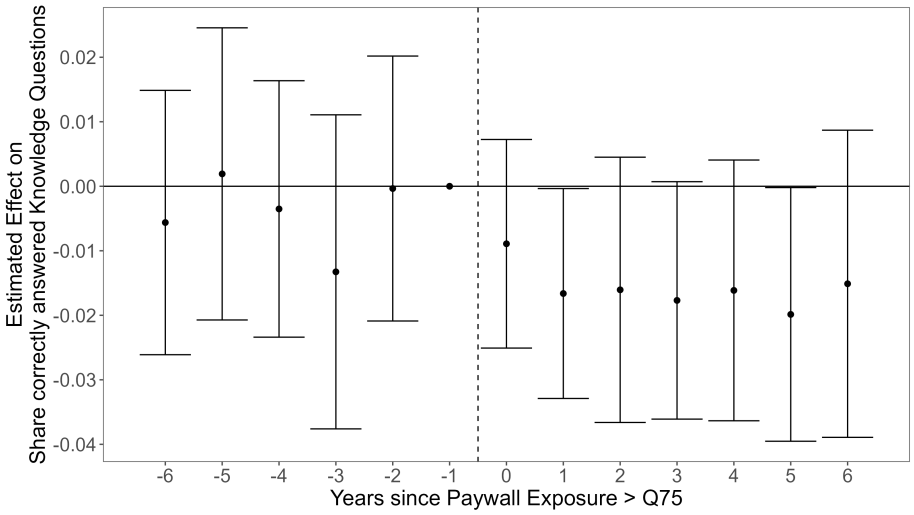

(b) County and state-by-year fixed effects

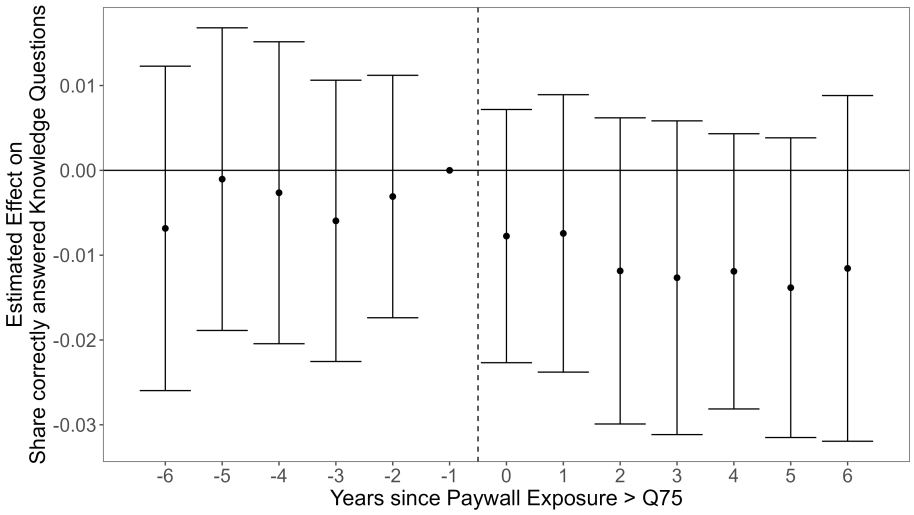

Notes: Yearly estimates for the effect of high exposure to paywalls on political knowledge. The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021), residualized on individual demographic controls, county fixed effects, and either year fixed effects (Panel (a)) or state-by-year fixed effects (Panel (b)). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Vertical bars show bootstrapped 95 percent uniform confidence intervals. Sample weights are included. Standard errors are clustered by county.

Table 1.4: Effect of paywalls on political knowledge

Dependent Variable:	Share Correct Answers				
	(1)	(2)	(3)	(4)	(5)
Paywall Exposure > Q75	-0.018*	-0.017***	-0.018**	-0.012**	-0.015***
	(0.008)	(0.005)	(0.006)	(0.004)	(0.004)
Dep. Var. Mean	0.745	0.745	0.745	0.745	0.745
County FE	✓	✓	✓	✓	✓
Year FE	✓	✓		✓	
State-Year FE			✓		✓
Individual Controls		✓		✓	
Individual Controls x Year			✓		✓
Observations	530,372	518,867	518,867	518,867	518,867

Notes: Estimates for the effect of high exposure to paywalls on political knowledge. The specification is motivated by Equation (1.3), estimated and aggregated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021). All estimations include county and year fixed effects. Where indicated, the outcome is additionally residualized on individual demographic controls – categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2) – as well as county fixed effects and either year or state-by-year fixed effects. Counties are treated once paywall exposure exceeds the 75$^{\text{th}}$ percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Further robustness

The findings are robust to varying the threshold that defines whether counties are classified as as high- or low-paywall exposure. Figure A1.5 presents estimates for the baseline specification under different threshold levels. For easier comparison, the figure also reports the associated shares of survey respondents in the never-treated group. The effects remain quantitatively similar and at least marginally significant for a wide range of control group sizes.

Heterogeneity

To better understand which types of political knowledge are affected by paywalls, I decompose the knowledge index by estimating eight separate regressions, each using as the outcome one of the index components. Figure 1.10 shows results, with Panel (a) reporting the specification with year fixed effects only, and Panel (b) including state-by-year fixed effects instead. Estimates in black indicate conventional 95 percent confidence intervals, while those in red apply Bonferroni adjustments that account (conservatively) for multiple hypothesis testing. The four upper rows correspond to questions about the majority party in state or federal legislative bodies, and the four lower rows pertain to questions regarding party affiliation given the name of a political representative.

The results show that the effects on political knowledge are driven by reduced knowledge of both party majorities and the party affiliations of elected representatives. This

Figure 1.10: Effect of paywalls on political knowledge - Index decomposition

(a) County and year fixed effects

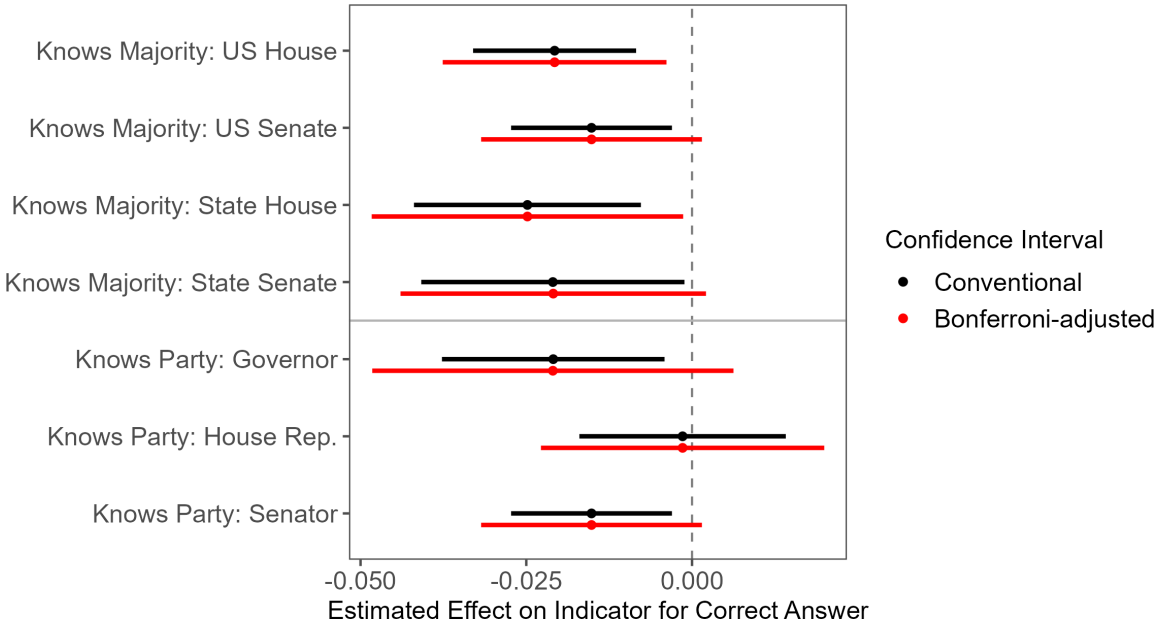

(b) County and state-by-year fixed effects

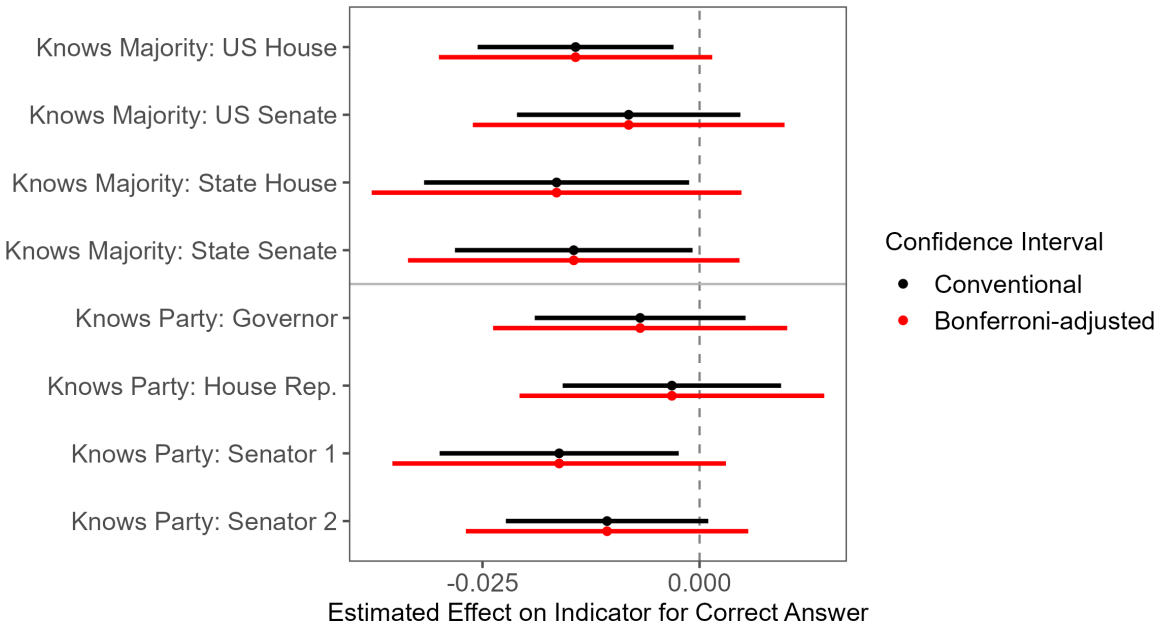

Notes: Each row represents an estimate for the effect of high exposure to paywalls on correctly answering the respective political knowledge question. Horizontal bars show bootstrapped 95 percent confidence intervals, estimated conventionally (black) or with Bonferroni adjustment (red). The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is an indicator for a correct answer (CES, 2006–2021), residualized on individual demographic controls, county fixed effects, and either year fixed effects (Panel (a)) or state-by-year fixed effects (Panel (b)). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors are clustered by county.

pattern is consistent across both specifications, though the estimates are somewhat attenuated in the more restrictive version. There is no clear evidence that knowledge of state-level politics declined more than knowledge of federal politics: While knowledge of majorities declined slightly more for state-level compared to national-level congressional elections, knowledge on House representatives (elected on district level) did not change at all. Overall, these findings align most closely with a general decline in attention to politics.

Which types of individuals are most affected by paywalls? To answer this question, Figure 1.11 shows results from estimating the baseline specification separately for subgroups of survey respondents defined by family income, education, or age. Individuals with low income and low education levels experience larger declines in political knowledge. This finding is consistent with the notion that these individuals face relatively higher monetary costs of purchasing a paywall subscription, and may derive lower subjective benefits from high-quality information, either due to less interest in or a lower perceived value of such content. However, effects seem to be similar for older and younger individuals.

Figure 1.11: Effect of paywalls on knowledge - Heterogeneity by individual characteristics

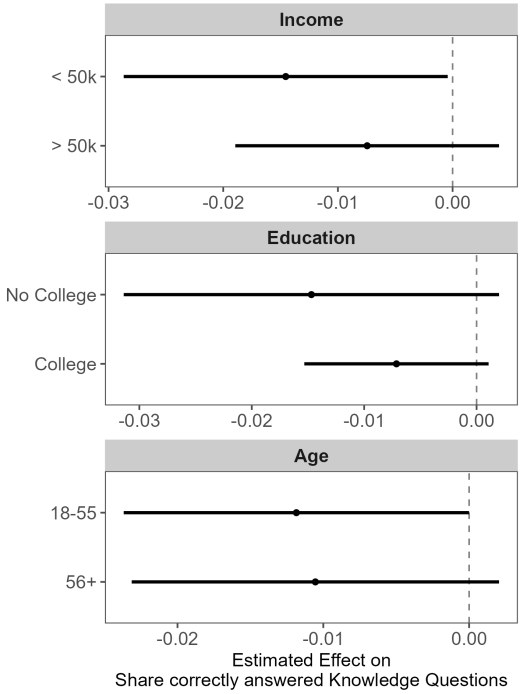

Notes: Estimates for the effect of high exposure to paywalls on political knowledge. Each panel reports coefficients from two separate estimations, corresponding to subgroups defined by yearly family income, reported education level, and age group, respectively. Horizontal bars show bootstrapped 95 percent confidence intervals. The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021), residualized on individual demographic controls, county fixed effects, and state-by-year fixed effects. Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors are clustered by county.

Mechanism

There are two distinct mechanisms that could explain the observed reduction in political knowledge due to paywalls. First, the declines could be driven by individuals who are unwilling to pay for a subscription, and instead either reduce their news consumption or switch to alternative sources on which they consume less political information. While the previous analysis in Section 3 demonstrates that many readers do switch away from paywalled news sources, it is not immediately clear whether these individuals are also those experiencing declines in knowledge. Hence, a second possibility is that newspapers altered their content along with the introduction of paywalls – for instance, by reducing political coverage or overall quality. In that case, the effect could be driven by individuals who continued consuming the now-paywalled newspapers.

To address this limitation, I complement the analysis with survey data that records whether individuals pay for news subscriptions. Specifically, I use the 2018 *Pew–Ipsos Local News Survey* of more than 34,000 US adults. This survey includes a question on whether respondents paid for a subscription to local news source in the past year, providing information on the types of individuals that are more likely to pay for news.

I estimate individuals' propensities to pay for news using categorical indicators for age, sex, marital status, education, income, ethnicity, census region, and party identification. I implement four prediction models: linear regression, logistic regression, random forest, and gradient-boosted random forest (XGBoost). For each method, I select the best-performing estimator based on the F1 score, a common metric that balanced precision and recall in the positive (paying) group. I tune any hyperparameters via a successive halving random search across a manually defined grid of approximately 10,000 candidate configurations, evaluated using four-fold cross-validation. To account for sampling probabilities as well as the smaller share of subscription purchasers in the sample, I apply the survey's sampling weights multiplied by inverse class weights.

Figure A1.6 shows the benchmark distribution of F1 scores for each model. XGBoost, the best-performing model, achieves an average F1 score of approximately 0.50, highlighting the difficulty of the classification task. Nonetheless, the prediction yields highly informative results. Performance considerably and consistently exceeds an uninformative ("dummy") model, which scores only at around 0.23. Cross-validated precision is 43 percent for predicted subscribers and 89 percent for predicted non-subscribers, compared to baseline rates of 14 and 86 percent, respectively. This represents a substantial improvement, particularly for the smaller and more relevant group of paying users.

All predictors contribute to the model's performance, though with varying intensities. Figure A1.7 displays impurity-based feature importance scores, which measure the improvement in the score of the gradient-boosted random forest attributable to each variable. Age emerges as the most important predictor, explaining almost half of the model's predictive power. However, other variables variables also play an important role, most notably, education and income.

I use the XGBoost model trained on the *Pew-Ipsos* data to predict subscription probabilities in the *CES*. As an initial check of external validity, I examine their correlation with two measures that reveal intensive engagement with news. Figure 1.12 displays binned scatter plots relating predicted subscription probabilities to these outcomes. Panel (a) plots these probabilities against indicators for whether the respondent reports following the news "most of the time," the highest of four response categories. Panel (b) shows the relationship with indicators for whether the respondent read a newspaper in the past 24 hours. Both figures reveal strong, monotonic, and positive correlations across the full distribution of predicted probabilities. These patterns confirm that the predicted subscription probabilities capture meaningful variation in real-world news engagement, consistent with higher actual propensities to purchase subscriptions.

Figure 1.12: Correlation of predicted probability to pay for news with proxies for news interest

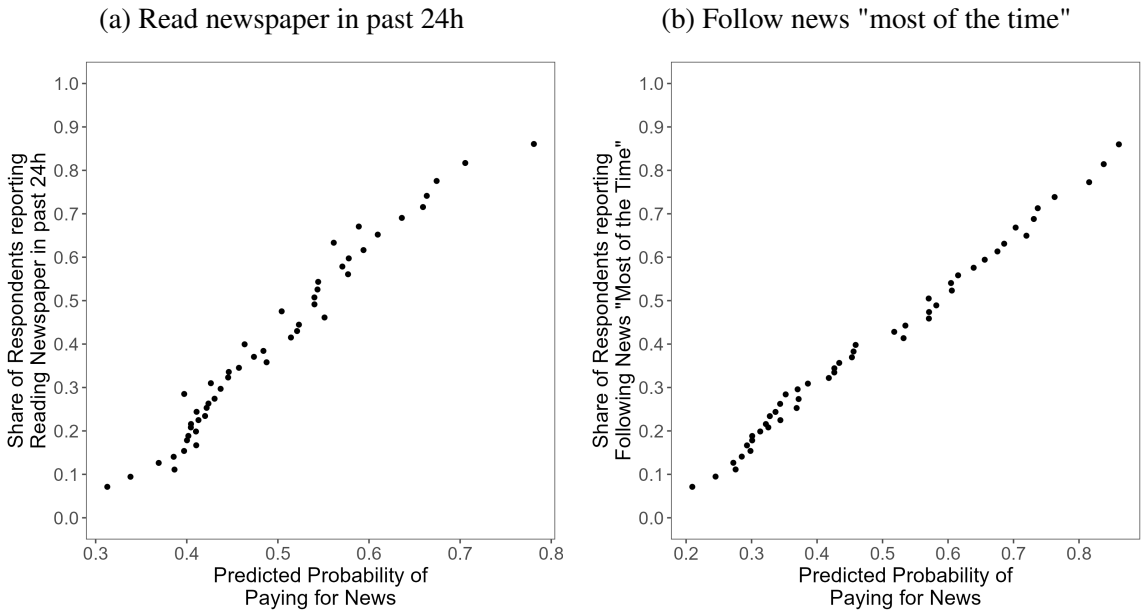

Notes: Binned scatter plots of CES respondents (2006–2021). The x-axis shows the predicted probability of paying for a news subscription, based on the best XGBoost classifier trained on the Pew Research Center (2018) survey. The y-axis in Panel (a) is an indicator for whether the respondent reported reading a newspaper in the past 24 hours; in Panel (b), for whether the respondent reports following the news "most of the time," the highest among four categories. Sample weights are included.

Using the predicted subscription probabilities, I repeat the previous analysis for the effect of paywalls on political knowledge as motivated by Equation (1.3), this time splitting the sample by whether survey respondents are predicted to pay for news. Table A1.6 provides motivational evidence by correlating paywall exposure with political knowledge separately for the two groups. Columns (1)-–(2) report estimates for respondents predicted *not* to pay for news, while Columns (3)—(4) report results for those predicted to pay. The negative correlation between paywalls and knowledge is entirely driven by the non-paying group. Among predicted subscribers, by contrast, the correlation is statistically indistinguishable from zero, despite only a modest loss in precision.

The event study estimates in Figure 1.13 confirm this pattern. The orange series corresponds to respondents predicted to pay for news, while the blue series reflects those predicted not to pay. The reduction in political knowledge is entirely driven by individuals predicted not to pay for news. Moreover, compared to the full sample, yearly effects are estimated with considerably higher precision. Instead, individuals predicted to pay for news are on a similar trajectory before treatment, but experience no changes in political knowledge after being exposed to paywalls.

Table 1.5 presents the corresponding average treatment effects. Among non-paying individuals, the estimated reduction in the share of correctly answered questions ranges from 1.5 to 1.9 percentage points, depending on the specification, which is slightly larger than the 1.2 to 1.8 percentage points estimated in the full sample. In contrast, for individuals predicted to pay for news, paywall exposure has no measurable effect on political knowledge in either specification.

These findings reject the hypothesis that the observed declines in political knowledge are driven by changes in the news reporting of paywalled news websites, such as content or slant. If such changes were responsible, one would expect the effects to be concentrated among the individuals who continue to access and engage most intensively with the newspapers, which are most likely those who purchase subscriptions. Instead, the results are most consistent with the alternative hypothesis: Individuals unwilling to pay for news may reduce their consumption of paywalled newspapers and either switch to alternative sources with less political content or lower their news consumption overall. The resulting decline in attention to politics may then lead to the observed reduction in political knowledge.

Figure 1.13: Effect of paywalls on political knowledge - Heterogeneity by paying for news

(a) County and year fixed effects

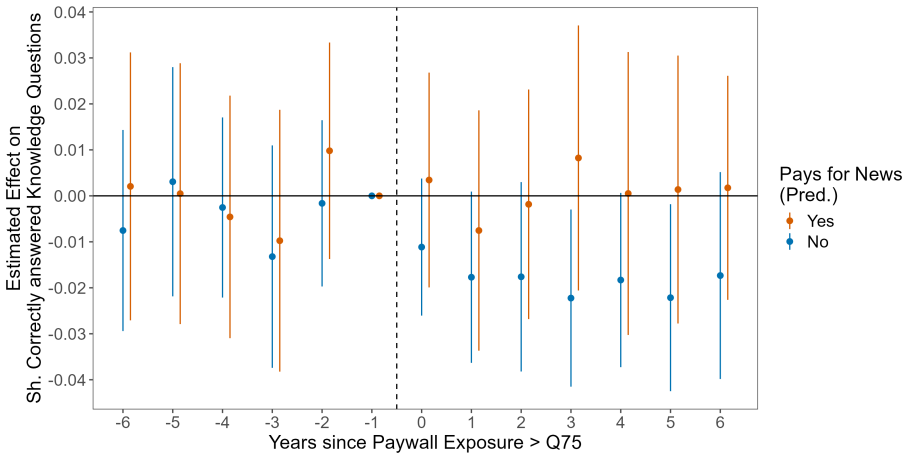

(b) County and state-by-year fixed effects

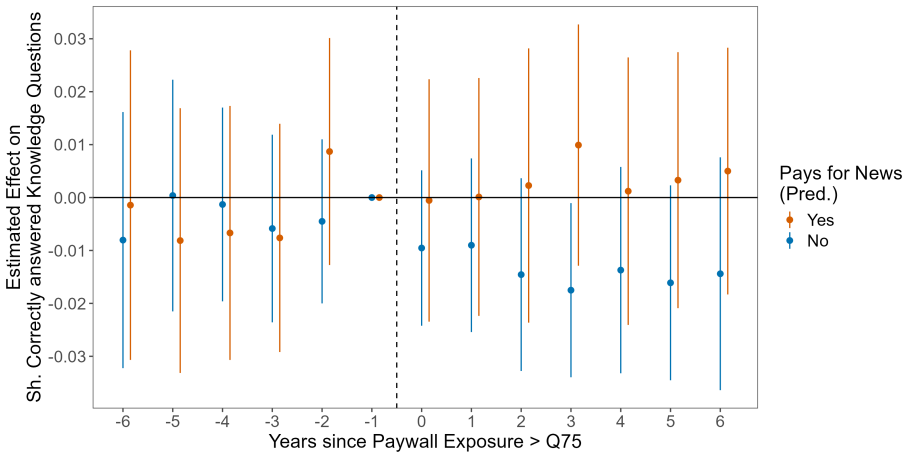

Notes: Yearly estimates for the effect of high exposure to paywalls on political knowledge, separately by whether individuals are predicted to pay for a news subscription. The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021), residualized on individual demographic controls, county fixed effects, and either year fixed effects (Panel (a)) or state-by-year fixed effects (Panel (b)). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Vertical bars show bootstrapped 95 percent uniform confidence intervals. Sample weights are included. Standard errors are clustered by county.

Table 1.5: Effect of paywalls on political knowledge - Heterogeneity by paying for news

Dependent Variable:	Share Correct Answers			
Paying for News (Pred.):	No		Yes	
	(1)	(2)	(3)	(4)
Paywall Exposure > Q75	-0.019**	-0.015**	-0.001	0.002
	(0.006)	(0.005)	(0.007)	(0.007)
Dep. Var. Mean	0.635	0.635	0.842	0.842
County FE	✓	✓	✓	✓
Year FE	✓		✓	
State-Year FE		✓		✓
Individual Controls	✓	✓	✓	✓
Observations	413,797	413,797	105,070	105,070

Notes: Estimates for the effect of high exposure to paywalls on political knowledge, estimated separately by whether individuals are predicted to pay for a news subscription. The specification is motivated by Equation (1.3), estimated and aggregated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021). All estimations include county and year fixed effects. Where indicated, the outcome is additionally residualized on individual demographic controls – categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2) – as well as county fixed effects and either year or state-by-year fixed effects. Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Magnitudes

In this subsection I discuss the magnitude of the estimated effect of paywalls on political knowledge. To facilitate comparisons with other studies in the media literature, I present a back-of-envelope calculation for the "persuasion rate" of paywalls, defined as the share of individuals who switched from the correct to the incorrect answer because of paywalls, among those exposed to paywalls.[17] Following the notation of DellaVigna and Gentzkow (2010), the persuasion rate for this setting is calculated as:

$$f = \frac{y_T - y_C}{e_T - e_C} \frac{1}{1 - y_0} \tag{1.4}$$

$$= \frac{0.012}{0.5 * 0.5} \frac{1}{0.745} = 0.048 \tag{1.5}$$

Here, $y_T - y_C$ is the change in the outcome between the treated and the control group in response to the treatment. To provide a conservative estimate, I rely on the smallest coefficient for the main effect which is 0.012. $e_T - e_C$ denotes the difference in the share of the population exposed to the treatment. I assume that around 25 percent of the voting-age population is exposed, resulting from the observations that around 50 percent of the survey respondents read newspapers online during the sample period, and that average

[17] In this setting, the term "persuasion rate" applies imperfectly, as individuals do not change their *opinion* in response to the treatment.

paywall exposure is around twice as high in the treated group compared to the control group. $1 - y_0$ represents the share of the population left to be persuaded. In this context, I use the average share of correct questions in the population at baseline, which is 0.745 (calculated using the years 2006 through 2010).

The resulting magnitude is 4.8 percent. On average, this means that approximately 4.8 percent among individuals who initially read a newspaper that implemented a paywall lost all their political knowledge of the type tested in the survey. As an equivalent interpretation, 9.6 percent answered only 3 instead of 6 out of the eight questions correctly, or only half as many as initially.

The estimated persuasion rate of 4.8 is in the lower range of estimates in the literature on the effects of news availability on political outcomes (DellaVigna and Gentzkow, 2010). For example, Enikolopov et al. (2011) find a persuasion rate of 7.7 percent from the availability of the independent anti-Putin TV station *NTV* on reducing the vote share for pro-Putin parties. Similarly, DellaVigna and Kaplan (2007) estimate an 11.6 percent persuasion rate from the availability of *Fox News* on the Republican vote share in presidential elections. Gerber et al. (2009) find a rate of 19.5 percent for free subscriptions to the *Washington Post* increasing intent to vote Democratic. Gentzkow et al. (2011) find a 12.9 percent persuasion rate for local newspaper readership in the late 1800s and early 1900s on voter turnout in presidential elections, while Gentzkow (2006) reports a 4.4 percent persuasion rate for early television exposure reducing congressional election turnout.

While informative, my calculation has some limitations. First, since I measure knowledge as an index comprised of several outcomes, the estimates reflect the effect of paywalls on the share of answers in the population, rather than on the number of individuals who lose knowledge. Therefore, the calculation implicitly assumes that the effect of paywalls on knowledge is homogeneous across individuals. Consequently, if reductions knowledge are concentrated among fewer individuals, the true persuasion rate is higher; if many individuals lose a smaller amount of knowledge, the rate is lower.

Second, the calculation should be interpreted as the intent-to-treat effect of being exposed to a paywall. It represents the net effect of two opposing mechanisms: increases in consumption by some individuals who purchase a subscription, and decreases in consumption by others who substitute away. Without more detailed data to disentangle these mechanisms, it is not possible to compute a persuasion rate specifically for the subset of readers who reduce their consumption due to paywalls.

4.5 Effects on electoral participation

Setup

To estimate the effects of paywalls on electoral participation, I use voting data from the *CES*, which is available in two forms: self-reported and validated. My primary outcome is participation in major elections – specifically, for president, governor, senators, and federal House representative. These measures are available in both self-reported and validated formats, enabling direct comparison. I construct an index of vote participation by calculating, for each respondent, the share of elections in which they participated, based on the elections held in their state and year. That way, the approach captures both the extensive and intensive margins of participation.

Because elections occur only every two years or less, I adapt the specification from Equation (1.3) to a two-year frequency, defining treatment status based on paywall exposure as of October in each election year. As in prior analyses, I apply the Callaway–Sant'Anna estimator, residualizing on individual characteristics as well as county and state-by-year fixed effects. Note that the latter are preferred because they effectively control for the number of elections in each state and year, which directly affects the outcome.

Main results

Table 1.6 reports results for different population subsets. Panel A uses self-reported turnout, and Panel B uses validated turnout. Columns (1)–(2) present estimates for the full sample. Including state-by-year fixed effects in Column (2) is important, as it absorbs variation in election exposure across states and years. While the estimated effects are negative, they are not statistically significant.

Motivated by the earlier finding that paywall effects on political knowledge are concentrated among individuals predicted not to pay for news, Column (3) restricts the sample to this group. Consistent with prior results, the estimated effects are larger in magnitude for both self-reported and validated participation. Moreover, they are statistically significant for self-reported participation.

Finally, Column (4) further restricts the sample by excluding the most partisan counties, defined as those with a partisanship index outside the 0.3-0.7 interval (0.5 indicates an evenly divided electorate between 2002 and 2010). The resulting more competitive counties are particularly relevant, as they feature stronger incentives for individuals to stay informed and are more likely to exhibit meaningful electoral consequences from shifts in participation. In this subsample, the estimated effects become more pronounced for both

self-reported and validated outcomes. The effect on self-reported participation is highly significant, while the effect on validated participation is significant at the 10 percent level.

A comparison between self-reported and validated voting data reveals evidence for misreporting. On average, self-reported turnout exceeds validated participation by approximately 10 percentage points. As noted by Ansolabehere and Hersh (2012), this discrepancy most likely reflects intentional overreporting driven by experimenter demand effects, although other factors such as faulty recall or errors in the validation procedure may also play a role. While such potential measurement error in either variable complicates the interpretation of absolute values, it affects my analysis only if *changes* in mismeasurement are correlated with paywall exposure, conditional on fixed effects and controls. In this respect, it is reassuring that effect sizes relative to the respective baseline are similar across both measures. Among respondents predicted not to pay for news, the average reduction in participation is roughly 2 percent for both self-reported and validated turnout.

Figure 1.14 presents event study estimates for the sample used in Column (4), restricted to non-paying individuals in less partisan counties. Panel (a) displays results for self-reported voting, while Panel (b) shows the corresponding estimates for validated participation. Both variables exhibit no discernible pre-treatment trends, but reveal a decline in participation during the first two post-treatment periods that remains stable afterwards. This pattern closely mirrors the event study results for political knowledge, supporting the interpretation that reduced information and lower turnout are connected. Taken together, the findings suggest that paywalls, by lowering political engagement, ultimately reduce participation in the democratic process.

Heterogeneity

Figure 1.15 and Table A1.7 present a decomposition of the vote participation index by election type. I estimate separate regressions using as the outcome a binary indicator for participation in each election. In addition to the core outcomes that are recorded both as self-reported and as validated variables, I examine participation for other elections available in the survey. I restrict the sample to individuals predicted not to pay for news.

The results show that the declines in participation are driven by local elections. Although all point estimates are negative or zero, the most robust declines across both self-reported and validated outcomes are observed House elections held at the congressional district level. In line with this result, participation in state-level congressional elections (self-reported) also sees large reductions. Instead, the results for elections of Senators and Governors are less clear, while the estimates for presidential elections are close to zero.

Table 1.6: Effect of paywalls on electoral participation

Dependent Variable:	Participation in Elections			
Paying for News (Pred.):	All		No	
County Partisanship Index:	All			0.3-0.7
	(1)	(2)	(3)	(4)
Panel A: Self-Reported				
Paywall Exposure > Q75	-0.0092	-0.0056	-0.0119**	-0.0147***
	(0.0068)	(0.0049)	(0.0060)	(0.0055)
Dep. Var. Mean	0.6072	0.6072	0.5736	0.5739
County FE	✓	✓	✓	✓
Year FE	✓			
State-Year FE		✓	✓	✓
Individual Controls	✓	✓	✓	✓
Observations	418,328	418,328	330,193	279,457
Panel B: Validated				
Paywall Exposure > Q75	-0.0142**	-0.0055	-0.0090	-0.0115*
	(0.0064)	(0.0051)	(0.0057)	(0.0063)
Dep. Var. Mean	0.5089	0.5089	0.4798	0.4814
County FE	✓	✓	✓	✓
Year FE	✓			
State-Year FE		✓	✓	✓
Individual Controls	✓	✓	✓	✓
Observations	383,040	383,040	300,286	253,940

Notes: Estimates for the effect of high exposure to paywalls on electoral participation. The specification is motivated by Equation (1.3), estimated and aggregated using the Callaway-Sant'Anna estimator. The outcome is the share of elections participated in among elections held for president, house representative, senators, and governor (CES, 2008–2020). Panel A uses self-reported voting; Panel B uses validated data from matching survey respondents to official records (Ansolabehere and Hersh, 2012). All specifications residualize the outcome on individual demographic controls – categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2) – as well as county fixed effects. Column (1) adds year fixed effects, and Columns (2)–(4) add state-by-year fixed effects. Columns (1)–(2) use the full sample. Column (3) restricts to individuals predicted not to pay for news. Column (4) further restricts to counties with a 2010 partisanship index between 0.3 and 0.7, (see Chenoweth et al., 2020). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 1.14: Effect of paywalls on participation in elections - Individuals not paying for news, Non-partisan counties

(a) Self-reported

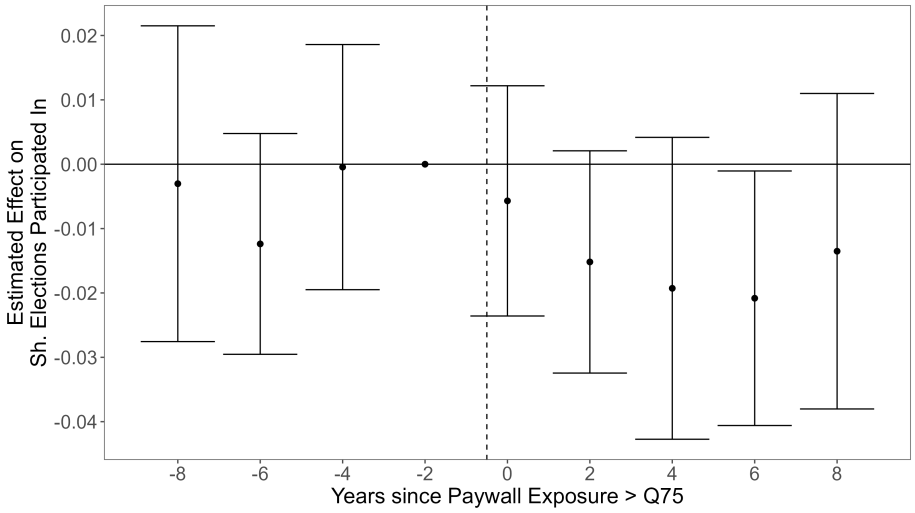

(b) Validated

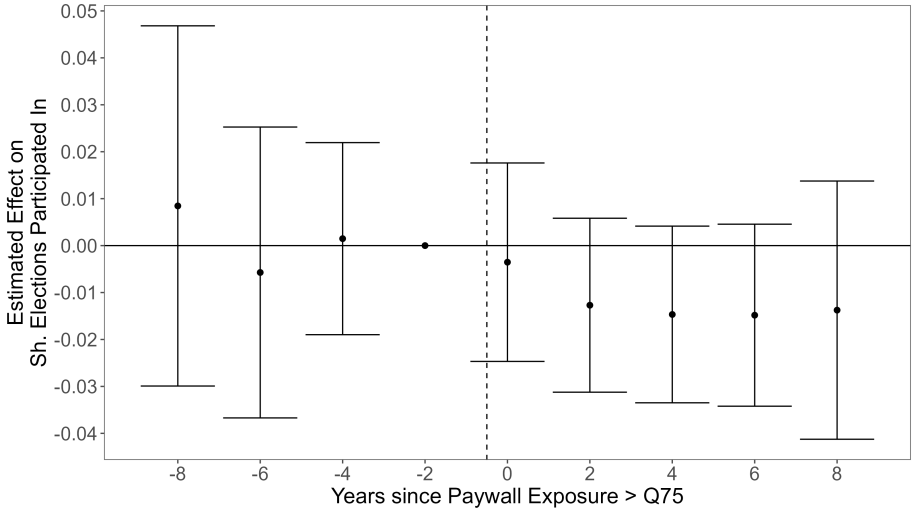

Notes: Bi-yearly estimates for the effect of high exposure to paywalls on electoral participation, among individuals predicted not to pay for a news subscription and residing counties with low partisanship. The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is the share of elections participated in among elections held for president, house representative, senators, and governor (CES, 2008–2020), residualized on individual demographic controls, county fixed effects, and state-by-year fixed effects. Panel (a) uses self-reported voting data, while Panel (b) uses validated data from matching survey respondents to official voting records (Ansolabehere and Hersh, 2012). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Vertical bars show bootstrapped 90 percent uniform confidence intervals. Sample weights are included. Standard errors are clustered by county.

These patterns may reflect two related mechanisms. First, presidential and gubernatorial races typically receive higher coverage than congressional elections across all news sources. Therefore, individuals substituting away from paywalled newspapers may have remained informed about these higher-profile races, while becoming less engaged with more local contests, leading to lower turnout. Second, local newspapers, which are more likely to cover local elections, were more heavily affected by paywalls relative to national outlets. If readers shifted away from these local sources, exposure to information about district-level races may have declined disproportionally, leading to lower participation.

Figure 1.15: Effect of paywalls on participation in elections - Index decomposition (Individuals not paying for news)

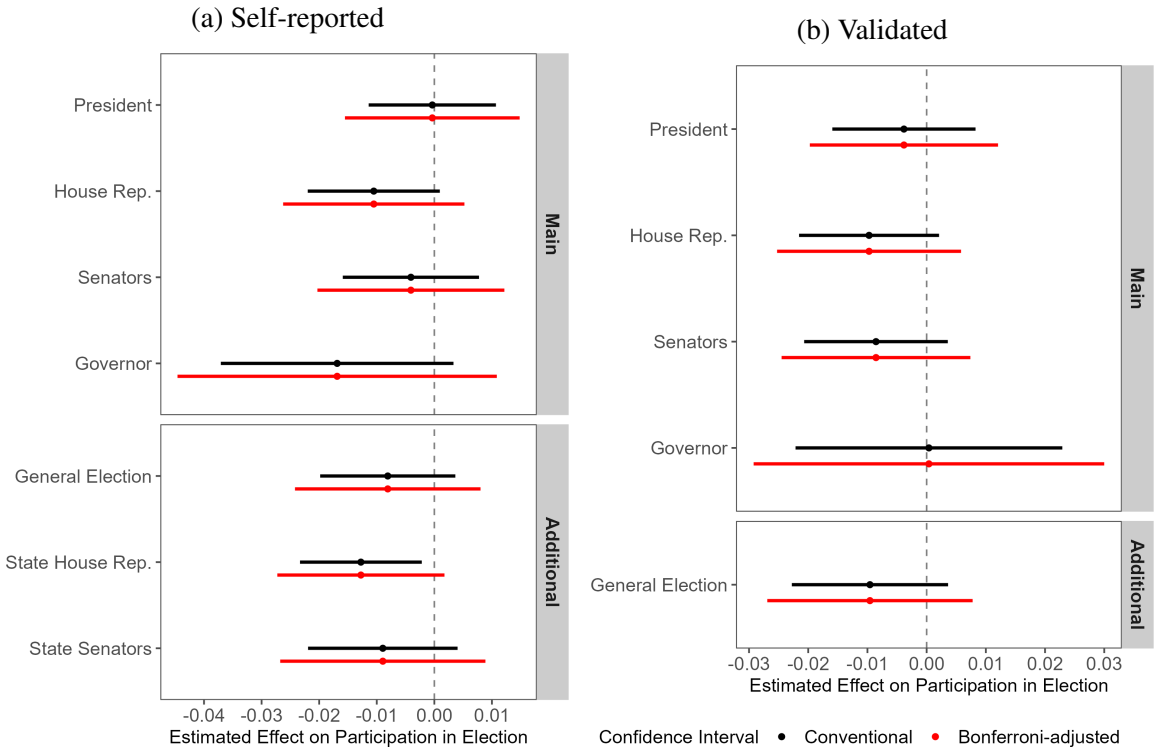

Notes: Each row represents an estimate for the effect of high exposure to paywalls on participating in the respective election, among individuals predicted not to pay for a news subscription. Panel (a) uses self-reported voting data, while Panel (b) uses validated data from matching survey respondents to official voting records (Ansolabehere and Hersh, 2012). In each panel, the top part refers to the elections included in the main index for participation, while the bottom part refers to other available data in the CES. Horizontal bars show bootstrapped 95 percent confidence intervals, estimated conventionally (black) or with Bonferroni adjustment (red). The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is an indicator for participating in the election (CES, 2008–2020), residualized on individual demographic controls, county fixed effects, and state-by-year fixed effects. Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors are clustered by county.

5 Conclusion

In this paper, I demonstrate that the introduction of paywalls on US newspaper websites reduced political knowledge and electoral participation. Paywalls led to substantial declines in online consumption of traditional newspapers. As a result, regions highly exposed to paywalls experienced declines in factual knowledge about contemporary politics, especially among less-educated and low-income individuals. These findings suggest that such groups are less willing or able to pay for digital news, leaving them disproportionately uninformed. In turn, paywalls lowered electoral participation among these populations, especially in politically competitive regions.

My findings underscore that news media play a critical role for the democratic process. In a media landscape dominated by private, profit-oriented news outlets, economic shocks can require business decision that ultimately affect voters' information and participation. Therefore, to counteract the potentially adverse political effects of such economic shifts, modern democracies should ensure that broad and equitable access to relevant information.

To better understand both the mechanisms behind my results as well as potential remedies, further research could explore how readers substitute paywalled websites, and examine supply-side changes to topics, slant, or quality that may accompany paywalls. Moreover, one could further investigate downstream effects on political opinions and party preference, particularly in local elections. Finally, it would be valuable to determine whether the displacement of readers from traditional media contributed to the rise of right-wing populism in the following years.

References

Angelucci, Charles and Julia Cagé, "Newspapers in times of low advertising revenues," *American Economic Journal: Microeconomics*, 2019, *11* (3), 319–64.

Ansolabehere, Stephen and Brian Schaffner, "CES Common Content, 2021," 2022.

__ **and Eitan Hersh**, "Validation: What big data reveal about survey misreporting and the real electorate," *Political Analysis*, 2012, *20* (4), 437–459.

Aral, Sinan and Paramveer S Dhillon, "Digital paywall design: Implications for content demand and subscriptions," *Management Science*, 2021, *67* (4), 2381–2402.

Ash, Elliott and Sergio Galletta, "How cable news reshaped local government," *American Economic Journal: Applied Economics*, 2023, *15* (4), 292–320.

Baker, Andrew C, David F Larcker, and Charles CY Wang, "How much should we trust staggered difference-in-differences estimates?," *Journal of Financial Economics*, 2022, *144* (2), 370–395.

Bartik, Timothy J, "Who benefits from state and local economic development policies?," 1991.

Besley, Timothy and Robin Burgess, "The political economy of government responsiveness: Theory and evidence from India," *The quarterly journal of economics*, 2002, *117* (4), 1415–1451.

Bhuller, Manudeep, Tarjei Havnes, Jeremy McCauley, and Magne Mogstad, "How the internet changed the market for print media," *American Economic Journal: Applied Economics*, 2024, *16* (2), 318–358.

Blanchard, Olivier Jean, Lawrence F Katz, Robert E Hall, and Barry Eichengreen, "Regional evolutions," *Brookings papers on economic activity*, 1992, *1992* (1), 1–75.

Cagé, Julia, "Media competition, information provision and political participation: Evidence from French local newspapers and elections, 1944–2014," *Journal of Public Economics*, 2020, *185*, 104077.

Callaway, Brantly and Pedro HC Sant'Anna, "Difference-in-differences with multiple time periods," *Journal of econometrics*, 2021, *225* (2), 200–230.

Campante, Filipe, Ruben Durante, and Francesco Sobbrio, "Politics 2.0: The multifaceted effect of broadband internet on political participation," *Journal of the European Economic Association*, 2018, *16* (4), 1094–1136.

Chenoweth, Megan, Mao Li, Iris N Gomez-Lopez, and Ken Kollman, "National Neighborhood Data Archive (NaNDA): Voter Registration, Turnout, and Partisanship by County, United States, 2004-2018," *Ann Arbor, MI: Inter-university Consortium for Political and Social Research[distributor]*, 2020, pp. 11–04.

Chiou, Lesley and Catherine Tucker, "Paywalls and the demand for news," *Information Economics and Policy*, 2013, *25* (2), 61–69.

Chung, Doug J, Ho Kim, and Reo Song, *The comprehensive effects of a digital paywall sales strategy*, Harvard Business School, 2019.

Cook, Jonathan E and Shahzeen Z Attari, "Paying for what was free: Lessons from the New York Times paywall," *Cyberpsychology, behavior, and social networking*, 2012, *15* (12), 682–687.

DellaVigna, Stefano and Ethan Kaplan, "The Fox News effect: Media bias and voting," *The Quarterly Journal of Economics*, 2007, *122* (3), 1187–1234.

__ **and Matthew Gentzkow**, "Persuasion: empirical evidence," *Annu. Rev. Econ.*, 2010, *2* (1), 643–669.

Djourelova, Milena, Ruben Durante, and Gregory J Martin, "The impact of online competition on local newspapers: Evidence from the introduction of Craigslist," *Review of Economic Studies*, 2024, p. rdae049.

Drago, Francesco, Tommaso Nannicini, and Francesco Sobbrio, "Meet the press: How voters and politicians respond to newspaper entry and exit," *American Economic Journal: Applied Economics*, 2014, *6* (3), 159–88.

Durante, Ruben, Paolo Pinotti, and Andrea Tesei, "The political legacy of entertainment TV," *American Economic Review*, 2019, *109* (7), 2497–2530.

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya, "Media and political persuasion: Evidence from Russia," *American Economic Review*, 2011, *101* (7), 3253–85.

Falck, Oliver, Robert Gold, and Stephan Heblich, "E-lections: Voting Behavior and the Internet," *American Economic Review*, 2014, *104* (7), 2238–65.

Ferraz, Claudio and Frederico Finan, "Exposing corrupt politicians: the effects of Brazil's publicly released audits on electoral outcomes," *The Quarterly journal of economics*, 2008, *123* (2), 703–745.

Gao, Pengjie, Chang Lee, and Dermot Murphy, "Financing dies in darkness? The impact of newspaper closures on public finance," *Journal of Financial Economics*, 2020, *135* (2), 445–467.

Gavazza, Alessandro, Mattia Nardotto, and Tommaso Valletti, "Internet and politics: Evidence from UK local elections and local government policies," *The Review of Economic Studies*, 2019, *86* (5), 2092–2135.

Gentzkow, Matthew, "Television and voter turnout," *The Quarterly Journal of Economics*, 2006, *121* (3), 931–972.

_ **and Jesse M Shapiro**, "Introduction of Television to the United States Media Market, 1946-1960," 2008.

_ **,** _ **, and Michael Sinkinson**, "The effect of newspaper entry and exit on electoral politics," *American Economic Review*, 2011, *101* (7), 2980–3018.

_ **,** _ **, and** _ , "Competition and ideological diversity: Historical evidence from us newspapers," *American Economic Review*, 2014, *104* (10), 3073–3114.

George, Lisa M and Joel Waldfogel, "The New York Times and the market for local newspapers," *American Economic Review*, 2006, *96* (1), 435–447.

Gerber, Alan S, Dean Karlan, and Daniel Bergan, "Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions," *American Economic Journal: Applied Economics*, 2009, *1* (2), 35–52.

Guriev, Sergei, Nikita Melnikov, and Ekaterina Zhuravskaya, "3g internet and confidence in government," *The Quarterly Journal of Economics*, 2021, *136* (4), 2533–2613.

Jr, James M Snyder and David Strömberg, "Press coverage and political accountability," *Journal of political Economy*, 2010, *118* (2), 355–408.

Kim, Ho, Reo Song, and Youngsoo Kim, "Newspapers' Content Policy and the Effect of Paywalls on Pageviews," *Journal of Interactive Marketing*, 2020, *49*, 54–69.

MRI-Simmons, "NCS Adult Study 12 Month," 2014.

Pattabhiramaiah, Adithya, S Sriram, and Puneet Manchanda, "Paywalls: Monetizing online content," *Journal of marketing*, 2019, *83* (2), 19–36.

Pew Research Center, "The State of the News Media 2013: An Annual Report on American Journalism," 2013. Accessed on January 23, 2025.

_ , "For Local News, Americans Embrace Digital but Still Want Strong Community Connection," 2018. Accessed on April 28, 2025.

_ , "State of the News Media 2023," 2023. Accessed on January 23, 2025.

Sant'Anna, Pedro HC and Jun Zhao, "Doubly robust difference-in-differences estimators," *Journal of econometrics*, 2020, *219* (1), 101–122.

Seamans, Robert and Feng Zhu, "Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers," *Management Science*, 2014, *60* (2), 476–493.

Strömberg, David, "Radio's impact on public spending," *The Quarterly Journal of Economics*, 2004, *119* (1), 189–221.

Sun, Liyang and Sarah Abraham, "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

U.S. Census Bureau, "American Community Survey 5-Year Estimates: Comparison Profiles 5-Year," 2010. Accessed on April 22, 2024.

World Bank, "Internet users for the United States [ITNETUSERP2USA]," 2025. Retrieved from FRED, Federal Reserve Bank of St. Louis.

Appendix

Figure A1.1: State of the US newspaper market in the 2000s

(a) Print circulation and share of internet users

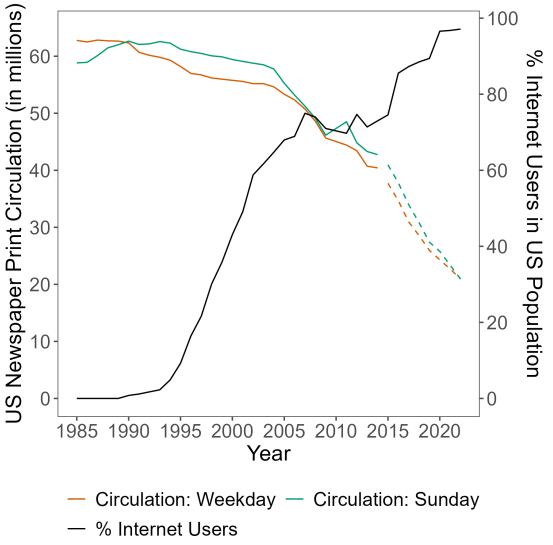

(b) Revenues by type

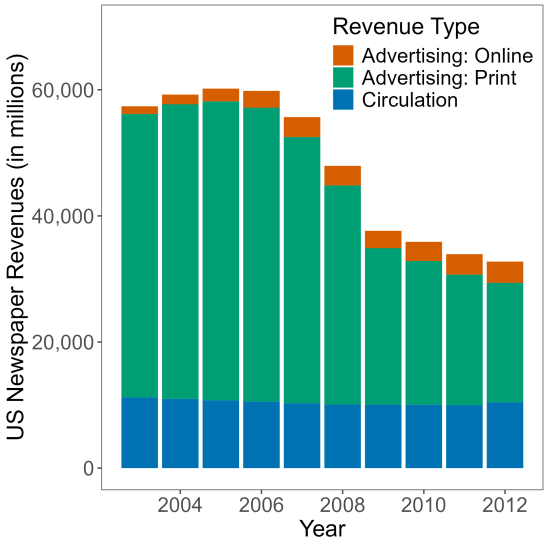

Notes: Panel (a) shows the total number of print copies sold of US daily newspapers by year. The orange and green line represent the circulation of the newspapers' weekday and Sunday editions, respectively. Dotted lines are estimates (Pew Research Center, 2023). The black line represents the share of the US population with access to the internet (World Bank, 2025). Panel (b) shows the composition of revenues among US daily newspapers by year. Orange and green bars represent revenues from online and print advertising, respectively, while the blue bars represent revenues from the sale of physical and virtual newspaper copies (Pew Research Center, 2013).

Figure A1.2: Effect of paywalls on website page views - Estimators

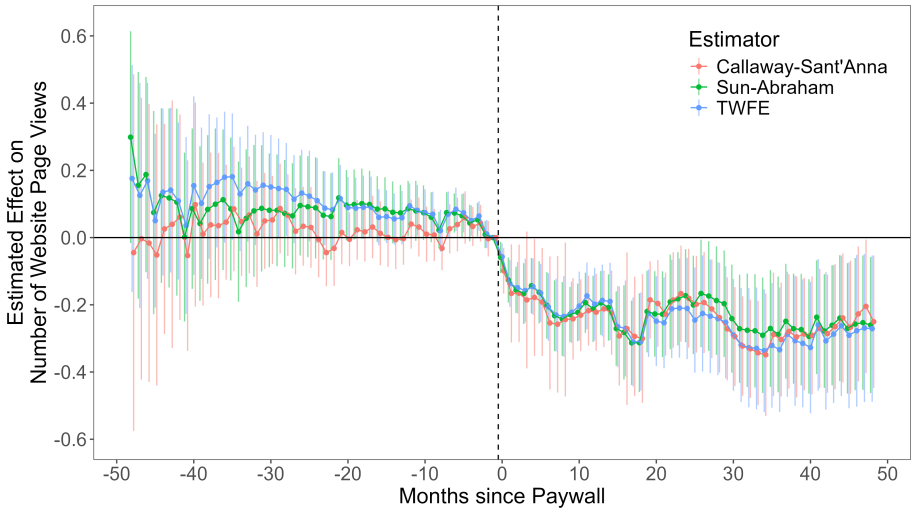

Notes: Monthly coefficients for β_τ from Equation (1.1), estimated either via Callaway-Sant'Anna, Sun-Abraham, or OLS (TWFE): Regression of monthly number of newspaper website views on indicators denoting the number of months since paywall implementation. The omitted category is the month before the paywall. Includes newspaper fixed effects, year-month fixed effects, and log population density, shares of three income buckets, college-educated share, and partisanship index for newspapers' audiences, interacted with year-month fixed effects. Vertical bars denote 95 percent pointwise confidence intervals. Standard errors are clustered by newspaper.

Table A1.1: Effect of paywalls on pageviews - Estimators

Dependent Variable:	Log(Pageviews)					
Estimator:	TWFE		Sun-Abraham		Callaway-Sant'Anna	
	(1)	(2)	(3)	(4)	(5)	(6)
Paywall	-0.258***	-0.273***	-0.267***	-0.279***	-0.257***	-0.280***
	(0.059)	(0.060)	(0.074)	(0.074)	(0.069)	(0.065)
Newspaper FE	✓	✓	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓	✓	✓
Market Controls		✓		✓		✓
Observations	6,912	6,912	6,912	6,912	6,912	6,912
R^2	0.957	0.960	0.978	0.981		

Notes: Coefficients for static (pre-post) version of Equation (1.1), estimated via Ordinary Least Squares (TWFE), Sun-Abraham, or Callaway-Sant'Anna: Regressions of monthly number of newspaper website views on indicator denoting active paywall, controlling for newspaper fixed effects and year-month fixed effects. Columns (2), (4), and (6) additionally control for average county characteristics at 2010-levels where newspaper is active, weighted by number of print copies sold, interacted with year-month fixed effects: Log population density, share of households with yearly income below $50,000 and above $100,000, share adults with college degree, and NaNDA partisanship index. Standard errors in parentheses are clustered by newspaper. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A1.2: Effect of paywalls on pageviews - Restrict sample

Dependent Var.:	Log(Pageviews)				
	Exclude Newspapers				Add Weights
	National	Top 15	Top 15 Paywalled	Bottom 5 Control	Weights
	(1)	(2)	(3)	(4)	(5)
Paywall	-0.282***	-0.332***	-0.357***	-0.320***	-0.273***
	(0.065)	(0.071)	(0.073)	(0.075)	(0.068)
Newspaper FE	✓	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓	✓
Market Controls	✓	✓	✓	✓	✓
Restricted Control Group			✓	✓	
Sample Weights					Log(Pageviews 2010)
Observations	6,624	5,472	5,472	6,432	6,912

Notes: Coefficients for static (pre-post) version of Equation (1.1), excluding different subsets of newspapers from the sample. Column (1) excludes the three national newspapers *New York Times*, *USA Today*, and *Wall Street Journal*. Column (2) excludes the largest 15 newspapers in 2010 by website page views. Columns (3) and (4) exclude the largest 15 newspapers in the treated group, and the bottom 5 newspapers in the control group, respectively. In both columns, treated observations pre-treatment are also excluded from the control group. Column (5) includes the log of 2010 page views as sample weights. Regressions of monthly number of newspaper website views on indicator denoting active paywall. All specifications include newspaper fixed effects and year-month fixed effects. Controls represent average county characteristics at 2010-levels where newspaper is active, weighted by number of print copies sold, interacted with year-month fixed effects: Log population density, share of households with yearly income below $50,000 and above $100,000, share adults with college degree, and NaNDA partisanship index. Estimated using Callaway-Sant'Anna estimator. Standard errors in parentheses are clustered by newspaper. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A1.3: Geographical variation of paywall exposure - Residualized

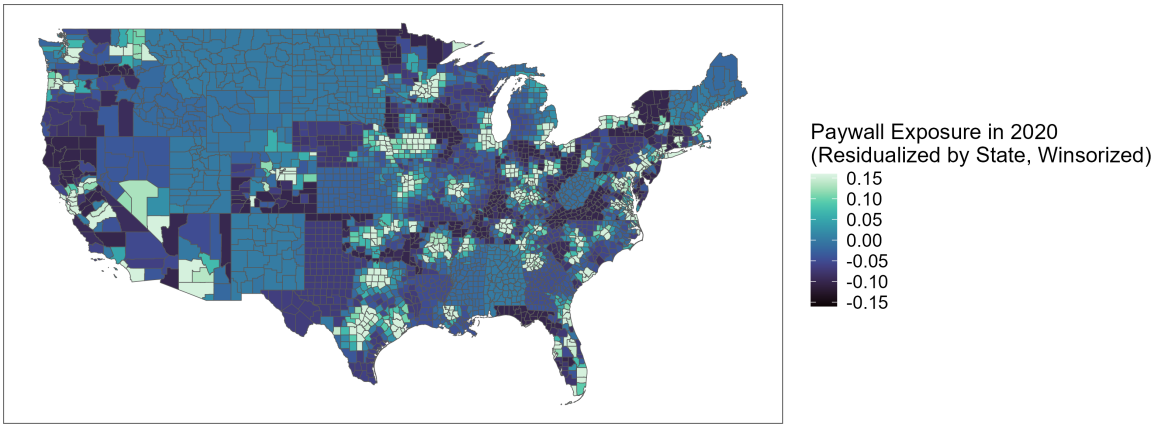

Notes: Paywall exposure by county in 2020, residualized by state and winsorized at 10 and 90 percent.

Figure A1.4: County-level predictors of paywall exposure

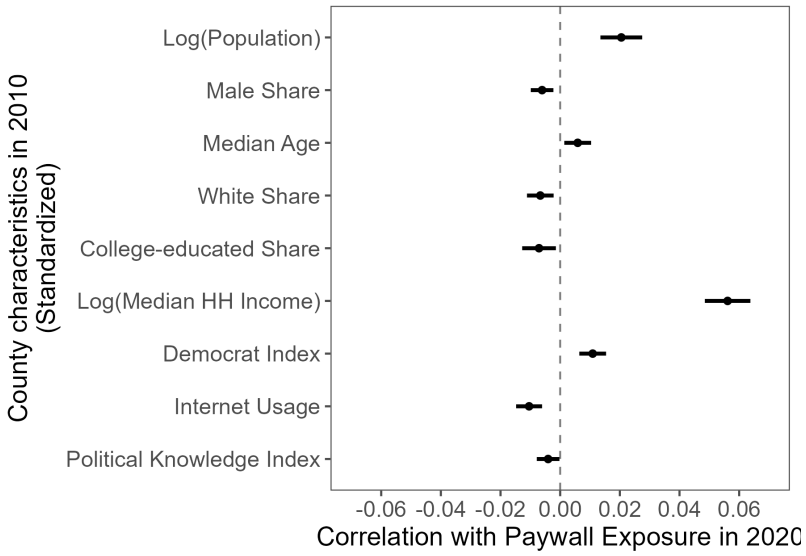

Notes: Multivariate regression of county-level paywall exposure in 2020 on county characteristics in 2010. All variables are standardized. Horizontal bars denote 95 percent confidence intervals.

Table A1.3: Correlation of (continuous) paywall exposure and political knowledge

Dependent Variable:	Share Correct Answers					
	(1)	(2)	(3)	(4)	(5)	(6)
Paywall Exposure	-0.047***	-0.025**	-0.022**	-0.024*	-0.016*	-0.016*
	(0.013)	(0.011)	(0.011)	(0.013)	(0.009)	(0.010)
Dep. Var. Mean	0.745	0.745	0.745	0.745	0.745	0.745
County FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓			
State-Year FE				✓	✓	✓
Indiv. Demographics		✓			✓	
Indiv. Demographics x Year FE			✓			✓
Observations	529,409	517,924	517,924	529,409	517,924	517,924
R^2	0.080	0.329	0.320	0.091	0.338	0.328

Notes: Estimates from regressing the share of correct answers to eight political knowledge questions on paywall exposure. All specifications control for fixed effects as indicated. Individual demographic controls include categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2). Data are from the CES 2006–2021. Survey weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A1.4: Correlation of (binned) paywall exposure and political knowledge

Dependent Variable:	Share Correct Answers					
	(1)	(2)	(3)	(4)	(5)	(6)
Paywall Exposure, 2nd Tercile	-0.001	-0.004	-0.002	-0.002	-0.003	-0.004
	(0.005)	(0.004)	(0.004)	(0.006)	(0.005)	(0.005)
Paywall Exposure, 3rd Tercile	-0.009	-0.008*	-0.007	-0.007	-0.008	-0.009*
	(0.006)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)
Dep. Var. Mean	0.745	0.745	0.745	0.745	0.745	0.745
County FE	✓	✓	✓	✓	✓	✓
Year FE	✓	✓	✓			
State-Year FE				✓	✓	✓
Indiv. Demographics		✓			✓	
Indiv. Demographics x Year FE			✓			✓
Observations	437,835	437,039	437,039	437,835	437,039	437,039
R^2	0.066	0.326	0.312	0.075	0.334	0.320

Notes: Estimates from regressing the share of correct answers to eight political knowledge questions on indicators for terciles of paywall exposure. All specifications control for fixed effects as indicated. Individual demographic controls include categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2). Data are from the CES 2006–2021. Survey weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A1.5: Effect of paywalls on political knowledge - Thresholds

(a) County and year fixed effects

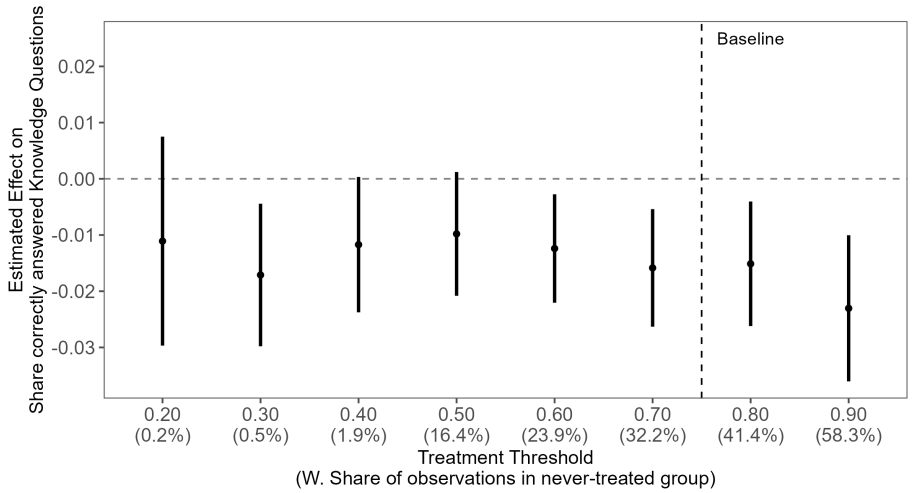

(b) County and state-by-year fixed effects

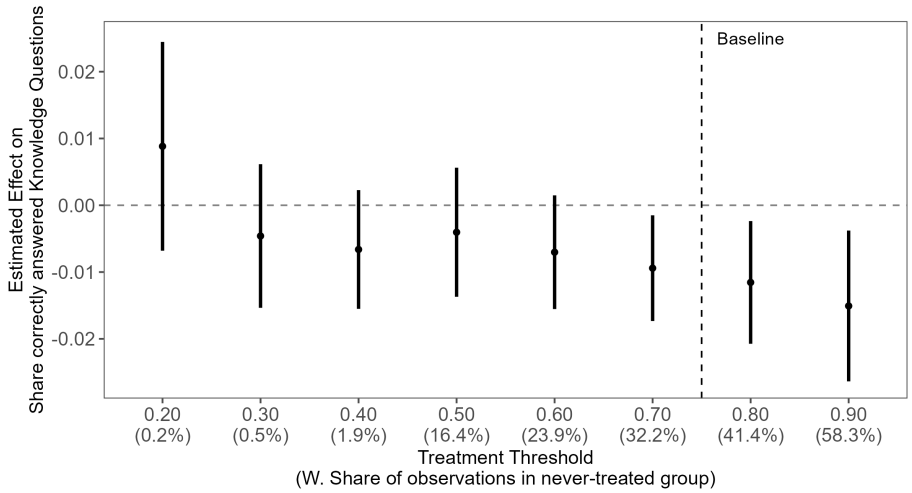

Notes: Estimates for the effect of high exposure to paywalls on political knowledge, for different percentile thresholds used to define when a county switches from low to high paywall exposure. The numbers in brackets indicate the sample-weighted share of observations in the never-treated group, across all years. The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021), residualized on individual demographic controls, county fixed effects, and either year fixed effects (Panel (a)) or state-by-year fixed effects (Panel (b)). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Vertical bars show bootstrapped 95 percent confidence intervals. Sample weights are included. Standard errors are clustered by county.

Table A1.5: Effect of paywalls on political knowledge - Heterogeneity by individual characteristics

Dependent Variable:	Share Correct Answers					
Subsample:	Income		Education		Age	
	$\leq$ 50k	$>$ 50k	No College	College	18-55	56+
	(1)	(2)	(3)	(4)	(5)	(6)
Paywall Exposure $>$ Q75	-0.015**	-0.007	-0.015*	-0.007*	-0.012**	-0.011
	(0.007)	(0.006)	(0.009)	(0.004)	(0.006)	(0.006)
Dep. Var. Mean	0.555	0.759	0.621	0.764	0.552	0.746
County FE	✓	✓	✓	✓	✓	✓
State-Year FE	✓	✓	✓	✓	✓	✓
Individual Controls	✓	✓	✓	✓	✓	✓
Observations	215,551	247,899	160,579	358,288	307,134	211,733

Notes: Estimates for the effect of high exposure to paywalls on political knowledge, estimated separately for different subgroups defined by yearly family income (Columns 1–2), highest attained educational degree (Columns 3-4), and age (Columns 5-6). The specification is motivated by Equation (1.3), estimated and aggregated using the Callaway-Sant'Anna estimator. The outcome is the share of correct answers by survey respondents to eight political knowledge questions (CES, 2006–2021), residualized on individual demographic controls – categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2) – as well as county and state-by-year fixed effects. Counties are treated once paywall exposure exceeds the 75$^{\text{th}}$ percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A1.6: Correlation of (continuous) paywall exposure and political knowledge - Heterogeneity by paying for news

Dependent Variable:	Share Correct Answers			
Pays for News (Pred.):	No		Yes	
	(1)	(2)	(3)	(4)
Paywall Exposure	-0.029**	-0.022**	-0.003	0.010
	(0.011)	(0.011)	(0.015)	(0.015)
Dep. Var. Mean	0.693	0.693	0.856	0.856
County FE	✓	✓	✓	✓
Year FE	✓		✓	
State-Year FE		✓		✓
Indiv. Demographics	✓	✓	✓	✓
Observations	413,797	413,797	105,070	105,070
R^2	0.297	0.307	0.285	0.308

Notes: Estimates from regressing the share of correct answers to eight political knowledge questions on paywall exposure, separately by whether individuals are predicted to pay for a news subscription. All specifications control for fixed effects as indicated, as well as individual demographic controls comprised of categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2). Data are from the CES 2006–2021. Survey weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A1.6: Prediction of paying for news - Benchmark of prediction algorithms

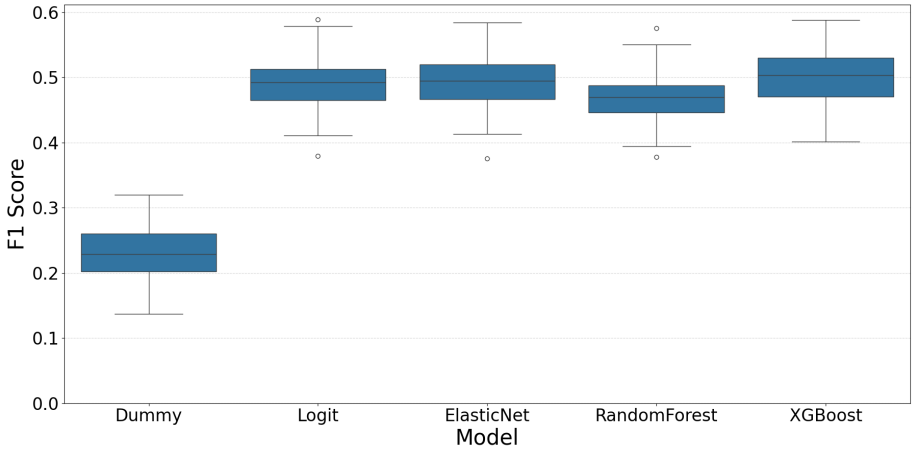

Notes: Distribution of F1 scores from 100-fold cross-validation of the best estimator in each model category. The dummy classifier predicts a positive outcome at random, with probability equal to the share of positive cases in the training data. For the other classifiers, the best estimator was selected via successive halving random search over a manually defined starting grid, using 4-fold cross-validation. The target is an indicator for whether an individual reports paying for a news subscription, based on the Pew Research Center (2018) survey of 34,518 respondents. Explanatory variables include categorical indicators for age (4 groups), education (6), family income (9), ethnicity (2), party preference (4), marital status (2), region (4), and sex (2). Sample weights are applied during training.

Figure A1.7: Prediction of paying for news - Feature importances

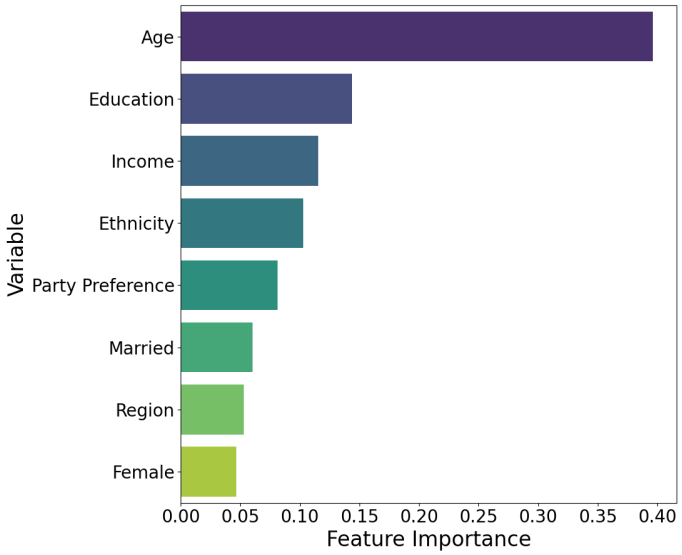

Notes: Impurity-based feature importances for variable categories from the best XGBoost model predicting whether an individual pays for news, using data from the Pew Research Center (2018) survey. Explanatory variables include categorical indicators for age (4 groups), education (6), family income (9), ethnicity (2), party preference (4), marital status (2), region (4), and sex (2). Sample weights are applied during training.

Table A1.7: Effect of paywalls on electoral participation - Index Decomposition (Individuals not Paying for News)

Dependent Variable:	Participation in Election						
Paying for News (Pred.):	No						
	President	House Rep.	Senators	Governor	General, any	State Rep.	State Sen.
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Self-Reported							
Paywall Exposure > Q75	-0.0021	-0.0139**	-0.0097	-0.0225**	-0.0114**	-0.0153**	-0.0121*
	(0.0060)	(0.0060)	(0.0068)	(0.0090)	(0.0058)	(0.0062)	(0.0063)
Dep. Var. Mean	0.6460	0.5430	0.5920	0.5690	0.7030	0.5450	0.5360
County FE	✓	✓	✓	✓	✓	✓	✓
State-Year FE	✓	✓	✓	✓	✓	✓	✓
Individual Controls	✓	✓	✓	✓	✓	✓	✓
Observations	142,277	279,101	187,454	130,002	219,784	245,259	241,568
Panel B: Validated							
Paywall Exposure > Q75	-0.0048	-0.0135*	-0.0150**	-0.0051	-0.0117		
	(0.0066)	(0.0072)	(0.0075)	(0.0104)	(0.0072)		
Dep. Var. Mean	0.5450	0.4550	0.4980	0.4720	0.5270		
County FE	✓	✓	✓	✓	✓		
State-Year FE	✓	✓	✓	✓	✓		
Individual Controls	✓	✓	✓	✓	✓		
Observations	142,277	253,584	168,268	110,144	279,457		

Notes: Estimates for the effect of high exposure to paywalls on electoral participation, among individuals predicted not to pay for a news subscription. The specification is motivated by Equation (1.3), estimated and aggregated using the Callaway-Sant'Anna estimator. he outcome is an indicator for participating in the election (CES, 2008–2020). Panel A uses self-reported voting; Panel B uses validated data from matching survey respondents to official records (Ansolabehere and Hersh, 2012). All specifications residualize the outcome on individual demographic controls – categorical indicators for sex (2 groups), age (5), ethnicity (2), education (6), employment status (6), family income (5), and children (2) – as well as county fixed effects and state-by-year fixed effects. Counties are treated once paywall exposure exceeds the 75$^{\text{th}}$ percentile. The control group consists of not-yet-treated and never-treated counties. Sample weights are included. Standard errors (in parentheses) are clustered by county. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A1.8: Effect of paywalls on participation in elections - Individuals not paying for news

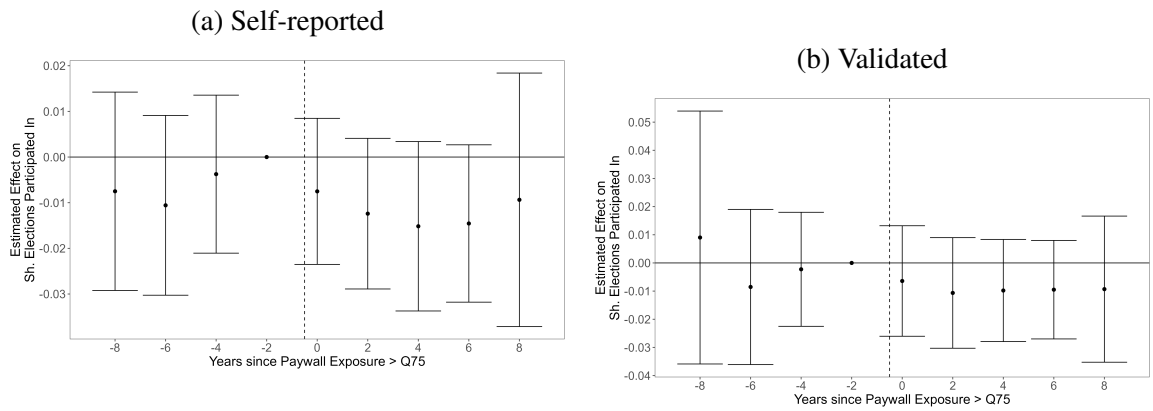

Notes: Bi-yearly estimates for the effect of high exposure to paywalls on electoral participation, among individuals predicted not to pay for a news subscription. The specification is motivated by Equation (1.3) and estimated using the Callaway-Sant'Anna estimator. The outcome is the share of elections participated in among elections held for president, house representative, senators, and governor (CES, 2008–2020), residualized on individual demographic controls, county fixed effects, and state-by-year fixed effects. Panel (a) uses self-reported voting data, while Panel (b) uses validated data from matching survey respondents to official voting records (Ansolabehere and Hersh, 2012). Counties are treated once paywall exposure exceeds the 75th percentile. The control group consists of not-yet-treated and never-treated counties. Vertical bars show bootstrapped 90 percent uniform confidence intervals. Sample weights are included. Standard errors are clustered by county.

Chapter 2

#Science: How Twitter Shapes Economic Research

Co-authored with Carlo Schwarz

Abstract

We study how Twitter usage affects the production of research in economics. We focus on 17,000 economists, for whom we collect detailed information on publications, citations, co-authors, and research topics from RePEc. We augment this data by linking economist to their personal Twitter accounts as well as PhD cohorts based on publicly available curricula vitae. To estimate causal effects, we develop a novel instrument for Twitter usage that exploits peer effects across PhD cohorts and research fields. Our results indicate that Twitter usage increases productivity by approximately one additional paper every 2.2 years for the average user. Although individual papers receive slightly fewer citations, higher overall productivity leads to an additional 8.7 citations per year. These gains are largest for intensive Twitter users, men, and more experienced economists. Key mechanisms behind these effects include expanded co-author networks and shifts toward research topics that resonate more on Twitter, such as themes related to Covid-19, cryptocurrency, and Trump. This shift in topics is associated with higher attention toward economists' work in mainstream media and the general public. Together, our findings suggest that Twitter significantly reshapes the direction of economists' work and amplifies its recognition both within and beyond academia.

Keywords: Science, Economic Research, Social Media, Information Aquisition
JEL Codes: D83, D91, L82, L86, O31

1 Introduction

Scientists require a deep knowledge of existing and ongoing research to make critical decisions about which research questions to pursue. Such choices, in turn, shape careers and influence the direction of the discipline overall. Social networks play a crucial role in this knowledge acquisition process. Researchers learn about research from their academic peers at conferences and seminars. Moreover, academic peers influence where research is published and how widely it is cited as editors and referees. In the past decade, online social media have reshaped academic networks by providing new platforms for scholars to share their work and engage in academic discourse. In economics, the social media platform Twitter has emerged as particularly influential, with the hashtag *#EconTwitter* serving as an important channel through which economists learn what others in the field are working on, and what they are interested in.

Despite Twitter's growing prominence in the profession, little is known about how it affects scientific networks and the production of knowledge. The influence of social media likely operates through several, potentially opposing, mechanisms. On one hand, Twitter may raise visibility, expand professional networks, and help researchers identify which topics are perceived as important. On the other hand, it may distract researchers, amplify incentives to follow short-term trends, or contribute to inefficient herding behavior. As the first global, accessible, real-time platform to host a substantial share of academic debate, Twitter represents a novel development in the organization of scholarly communication. Understanding its impact is thus essential to assess how digital platforms shape the direction of scientific knowledge production.

In this paper, we explore these questions by examining research in the field of economics over the past two decades. Our analysis is based on an extensive data collection effort to link 17,000 economists to their personal Twitter profiles and collect publicly available curricula vitae from personal websites. Based on these new data, we estimate the causal effect of Twitter using a novel instrumental variable strategy that allows us to overcome the endogeneity of Twitter adoption. This enables us to provide the first causal estimate of the impact of Twitter usage on economists' academic output, scientific networks, topics of research, and the visibility of their work both inside and outside academia.

For our analysis, we construct a comprehensive dataset linking economists to their academic output, Twitter activity, and early-career peer networks. Our sample consists of researchers with profiles on RePEc, a database of scholars and research papers in economics that provides detailed information on publications, citations, co-authors, and re-

search topics. We identify personal Twitter accounts through a combination of sources, including personal websites, targeted Google and Twitter searches, and external datasets, hand-validating all matches with the assistance of research assistants. This procedure allows us to recover the date of Twitter sign-up as well as the usage intensity of the platform. To capture early career networks, we collect researchers' publicly available curricula vitae from RePEc, personal websites, and external datasets to infer PhD institutions and graduation years of researchers. Finally, we use Altmetric data to measure attention to researchers' work in the broader public.

The key empirical challenge when trying to estimate the effect of Twitter usage on research is the endogeneity of the decision to join Twitter. As a result, any naive comparison between Twitter users and non-Twitter users would almost certainly be biased. To overcome this endogeneity issue, we develop a novel instrumental variables strategy that exploits plausibly exogenous variation in indirect peer pressure to sign up for Twitter based on variation in research field composition of PhD cohorts. More specifically, we instrument each researcher's Twitter usage with the indirect Twitter pressure their PhD peers are exposed to through their respective research fields.

Put differently, we leverage the fact that PhD cohorts with more peers working on financial markets, as opposed to monetary policy, experience higher Twitter adoption rates, thereby generating plausibly exogenous variation in indirect Twitter pressure on their peers. In this way, our instrument does not rely on any individual's decision to sign up for Twitter, which is endogenous, but rather only exploits variation in the field composition across PhD cohorts. Importantly, our instrument accounts for any effect of the PhD institution itself and the research field of economists. We show that our instrument strongly predicts individual Twitter usage, which is robust to various choices of construction. Moreover, the timing of adoption confirms that indeed the composition of peers drives adoption.

In our main specification, we instrument the number of years of Twitter usage with cohort-specific Twitter pressure in a two-stage least squares framework. We control for graduation year, cohort size, PhD institution, primary research field, and, in some specifications, pre-Twitter publication or citation counts. Our approach identifies the local average treatment effect of Twitter usage for researchers who joined the platform specifically due to Twitter usage among their PhD peers. The key identifying assumption is that, conditional on these controls, no omitted variable jointly affects researchers' outcomes and the Twitter pressure faced by their PhD colleagues.

Our analysis proceeds in three parts. In the first part, we show that Twitter usage increases the number of economists' publications and received citations. We estimate that Twitter users on average published one additional paper every 2.2 years. These increases

are present both for journals with low and high impact factors, albeit of smaller absolute magnitude for more prestigious journals. To examine whether the productivity gains also lead to more academic recognition, we next turn to the effects on citations. We find that the average Twitter user receives an additional 8.7 citations per year. Interestingly, these *increases* in citation mask a slight *decrease* in citations per individual paper, which is being offset by economists' higher research output. Additionally, we establish that the increases in citations stem from a higher share of citations by other researchers on Twitter.

We additionally explore conduct heterogeneity analyses based on researchers' intensity of Twitter usage, gender, and career stage, to understand to what extent the benefits of Twitter usage differ by type. We find that economists with more followers, more following accounts, and more Tweets experience larger gains in publications and citations. Moreover, although the effects of Twitter usage are positive for both men and women, men benefit significantly more. Finally, economists who received their PhD before the year 2000 see larger increases in both publications and citations, while younger economists experience increases only in publications, but not in citations.

In the second part of the paper, we investigate whether the observed increase in research productivity can be explained by changes in co-authorship and referencing behavior. This allows us to shed more light on the effects of Twitter usage on research networks and the acquisition of knowledge about existing research. First, we show that economists using Twitter expand their network of co-authors, gaining on average 0.44 additional co-authors per year. Strikingly, most of these new collaborations involve economists who are also active on Twitter. We also observe modest effects on co-author quality with increases in the number of Top 5 and Top General Interest journal publications of co-authors. Second, we document that Twitter usage leads economists to reference more papers in their average publication. These increases again are driven by references to papers authored by other Twitter users, and are distributed equally across all levels of journal quality. This suggests that Twitter shapes the set of papers economists are aware of.

Motivated by the findings in the previous part, the third part of the paper studies how Twitter usage influences economists' choice of research topics. We find that Twitter usage increases the similarity of research to other Twitter users while decreasing the similarity to research of non-Twitter users. Note that this finding holds conditional on an economist's research field. We also show that Twitter usage increases the likelihood of economists working on trending or controversial themes such as Covid-19, cryptocurrencies, or Trump. Finally, we document that this shift in research topics coincides with increased attention in the broader media landscape, measured by mentions in mainstream media, blogs, and Twitter posts based on data from Altmetric.

Taken together, our results suggest that Twitter usage increases the recognition economists receive for their work, both within academia and in the broader public. This effect appears to operate primarily through two channels. First, Twitter seems to facilitate finding co-authors, leading to higher productivity. Second, economists choose research topics that resonate more with audiences on Twitter and beyond. However, the ability to realize these benefits likely depends on researchers' success in building a visible presence on the platform. In practice, this may advantage more experienced economists who are already established in the field and better positioned to attract a large follower base.

Related literature

Our findings contribute to the existing literature in at least three ways. First, our findings contribute to the literature on the economics of science. Previous studies have explored many important drivers of innovation, such as the relevance of patenting rights (Scotchmer, 1991; Moser and Voena, 2012; Galasso and Schankerman, 2015), the geography of knowledge spillovers (Jaffe et al., 1993; Peri, 2005), and productivity gains to collaboration (Wuchty et al., 2007; Jones et al., 2008). We contribute to this field by shedding light on how modern information technologies shape the direction of scientific progress, which has so far been overlooked.

Second, we relate to the literature on the effects of social media, that has been shown to impact voting, polarization, and protests (Bond et al., 2012; Tucker et al., 2018; Zhuravskaya et al., 2020; Fujiwara et al., 2021), hate crimes (Müller and Schwarz, 2021, 2022), as well as well-being and mental health (Allcott et al., 2020; Braghieri et al., 2022). However, evidence regarding the effects on academic research is limited. Tonia et al. (2016) randomizes social media exposure of 130 articles in the International Journal of Public Health published between 2012 and 2014 and finds no effects on downloads and citations. Instead, observational studies find a positive correlation between social media coverage and subsequent citations (Eysenbach et al., 2011; Haustein et al., 2014; Smith et al., 2019; Sathianathen et al., 2020; Özkent, 2022). We contribute to this literature by providing the first large-scale causal evidence for the effects of social media on scientific output, confirming the cross-sectional correlations found in other studies. Moreover, we are the first to explore how social media affects the *direction* of academic research.

Third, we relate to previous research on information acquisition and herd behavior. Theoretical papers have explored the origins of herd behavior (Banerjee, 1992; Bikhchandani et al., 1992) and the dynamics of information acquisition in networks (Acemoglu et al., 2010, 2011). Empirical studies have provided evidence for herd behavior in a wide range of contexts, for example in online product choice of private consumers (Huang and

Chen, 2006; Babić Rosario et al., 2016) and investing decisions of professional financial investors (Scharfstein and Stein, 1990; Welch, 2000; Sias, 2004). We contribute to this literature by providing the first evidence on herd behavior among scientists.

2 Data

For our analysis, we built the largest existing database of economists, their Twitter usage, and their academic output. We identify personal Twitter accounts using a combination of information from personal websites, targeted web searches, and other existing datasets. We further augment this data with PhD institutions and graduation years by collecting CVs from personal webpages. We describe the data collection steps in the following.

2.1 Publication Data

The starting point of our data collection is RePEc (**Re**search **P**apers in **Ec**onomics), which provides a comprehensive online database of publications and author profiles in economics. RePEc is a decentralized, collaborative, non-profit initiative that aims to enhance the accessibility of research in economics and related fields. Established in 1997, RePEc aggregates research papers and metadata from over 2,000 participating archives. Publishers provide standardized, machine-readable files describing the works in their journals or working paper series, which are made accessible through various RePEc services. A great advantage of RePEc is that it allows authors to create and maintain their personal research profiles, allowing them to link their work and track related statistics such as views and citations. In this way, RePEc can avoid errors in author disambiguation of papers that occur in other publication databases.

The database is organized in text files as shown Figure 2.1(a), and publicly available for download. By parsing these files, we obtain information on around 82,000 authors who produced more than 1.3 million published and working papers in more than 3,000 journals and working paper series. For each author, we observe current affiliation(s), and, if provided, links to their personal website and Twitter account. Moreover, we obtain their full history of published and working papers. In particular, the data record title, abstract, names of all co-authors, publishing journal or working paper series, and date of publication. For a subset of papers, we also have information on keywords and Journal of Economic Literature (JEL) topic codes chosen by the authors, as well as DOIs.[1] Finally, we observe the directed graph of citations between papers as of 2022.

[1] *Digital Object Identifier (DOI)*, a unique and persistent identifier that is widely used to link digital objects across different sources.

We supplement the RePEc data using OpenAlex, a public database that tracks publications and citations of scholars across all disciplines. We link authors based on first and last name as well as paper titles. While OpenAlex is useful for validation, RePEc is better suited to our analysis as its exclusive focus on economics ensures higher data quality for this discipline.

In addition, we capture the attention paid to economists' publications in the broader public based on data from Altmetric for roughly 130,000 papers for which we observe the DOI. Altmetric records the number of mentions across various platforms such as news media, blogs, and social media. Mentions on Twitter are classified by user type, including academics, journalists, and the general public, based on keywords in users' biographies. In addition to the number of individual mentions, Altmetric reports a composite attention score, with weights that reflect the relative importance of the source (e.g., mentions in the news receive a larger weight than social media posts).

Figure 2.1: Example Data

(a) RePEc Data

```
Name-First: Joshua
Name-Middle: D
Name-Last: Angrist
Name-Full: Joshua Angrist
Workplace-Name: Massachusetts Institute of Technology (MIT)
/ Economics Department
Workplace-Institution: RePEc:edi:edmitus
Workplace-Share: 60
Workplace-Name: National Bureau of Economic Research (NBER)
Workplace-Institution: RePEc:edi:nberrus
Workplace-Share: 40
Email: angrist@mit.edu
Homepage: http://web.mit.edu/angrist/www/
```

(b) Twitter Data

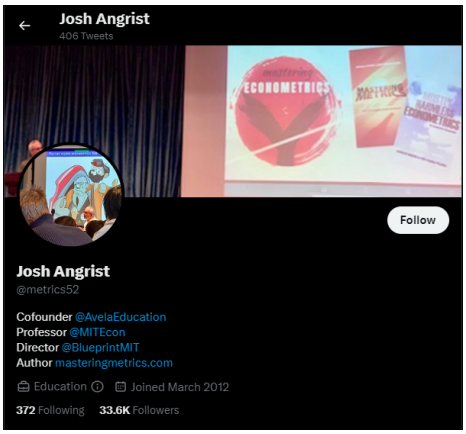

Notes: Panel (a) presents an example RePEc database file. Panel (b) shows a screenshot of a Twitter profile.

2.2 Twitter profiles

As the second step of our data collection, we identify Twitter accounts of economists using a five-step procedure: 1) RePEc, 2) information from professional websites of economists, 3) Google searches, 4) Twitter profile search, and 5) information on Twitter profiles from Mongeon et al. (2023). Afterwards, all potentially matching Twitter accounts were hand-validated with the help of research assistants. The overall procedure is visualized Figure 2.2 and described in the following. Also note that the data collection in large part took place before Elon Musk's takeover of Twitter in October 2022, and we therefore also observe Twitter profiles of users who might have deleted their accounts afterwards.

Figure 2.2: Data collection process for Twitter profiles and PhD information

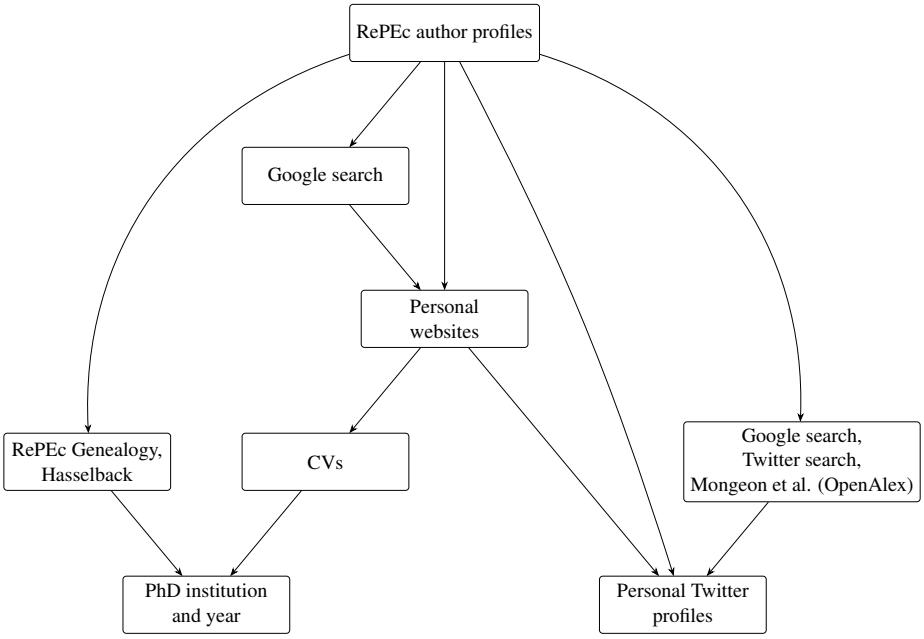

Notes: Schematic visualization of data collection process for identifying researchers' PhD information and personal Twitter profiles. Nodes represent datasets, and edges indicate usage or combination of data.

First, we utilize all Twitter profiles registered by authors on their RePEc profiles. This provides us with Twitter profiles for 3,256 economists in the RePEc data. The full list is available here: RePEc-registered economists on Twitter.

Second, we extract links to Twitter profiles from personal websites. In addition to websites linked on researchers' RePEc profiles, we retrieve additional websites using Google's Custom Search API. For each researcher, we run queries comprised of their first name, last name, and the keyword `economist`. We retain URLs from domains commonly used for personal pages (Google Sites, Wordpress, GitHub, etc.) or domains similar to the researcher's name. Research assistants manually validate that the identified website indeed belongs to the correct economists from the RePEc data. Afterwards, we use a custom scraper to visit all pages and subpages on the identified sites and extract embedded links to Twitter user profiles.

Third, we run additional Google's Custom Search API queries of the form first name, last name, while restricting the search to Twitter.com, and retain all found Twitter user profiles. This provides us with a large set of additional candidate accounts.

Fourth, we further augment the set of candidate accounts by using the Twitter Graph API, which was still active at the time of data collection, to directly search for profiles based on the first and last name of each economist. In the next step, we filter the set of

candidate Twitter accounts we have obtained in steps three and four for name similarity between RePEc and Twitter profiles, as well as keywords related to economics and academic positions in the biography (short description) of these Twitter profiles.

Fifth, we incorporate a recent dataset by Mongeon et al. (2023), which links OpenAlex authors to Twitter profiles. We match around 6,500 of these authors with Twitter profiles to our RePEc data. [2]

As a final step, we manually validate all matched Twitter accounts with the help of research assistants. The overall matching procedure is designed to minimize false positives in the matching between Twitter accounts and RePEc profiles: In addition to a similar name on the RePEc and Twitter profiles, we require clear evidence that the Twitter account belongs to the same person, based on the RePEc profile or personal website. In particular, we instructed research assistants to look for a matching profile photo, a link between the Twitter account and the personal website, or a mention of the researcher's current academic position in the Twitter biography. We ensure that no Twitter account is linked to more than one author. For researchers with multiple accounts, we retain all of them.

In total, we identify around 12,500 personal Twitter accounts among the economists in our sample. For each account, we collect the creation date, the total number of tweets, likes, and retweets, follower and following counts, and the full history of tweet content and timing. Figure 2.1 shows an example for the RePEc author file and the corresponding Twitter profile for Nobel Memorial Prize Winner Joshua Angrist.

2.3 PhD institution and graduation year

As the last step of our data collection, we obtain information on economists' PhD institutions and graduation years to identify peers during doctoral studies, which we will use in our identification strategy. We collect information based on three different sources.

First, we use data from the RePEc Genealogy project, which records PhD institutions and graduation years for a subset of authors with RePEc profiles. The database currently includes around 24,000 economists, even though only around 15,000 entries contain both variables. For example, if someone indicated a RePEc user as their advisor in the RePEc Genealogy Project, the advisor's profile will be part of the Genealogy Project, but the advisor's graduation date is not available.

[2] Much of our collection predates the 2023 API restrictions that followed Twitter's change in ownership. For a small number of accounts identified afterward, we supplement profile data from Garg and Fetzer (2025), who generously shared it with us.

Second, we incorporate information on PhD institutions and graduation years of around 3,000 US-based economics faculty from J. R. Hasselback's Faculty Directories.[3] The information was hand-collected by James R. Hasselback in five cohorts 1999-2000, 2001-2002, 2003-2004, 2006-2007, 2020-2021, and contains information on economists in 130 U.S. Economics Departments, even though not all of them list PhD institutions and graduation years.

Third, we collect publicly available CVs from the personal websites identified either from the RePEc data or our Google searches in our collection of Twitter profiles. Using a custom webscraper, we loop over all pages and subpages of economists' professional websites and identify links to CVs based on a set of keywords commonly associated with academic CVs (e.g., CV, curriculum vitae). In cases where the automated collection fails, research assistants were instructed to locate and retrieve the files manually if available. Afterwards, we parse the documents using a large language model to extract relevant career information, including educational history and academic positions, along with associated dates. To avoid hallucinations by the large language model, we use a reverse string matching procedure to ensure that the extracted PhD information is indeed contained in the CVs. For researchers who do not post a CV but provide career details in plain text on their website, research assistants record the information by hand. This procedure yields the CV for more than 9,000 economists, among them 7,700 in PDF format.

We validate the consistency of the collected information by cross-checking the records for researchers who appear in multiple sources. Moreover, we ensure that names match across data sources and that each CV is linked to a single researcher. In total, we recover both the PhD institution and the graduation year for more than 17,000 economists. To the best of our knowledge, this provides us with the largest existing database of economists and their PhD institutions.[4]

[3] We thank Advani et al. (2025) for providing the data matched to RePEc authors.
[4] For comparison, Koffi et al. (2024) hand-collect CVs for around 4,500 economists.

2.4 Outcome measures

For the empirical analysis, we create outcome measures that capture economists' output, co-authorship, citation network, as well as research topics.

Scientific output

The first set of outcome variables captures standard measures of scientific output. For each economist, we count the total number of publications and separately by journal quality. We distinguish between the commonly used "Top 5" (QJE, AER, JPE, ECMA, RESTUD) and "Top General Interest" journals[5], as well as the top 50 and top 100 journals based on RePEc's aggregate impact ranking. Moreover, we measure citations accrued by 2022 for papers published in each year. As above, we count both total citations and those coming from top-ranked journals. These measures aim to provide a differentiated account of the productivity and scientific impact of a researcher.

Scientific network

In addition, we compile outcomes that capture connections between researchers. For each author, we measure the number of unique co-authors and the number of references. By combining these variables with information on when researchers joined Twitter, we can assess each researcher's connection to the Twitter network. Specifically, we measure the share of references to papers with at least one Twitter-using author, and the share of Twitter users among (unique) co-authors.

Research topics

Further, we characterize the topics economists are working on in three ways. First, we use papers' JEL codes that jointly categorize the fields of each paper. We use the 20 one-character JEL codes (A-Z) as well as the 190 two-character JEL codes (A1-Z3). Second, we search for keywords related to "hot" topics in papers' abstracts, such as *Covid-19*, *Crypto*, or *Gender*. Third, we characterize the similarity of research topics by quantifying the similarity of abstracts using document embeddings. We construct document embeddings generated by Doc2Vec (Le and Mikolov, 2014) to represent each abstract as a

[5]The "Top General Interest" journals include the Top 5 plus the American Economic Journal (Applied Economics, Macroeconomics, Microeconomics, and Economic Policy), Economic Journal, Journal of the European Economic Association, and Review of Economics and Statistics.

50-dimensional vector, and quantify the closeness between abstracts using cosine similarity. Then, we compute the average similarity of each abstract to the 100 most similar abstracts published in the same year, among papers published by authors who use Twitter and authors who do not, respectively. These measures reveal how similar the research by economists on Twitter is to other economists on Twitter, relative to research by economists outside the Twitter network.

2.5 Sample Construction and Descriptive Statistics

Our analysis is based on the economists in the RePec data with at least one journal publication between 2006 and 2022. We restrict our analysis to authors with at least one publication in this time frame to avoid using information from retired or dead economists. The time frame exactly covers the launch of Twitter in March 2006 until the takeover of the platform by Elon Musk in October 2022, which was followed by large platform changes and an exodus of many economists from the platform to BlueSky and other networks. Hence, we capture almost the entire life cycle of the social media platform and its use by economists.

For our instrumental variable strategy, which we describe in the next section, we additionally need to restrict our analysis to economists for whom we have information on the PhD institution and year. This allows us to use economists' academic network of researchers to identify causal effects. This restriction also helps us to define a precise start year to include researchers in our data. Economists enter our sample in their graduation year, and leave it after their last observed paper.

Table 2.1 summarizes characteristics of the final sample. For comparison, Column (1) reports numbers for the full sample of RePEc economists, Column (2) includes only researchers active between 1995 and 2022, and Column (3) further restricts to researchers with information on PhD institution and year. The final sample contains 17,202 economists who graduated from more than 1,000 PhD institutions and produced almost 600,000 published and working papers that received close to 12 million citations. Among these researchers, we have found their personal website for 87.4 percent, and the full CV in PDF format for 34.5 percent. Moreover, at least 29.9 percent of these researchers have signed up for Twitter.[6]

[6]Successful researchers are more likely to join Twitter, as shown in Figure A2.1, which plots the share of Twitter users over time by quartiles of the number of citations in the previous three years.

Table 2.1: Sample composition

	Sample		
	All RePEc	**1995-2022**	**Final**
	(1)	(2)	(3)
Panel A: Share of researchers linked to:			
Website	0.602	0.678	0.874
CV	0.093	0.112	0.345
RePEc Genealogy	0.214	0.258	0.773
OpenAlex	0.744	0.922	0.973
PhD Institution & Year	0.231	0.277	1.000
Twitter Account*	0.148	0.184	0.299
Panel B: Number of unique observations:			
Citations	22,273,206	18,266,839	11,870,391
Papers	1,331,566	1,153,468	573,869
Researchers	82,183	65,352	17,202
PhD Institutions	1,088	1,057	1,021

Notes: This Table summarizes how we construct our final estimation sample. Column (1) includes all economists with a profile on RePEc. Column (2) restricts to economists with at least one observed paper between 1995 and 2022. Column (3) further restricts to economists for which we observe both PhD institution and graduation year. The share of users linked to their Twitter accounts is representative only in Column (3) because some data collection steps conditioned on having full PhD information.

For our analysis, we create a cross-section of economists measuring their average outcomes over our observation period. Table 2.2 presents summary statistics. The average researcher publishes roughly one paper and receives around 33 citations per year. Less than 1 out of 2 papers are published in top 100 journals, and 1 in 25 are published in the Top 5. Around 1 in 6 citations originate from a paper published in a Top 100 journal, and less than 1 in 200 come from a Top 5 paper. Researchers reference around 39 papers in a typical year. Note that we include citations and references that involve researchers outside our sample, which is why the two numbers need not be equal. Over the entire sample period, only a small fraction of references point to papers authored by at least one Twitter user. We count around 1.5 unique co-authors per year, again with a small minority being Twitter users when considering all years. Altmetric records slightly less than one mention of researchers' work per year, most of them on Twitter. Four out of five of these mentions originate from members of the general public as opposed to scientists, according to Altmetric's estimate.

Table 2.2: Descriptive Statistics

	Mean	SD	Min	Max	N
General					
Year	2012.81	4.93	1995	2022	17,202
PhD graduation year	2004.09	11.93	1942	2022	17,202
Twitter sign-up year	2014.27	4.00	2006	2023	5,149
Years on Twitter	1.86	3.56	0	16	17,202
Female	0.22	0.42	0	1	14,026
Publications					
All	1.00	1.03	0	29	17,202
Top 100	0.38	0.44	0	7	17,202
Top General Interest	0.07	0.16	0	4	17,202
Top 5	0.04	0.12	0	3	17,202
With JEL code	0.31	0.46	0	15	17,202
With DOI	0.13	0.31	0	6	17,202
Citations					
All	32.71	81.50	0	2117	17,202
Top 100	5.44	13.91	0	438	17,202
Top General Interest	0.23	0.74	0	17	17,202
Top 5	0.14	0.48	0	12	17,202
Network					
References	39.00	49.90	0	921	17,202
Sh. Twitter users among references	0.10	0.11	0	1	16,711
Unique co-authors	1.56	1.71	0	34	17,202
Sh. Twitter users among co-authors	0.17	0.24	0	1	15,831
Altmetric					
All	0.88	7.97	0	426	17,202
News media	0.05	0.80	0	62	17,202
Social media	0.63	6.15	0	423	17,202
Twitter	0.62	6.01	0	423	17,202
Twitter, by general public	0.47	5.19	0	421	17,202
Twitter, by scientist	0.12	1.06	0	87	17,202
Aggregate Attention Score	1.03	9.77	0	610	17,202

Notes: Summary statistics at the scientist level, representing the average value across all years observed in the sample.

3 Empirical Strategy

The key challenge in estimating the causal effect of Twitter usage on the research of economists lies in the endogenous nature of the decision to sign up for Twitter. In other words, researchers may adopt Twitter for reasons that are correlated with outcomes but nearly impossible to control for. Due to this endogeneity, any naive regression of research outcomes on Twitter usage would almost certainly be biased. On the one hand, the estimates could be biased upward if more productive researchers signed up for Twitter. On the other hand, the estimates could be biased downwards if researchers signed up for Twitter at points in time when they move away from research to a more policy-focused career.

Due to this endogeneity, the effect of Twitter usage has proven challenging to investigate empirically. We propose to estimate the causal impact of Twitter usage based on a novel instrumental variables strategy exploiting plausibly exogenous variation in *indirect* Twitter pressure through economists' peer network from their doctoral studies. We describe the construction and plausibility of our instrument in the following section and present the first stage evidence.

3.1 Instrument for Twitter Usage based on PhD Network and Academic Field

Our strategy is based on the idea that the utility from using social media platforms increases if more peers join the same platform. Indeed, existing research shows that individuals' decisions to join social media platforms strongly correlate with their peers' usage (e.g., Aridor et al., 2024). In the context of academic research, two particularly relevant networks are 1) classmates from the same PhD cohorts and 2) researchers in the same subfields of economics, whose collective interests and evaluations ultimately determine the academic impact and recognition of one's research.

Our instrument exploits a combination of two variations in these two key networks. First, building on the long-standing literature on peer effects (e.g., Manski, 1993; Goldsmith-Pinkham and Imbens, 2013; Angrist, 2014), we exploit the fact that the characteristics of the members of one's PhD cohort can influence one's decisions. Second, we exploit the network structure of Twitter for identification using the indirect effect of "peers of peers" (e.g., Bramoullé et al., 2009; Lin, 2010; De Giorgi et al., 2010). In combination, these sources of variation allow us to instrument for economists' own Twitter usage with the field-specific pressure to use Twitter faced by their PhD colleagues.

Figure 2.3: Motivational Evidence: Twitter Usage by Individual and in Network

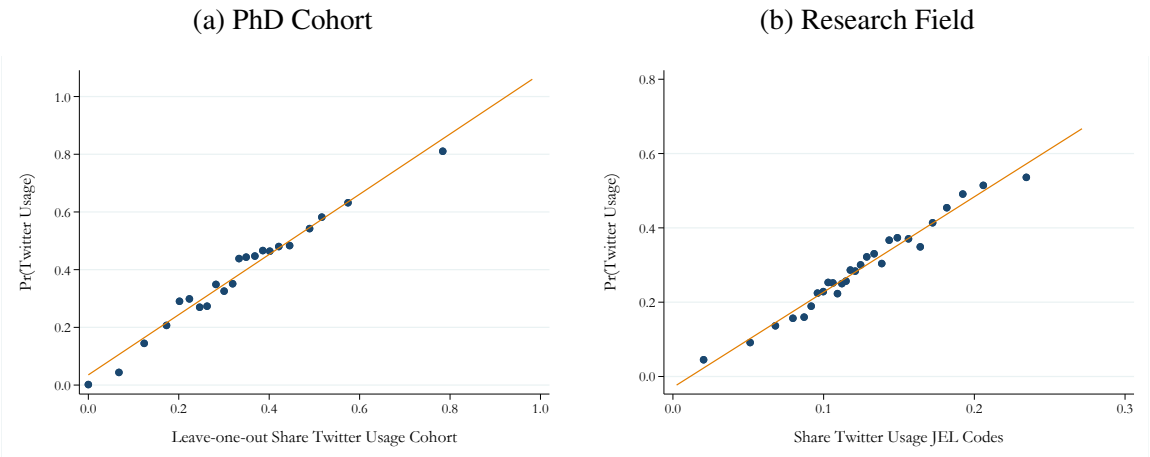

Notes: Binned scatter plots on researcher level correlating an indicator for ever using Twitter and the (leave-one-out) share of Twitter users in the PhD cohort (Panel (a)) and among the researchers with the same most frequent two-character JEL code (Panel (b)).

Using our data, we provide motivational evidence for the fact that both the Twitter usage in one's PhD cohort as well as one's research field are good predictors of an economist's own decision to join Twitter. Figure 2.3 Panel (a) presents a binned scatterplot of researchers' Twitter usage against the leave-one-out share of Twitter users in their PhD cohort (excluding the researcher themself). In Panel (b), we show a binned scatterplot of researchers' Twitter usage against the leave-one-out share of Twitter users in their respective research field. In both cases, we find a strong and positive correlation between an economist's own decision to join Twitter and the adoption among peers.

Building on this motivational evidence, we propose an empirical strategy in the spirit of Bramoullé et al. (2020), and instrument each researcher's Twitter usage with the Twitter usage of "peers of peers", defined as researchers who are two network links away. This strategy has several key advantages. First, by constructing the instrument without relying on the direct Twitter adoption decisions of either the researcher or their immediate colleagues, we avoid issues related to the mutual influence of peers on each other that are at the heart of the reflection problem (Manski, 1993). Second, as we use economists' early-career networks formed during doctoral studies – which are not directly chosen by a researcher and remain fixed throughout their career – we preclude any endogenous changes in network links and the composition of PhD colleagues' fields. Third, as we only use variation in Twitter usage in fields that an economist does not work in themself, we overcome concerns about correlated shocks that could drive Twitter adoption as well as output.

The construction of our instrument proceeds in three steps. As a first step, we determine PhD peers using information on researchers' doctoral institutions and graduation years. In our baseline measure, we construct a PhD cohort includes all scholars who graduated from

the same institution within five years before or after the researcher's own graduation. This choice is motivated by the average 6-year length of economic PhD programs and therefore should cover all other PhDs economists potentially interact with during their PhDs. As we show in robustness tests, later this choice is not material to any of our findings.

As a second step, we identify the research fields of all economists and their PhD peers based on the JEL codes indicated in their research papers. In our baseline approach, we classify each economist's main fields based on the three most frequently used two-character JEL codes (i.e., A1–Z3, 190 distinct codes) among their research papers in their entire career. By doing so, we assume that economists' main fields are stable over time and thus unaffected by career stage or Twitter adoption. We consider this assumption justified for two reasons. First, economists typically specialize in one or two core fields and work on them throughout their careers. Second, by drawing on authors' work over all available years, we capture long-run field affiliations rather than transient research interests. In robustness checks, we again confirm that our instrument is robust to various alternative ways of defining the research fields of economists, employing among others the more aggregated one-letter JEL codes (A–Z, 20 codes).

As a third step, we construct our instrument, capturing each researcher's indirect Twitter Pressure as the average share of Twitter usage in the fields of their PhD colleagues (excluding the researcher's own fields). To facilitate exposition, we present the formal construction of the instrument for the simplified case in which each researcher contributes to exactly one field. The general version, which accounts for researchers contributing to multiple fields with varying intensity, is detailed in Appendix A.2 In the simplified case, indirect Twitter Pressure for researcher i in year t is defined as:

$$Twitter\ Pressure_{i,t} = \frac{1}{|Cohort_i| - 1} \sum_{\substack{c \in Cohort_i, \\ c \neq i}} \sum_{\substack{f \in Fields, \\ f \neq Field_i}} I(Field_c = f) \times Sh.\ Twitter\ Users\ in\ Field_t^f \tag{2.1}$$

The summation runs over the set of peers in economists i's cohort ($\sum_{\substack{c \in Cohort_i, \\ c \neq i}}$) and over the set of research fields excluding i's own one ($\sum_{\substack{f \in Fields, \\ f \neq Field_i}}$). The indicator $I(Field_c = f)$ equals one if colleague c is active in field f, and zero otherwise. $Sh.\ Twitter\ Users\ in\ Field_t^f$ denotes the share of researchers in field f who have a Twitter account in year t.[7] In other words, the resulting instrument captures the average share of Twitter usage in fields represented among economists' PhD peers, excluding their own area of specialization.

[7] We compute the share of Twitter users in each field as a leave-one-out measure, excluding colleague c from the calculation.

3.2 Estimation

To estimate the causal effect of Twitter usage, we make use of our measure of indirect Twitter pressure in a two-stage least squares (2SLS) specification, in which we instrument for the number of years that a researcher has been on Twitter. More specifically, we estimate the following cross-sectional regression:

$$y_i = \alpha + \beta \widehat{Years\ on\ Twitter}_i + \mathbf{X}'\theta + \varepsilon_i \ . \tag{2.2}$$

where y_i denotes an outcome variable for researcher i (e.g., academic output, citations, or research topics) in the period 2006 to 2022. The primary variable of interest is $Years\ on\ Twitter_i$ and captures the number of years an economist spent on Twitter until 2022.[8] In our regression, we additionally account for a large set of observable controls $\mathbf{X}'$. In particular, we control for PhD year fixed effects to absorb career stage and cohort-level productivity shocks, the number of peers in one's PhD cohort, PhD institution fixed effects to account for institutional characteristics, and field fixed effects to capture differences across research areas.

In the first stage, we instrument $Years\ on\ Twitter_{\ i}$ using average Twitter Pressure individual i faced between 2006 and 2022. This gives rise to the following first-stage equation:

$$Years\ on\ Twitter_i = \omega + \lambda\, Twitter\ Pressure_i + \mathbf{X}'\gamma + \eta_i \tag{2.3}$$

where $Twitter\ Pressure_i$ is the average of researcher i's Twitter Pressure aggregated over all years and $\mathbf{X}'$ is the above described set of control variables. In both specifications, we cluster standard errors at the PhD institution $\times$ PhD year level, which is the relevant level of variation of our instrument.

We estimate our regressions using a cross-sectional data structure as it circumvents at least two econometric issues that would arise in a panel setting. First and most importantly, our instrumental variable is both continuous and staggered, with repeated treatments, and to the best of our knowledge, no estimators exist for a staggered-IV-difference-in-differences design. Second, in a panel data structure, we would need to account for the autocorrelation of outcomes as well as the dynamic dependencies of treatment exposure explicitly, using lags of treatment and outcome variables. The cross-sectional data structure allows us to overcome both issues simultaneously, providing a transparent research

[8] In alternative specifications, we also use an indicator for Twitter adoption, which captures the extensive instead of intensive margin of using Twitter.

design. By using the $Years\ on\ Twitter_i$ as our main treatment variable, we preserve variation in adoption timing. In addition, we will provide evidence for the timing of our treatment effects using a reduced-form staggered difference-in-differences regression.

3.3 First Stage Evidence

We start by confirming the relevance of our instrumental variable strategy using a binned scatter plot in Figure 2.4. We find a strong positive association between the amount of Higher Twitter Pressure (x-axis) and the likelihood of an economist joining Twitter (y-axis). The Figure also confirms an approximate linear relationship between the amount of Twitter pressure and the likelihood of Twitter usage.

Figure 2.4: First Stage: Network Twitter Pressare and Twitter Adoption

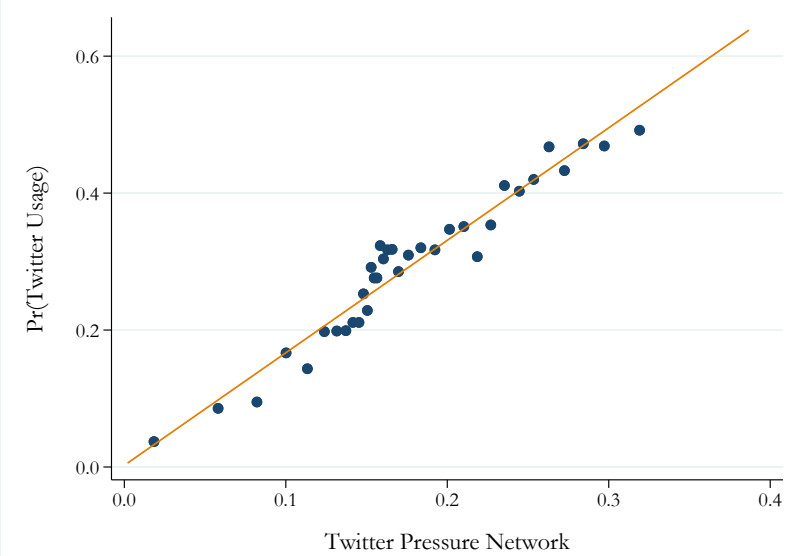

Notes: Binned scatter plot on researcher level correlating an indicator for ever using Twitter with Twitter Pressure. PhD cohorts include researchers graduating from the same institution within 5 years, and fields for each researcher are defined by the three most frequent two-character JEL codes indicated on their papers.

Next, we more formally investigate the strength of the relationship between the years of Twitter usage and our instrument in Table 2.3. To ease interpretation, we standardize the Twitter Pressure instrument to have a mean of 0 and a standard deviation of 1. The estimates confirm a strong and positive relationship between our instrument and the years of Twitter usage. We find that, on average, a one standard deviation increase in the amount of Twitter Pressure is associated with a 0.8-year increase in Twitter usage. This increase is substantial given that the average years of Twitter usage, including non-users, is around 1.9.

The first stage coefficient is also remarkably stable to the inclusion of additional control variables. In Column (1), we only control for the number of years since the completion of the PhD in our data using fixed effects. In Column (2), we additionally include indicators for the number of students in the same PhD cohort. In Column (3), we include over 1,000 fixed effects for each PhD institution. These fixed effects account for any difference in Twitter usage that can be explained by students from a particular university being more likely to join Twitter in general. Lastly, in Column (4), we include controls for the JEL codes economists are working in. These controls will account for any effect on Twitter usage that is driven by higher Twitter pressure in a particular research field. None of these controls makes any difference for the magnitude or the significance of our first-stage estimates. If anything, the estimates increase slightly in size in comparison to Column (1).

Table 2.3: First Stage: Network Twitter Pressure and Twitter Adoption

	Dep. Var.: Years of Twitter Usage			
	(1)	(2)	(3)	(4)
Twitter Pressure (std.)	0.784***	0.936***	0.903***	0.890***
	(0.028)	(0.028)	(0.031)	(0.031)
Observations	15,417	15,417	15,417	15,417
Mean of DV	1.94	1.94	1.94	1.94
R^2	0.03	0.03	0.03	0.03
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: Regressions of the number of years of Twitter usage on standardized Twitter Pressure, including fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Intensive vs. Extensive Margin

We also investigate if our instrument predicts both the extensive and the intensive margin of Twitter usage. Regarding the former, we replace the dependent variable in the first-stage regressions with an indicator that is equal to 1 if an economist uses Twitter. Similar to our baseline estimates, we find a strong and highly robust positive association between our instrumental variable and the decision to join Twitter (see Table 2.4). On average, a one standard deviation increase in Twitter pressure increases the probability that an economist joins Twitter by between 11.7 and 14.1 percentage points. Even the lower bound represents around a 40% increase relative to the mean.

Table 2.4: Robustness First Stage: Extensive Margin

	Dep. Var.: I[Twitter Usage]			
	(1)	(2)	(3)	(4)
Twitter Pressure (std.)	0.117***	0.141***	0.133***	0.131***
	(0.004)	(0.004)	(0.004)	(0.004)
Observations	15,417	15,417	15,417	15,417
Mean of DV	0.29	0.29	0.29	0.29
R^2	0.04	0.04	0.04	0.04
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: Regressions of an indicator for ever using Twitter on standardized Twitter Pressure, including fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

To study the intensive margin, we restrict our estimation sample to the economists who have joined Twitter and, among these, investigate the relationship between Twitter pressure and the years of Twitter usage (see Table 2.5). We find that even among Twitter users, higher Twitter pressure is associated with a higher number of years on Twitter. In other words, individuals who face higher Twitter pressure join Twitter earlier. A one-standard-deviation increase in Twitter pressure is associated with approximately one additional year of Twitter usage.

Table 2.5: Robustness First Stage: Intensive Margin

	Dep. Var.: Years Twitter Usage			
	(1)	(2)	(3)	(4)
Twitter Pressure (std.)	0.585***	0.930***	1.133***	1.157***
	(0.070)	(0.092)	(0.106)	(0.106)
Observations	4,462	4,459	4,308	4,308
Mean of DV	6.69	6.69	6.70	6.70
R^2	0.01	0.02	0.02	0.03
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: Regressions of the number of years of Twitter usage on standardized Twitter Pressure within the subsample of economists who have ever used Twitter. The specifications include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Robustness

We probe the robustness of our first-stage results with regard to several different margins. First, in Appendix Table A2.1 we show that our instrument is robust to different definitions of PhD cohorts. In particular, we vary the size of the cohort window we consider for the instrument construction from ± 6 to ± 1 years around the researcher's graduation. This perturbation has no impact on any of our findings. Second, in Appendix Table A2.2, we confirm the robustness of our findings to alternative definitions of researchers' fields. The relationship remains positive and significant for both one- and two-character JEL codes, and independent of the number of fields we consider per researcher. Third, Appendix Table A2.3 shows that the results are robust to clustering standard errors by PhD cohort, institution, or graduation year. Fourth, Appendix Table A2.4 shows that the instrument is robust to using alternative functional forms, including using the percentiles, inverse hyperbolic sine-transformed, and discretized variants split at the 50^{th}, 75^{th}, or 90^{th} percentile. Fifth, we show that our instrument holds in the sample of economists who completed their PhD before 2006 (see Appendix Table A2.5) and who are therefore consistently observed in all years of our sample period.

3.4 Exclusion Restriction

As with any instrumental variable strategy, we require the exclusion restriction. In our particular case, the indirect Twitter pressure through the PhD cohort can only affect academic outcomes via its impact on Twitter usage. While this assumption is inherently untestable, it strikes us as highly plausible in our context for at least three main reasons.

First, given that we are using the indirect effect of Twitter pressure in other research fields, we do not use information on the extent of Twitter usage in economists' own PhD cohort or in the research field. Hence, it is highly unlikely that economists can directly influence their exposure to the instrument. This assumption is particularly credible for economist who completed their PhD before 2006 and for whom the Twitter pressure therefore only materializes after they had already begun their work at a different institution. As we have shown previously, our instrumental variable also holds for this subsample.

Second, due to the large number of controls, it appears unlikely that an omitted variable jointly affects Twitter Pressure and researchers' outcomes. For example, our most restrictive specification includes fixed effects for both PhD institutions and research fields of each economist. Therefore, our instrument only relies on residual variation from the differential composition of PhD cohorts within the same PhD institution and research field.

Third, we can investigate the timing of the impact of our instrument relative to the decision of an individual to adopt Twitter. To do so, we estimate a staggered difference-in-difference specification in which researchers are treated once Twitter Pressure exceeds the 25th percentile. The estimates, obtained using the robust estimator by De Chaisemartin and D'Haultfoeuille (2022), are shown in Figure 2.5. We find no evidence for the differential adoption of Twitter in the pre-period, but significant increases in the likelihood of joining Twitter once Twitter pressure rises beyond the 25th percentile. This provides strong evidence for the idea that individuals would not have joined Twitter in the absence of Twitter pressure.

Based on these multiple pieces of evidence, we believe that it is credible that our instrument provides plausibly exogenous variation in Twitter usage, which we will exploit in the following to estimate the effect of Twitter usage on the research of economists.

Figure 2.5: Dynamic Effect of Twitter Pressure on Twitter Adoption

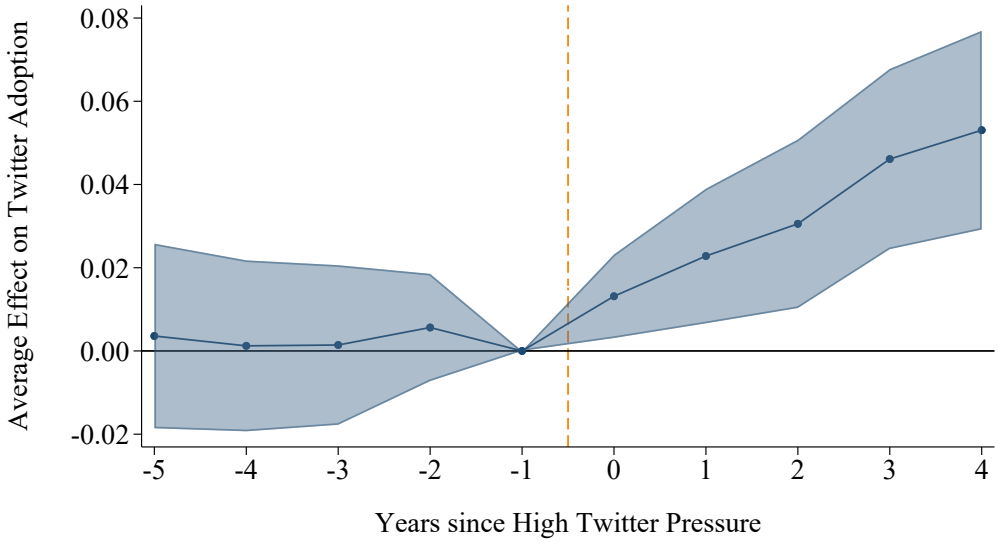

Notes: Point estimates and 95 percent confidence intervals for yearly effect of high Twitter Pressure on indicator for using Twitter in year, estimated using the heterogeneity-robust estimator by De Chaisemartin and d'Haultfoeuille (2020) in author-year panel. High Twitter Pressure is an indicator switching to one once Twitter Pressure ($\pm$5-year cohort, Top 3 JEL2 codes) exceeds the 25th percentile.

4 Effect of Twitter on Scientific Output

In the first part, we start by investigating the effect of Twitter usage on the number of publications of economists, both in low- and high-ranked journals. We further study the impact on the number of citations received by economists. Additionally, we analyze heterogeneity conditional on gender, age, and the intensity of Twitter usage.

4.1 Publications

We start by estimating Equation (2.2), where the dependent variable is the average number of publications of an economist in all journals covered by the RePEc data. Table 2.6 presents the effect of Twitter usage on the number of yearly publications, using Twitter Pressure as an instrument for the number of years on Twitter. As for the first stage estimates, Column (1) includes fixed effects for years since the completion of the PhD, Column (2) adds fixed effects for the the number of students in the $\pm$5-year PhD cohorts used to construct the instrument, Column (3) further includes fixed effects for the PhD institutions, and Column (4) finally controls for each researcher's JEL code shares.

We find that Twitter usage substantially increases economists' publication output. The estimates are positive and significant across all specifications, and the magnitude increases slightly with the inclusion of additional control variables. Given that the average number of years of Twitter usage (including non-users) is 1.9, the estimate of 0.239 in the most restrictive specification suggests that Twitter users on average published one additional paper every 2.2 years.

Table 2.6: Effect of Twitter Usage on Number of Publications

	Dep. Var.: Nr. of Publications per Year			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.147***	0.214***	0.236***	0.239***
	(0.016)	(0.013)	(0.015)	(0.015)
Observations	15,417	15,417	15,417	15,417
Mean of DV	1.07	1.07	1.07	1.07
F-Stat (KP)	775.82	1082.18	857.43	824.05
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of publications per year. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

While the increase in publications is very stable across specifications, an obvious question regards the quality of these publications, as the publication of Twitter users could either appear in high-quality journals with low acceptance rates or lower-ranked journals with high acceptance rates. We investigate this question by replacing the dependent variable with the number of publications in journals of varying quality. More specifically, we consider publication in the Top 100, Top 50, Top General Interest, and Top 5 journals.[9]

We visualize the estimates from these regressions in Figure 2.6 and report the precise regression estimates in Appendix Table A2.6. Interestingly, we find that Twitter usage is associated with an increase in the number of publications across all journal categories. The coefficients are decreasing with increasing journal quality: 0.239 (all publications), 0.091 (Top 100), 0.066 (Top 50), 0.02 (Top General interest and 0.013 (Top 5). In other words, the average Twitter usage is associated with around one additional Top 5 publication over the course of a 40-year career.

Figure 2.6: Heterogeneity by Publication Quality

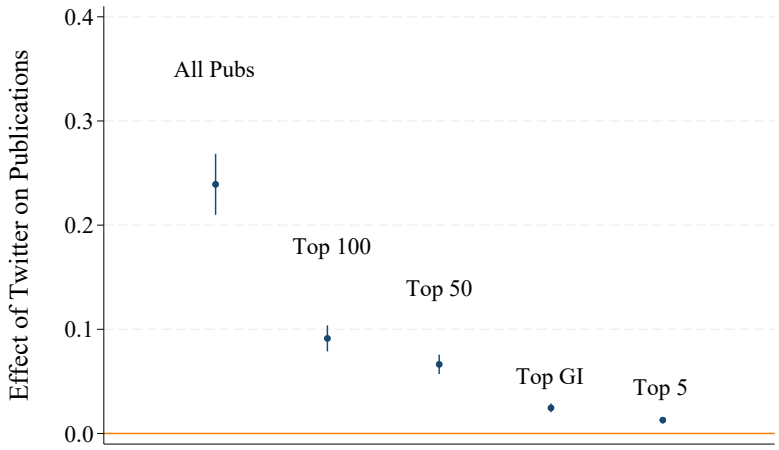

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications in different brackets of journal impact factor. The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

One concern with the previously presented estimates could be that our instrument is correlated with the pre-existing quality of economists. To address this concern, we can control for the average number of publications of economists in the period before the launch of Twitter. More specifically, we control non-parametrically for pre-period publication levels by including an indicator variable for each value of the average number of publications. Note that this approach cuts our sample by more than half, as we can only estimate the specification for people who finished their PhD before 2006.

[9]We discuss the precise definition of these categories in Section 2.4.

The results are reported in Appendix Table A2.7. We find that the inclusion of this flexible control for pre-period productivity has hardly any bearing on our results. The estimate in Column (4) of 0.193 is very close to our original estimate of 0.239, and across all specifications, the estimates continue to be large, positive, and statistically significant. This finding suggests that our instrument captures variation in Twitter usage that is unrelated to the inherent productivity of economists.

Taken together, these results suggest that Twitter usage significantly increased economists' research output, both in terms of the total number of publications and in the ability to place their work in high-quality journals.

4.2 Citations

We next turn our attention to the number of citations economists receive to understand the impact of published papers and their recognition in the academic community. We find that Twitter usage is also associated with an increase in the average number of citations of economists (see Table 2.7). The estimates are again remarkably stable, independent of the included control variables. The magnitude in Column (4) implies that the average Twitter user receives an additional 8.7 ($4.615 \cdot 1.9$) citations per year.

Table 2.7: Effect of Twitter Usage on Number of Citations

	Dep. Var.: Nr. of Citations per Year			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	7.443***	5.454***	4.879***	4.616***
	(0.640)	(0.623)	(0.683)	(0.689)
Observations	15,417	15,417	15,417	15,417
Mean of DV	29.95	29.95	29.95	29.95
F-Stat (KP)	775.82	1082.18	857.43	824.05
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of citations received per year. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

To rule out that this result is driven by pre-existing citation differences, we again repeat this analysis while non-parametrically controlling for the number of pre-Twitter citations among researchers that were active before 2006 (see Table A2.8). Our findings are again robust to the inclusion of these highly flexible controls, and the results are broadly similar, suggesting on average 5.4 ($2.855 \cdot 1.9$) additional citations per year for Twitter users.

The increase in citations could stem from two different margins. First, as we have shown in the previous part, Twitter users publish more papers than non-Twitter users. Second, each individual publication of Twitter users could receive more citations. To differentiate between these two margins, we additionally investigate the average number of citations per publication (see Table 2.8). We find that the higher number of citations for Twitter users appears to stem exclusively from the higher number of publications, and not from additional citations for each publication. Except for the least restrictive specification, the estimates for the number of citations per publication are negative and become more significant when including more control variables. The estimates suggest that the average publication of Twitter users receives 2.9 fewer citations per publication than non-Twitter users. Taken together, although individual papers become less impactful on average, the increase in publications offsets this reduction, resulting in a net increase in citations.

Table 2.8: Effect of Twitter Usage on Number of Citations per Publication

	Dep. Var.: Nr. of Citations per Publication			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	2.340***	-0.617	-1.138*	-1.520**
	(0.578)	(0.581)	(0.624)	(0.632)
Observations	14,588	14,588	14,562	14,562
Mean of DV	27.97	27.97	28.01	28.01
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of citations per publication per year. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

An often suggested advantage of Twitter usage is that it allows researchers to advertise their research to their peers online. We investigate this hypothesis by analyzing who is citing Twitter users. For this, we calculate the share of citations that stem from papers that are written by Twitter users.[10] The results from this analysis are presented in Table 2.9. The estimates are consistently positive, independent of the included control variables. The coefficient in Column (4) implies that, on average, Twitter users receive 15% more citations from papers authored by economists who are also on Twitter.

[10] A paper is considered to be authored by a Twitter user if at least one of the authors uses Twitter at the time of publication.

Table 2.9: Effect of Twitter Usage on Share of Citations from Twitter

	Dep. Var.: Share of Citations from Twitter			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.086***	0.084***	0.082***	0.081***
	(0.004)	(0.003)	(0.003)	(0.004)
Observations	14,665	14,665	14,635	14,635
Mean of DV	0.33	0.33	0.33	0.33
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average yearly share of citations originating from papers with at least one Twitter-using author. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This result may reflect two mechanisms. First, researchers joining Twitter may promote their work to a broader audience, increasing visibility among other Twitter users. Second, Twitter usage may shift research topics toward those that are of greater interest to other economists active on the platform. We return to the latter mechanism in the last part of the paper.

4.3 Effect Heterogeneity

We next study the heterogeneity of the effects of Twitter usage across three dimensions: 1) the intensity of Twitter usage, 2) gender, and 3) the career age of an economist. The first is motivated by the fact that we would expect the effect of Twitter usage to be larger for users who use the platform more frequently. The second stems from the hypothesis that Twitter might allow minorities to reach greater visibility. Lastly, we hypothesize that the effects of Twitter usage could differ by career stage.

First, we investigate heterogeneity by the degree of Twitter usage. A key reason why certain economists benefit more from Twitter may be their differential ability to achieve prominence on the platform. To study this, we estimate effects separately for Twitter users with high and low follower counts, defining high Twitter reach as having more followers than the 75$^{\text{th}}$ percentile (around 1,500 followers). Since this analysis includes only Twitter users, the variation exploited here represents the intensive margin of Twitter usage driven by our instrument. Figure 2.7 shows that among economists with high Twitter reach, more years on Twitter translate into more publications and citations than among economists with fewer followers.

Figure 2.7: Heterogeneity by Twitter Usage Intensity (Follower)

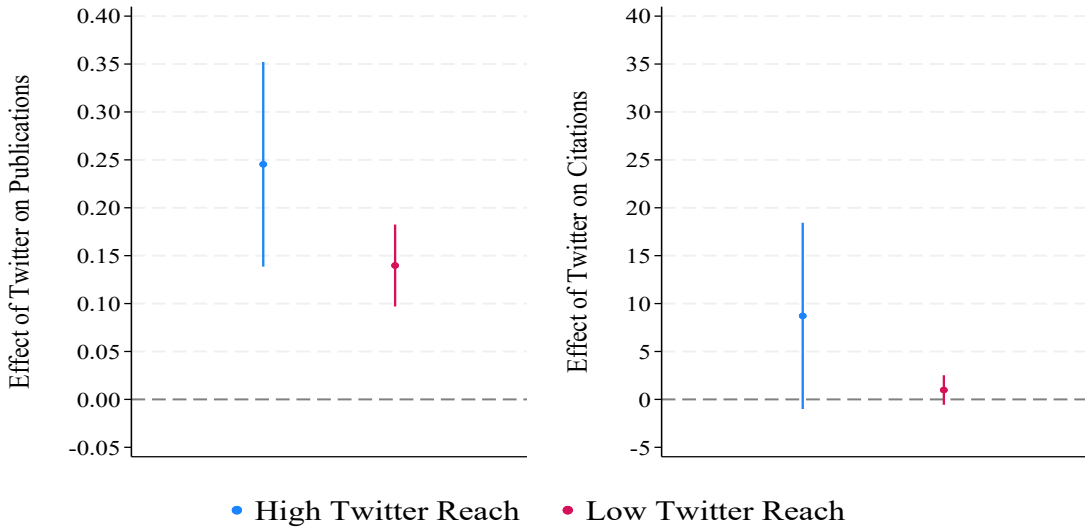

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications and citations among Twitter users. Blue and red represent results from separate regressions in the subgroups of users with high and low Twitter reach respectively, defined by whether the number of followers exceeds the 75$^{\text{th}}$ percentile ($\sim 1,500$). The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

In Appendix Figures A2.3 and A2.4, we confirm that similar patterns emerge if we instead use an alternative measure of Twitter usage intensity based on the total number of tweets or the number of *followed* accounts. We broadly find stronger effects for the individuals who appear to use Twitter more intensively.

Second, we split our sample by gender and investigate whether men or women benefit more from Twitter usage (see Figure 2.8). The results suggest that the average effect of Twitter usage is larger among men, both for the number of publications and the number of citations. This suggests that, in contrast to popular narratives, social media usage is unlikely to narrow the gender gap in economics.

Third, we test whether more experienced economists benefit more from using Twitter. Figure 2.9 presents effects on publications and citations separately for economists who received their PhD before and after the year 2000. For publications, the effect of Twitter usage is larger for experienced economists. Among citations, the positive impact is driven entirely by experienced economists, while for younger economists, the estimated effect is statistically indistinguishable from zero.

Figure 2.8: Heterogeneity by Gender

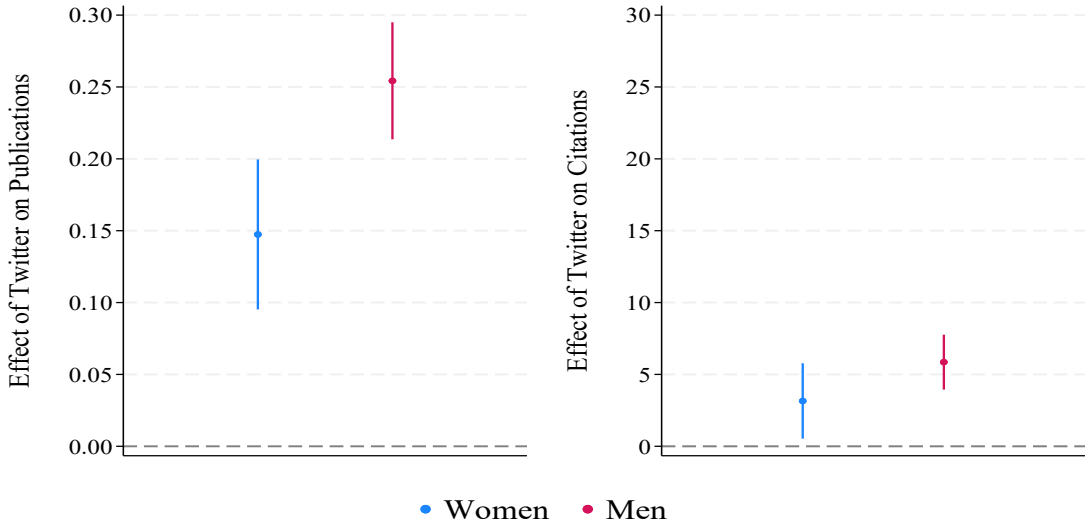

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications and citations. Blue and red represent results from separate regressions among women and men, respectively. The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

Figure 2.9: Heterogeneity by Age

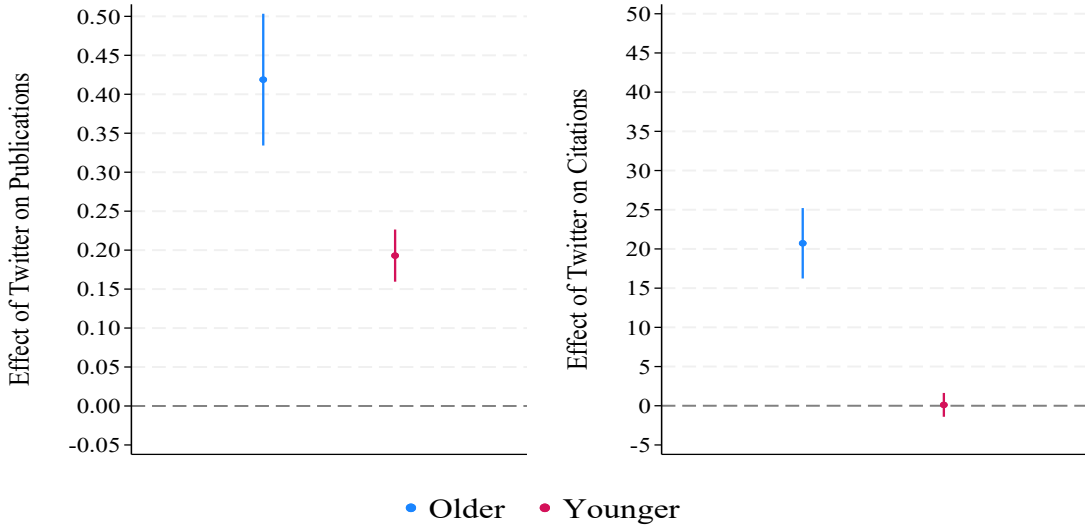

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications and citations. Blue and red represent results from separate regressions within the subgroups of economists who received their PhD before and after the year 2000, respectively. The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

5 Effect of Twitter on Scientific Network

In the second part of the paper, we examine the extent to which Twitter influences the production of papers, as measured by the number of unique co-authors and references in papers. These results thereby shed light on the mechanism underlying the increased research productivity of Twitter users, which we documented in the previous part.

5.1 Co-authorship

One of the key factors that shape the production of papers are co-authorship opportunities. Hence, we investigate whether Twitter affects the number of co-authors economists are working with. For this analysis, we replace the dependent variable in Equation (2.2) with the average number of unique co-authors each economist in our data worked with per year between 2006 and 2022. The results in Table 2.10 show that Twitter usage increases in the number of co-authors economists are working with. The estimates are positive and significant at the 1% level across all specifications, and the coefficients are remarkably stable, independent of the used control variables. On average, an additional year of Twitter usage leads to 0.23 more co-authors (see Column (4)). Given the observed average of 1.9 years of Twitter usage (including non-users), this estimate implies approximately 0.44 additional co-authors per year. These results suggest that the usage of Twitter facilitates the search for suitable co-authors, which explains some of the previously observed higher publication output.

Table 2.10: Effect of Twitter Usage on Number of Co-authors

	Dep. Var.: Number of Co-authors			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.300***	0.252***	0.237***	0.229***
	(0.024)	(0.022)	(0.022)	(0.023)
Observations	15,289	15,289	15,286	15,286
Mean of DV	1.72	1.72	1.72	1.72
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average yearly number of distinct co-authors. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

As further evidence for the effect of Twitter on co-authorship patterns, we investigate whether the documented rise in co-authors is driven by economists who also use Twitter. In Table 2.11, we estimate regressions with the share of co-authors who also use Twitter as the dependent variable. In line with this hypothesis, we observe that the share of co-authors who use Twitter increases by 20 percentage points ($0.108 \cdot 1.9$) – nearly doubling the baseline rate. This provides further evidence that Twitter facilitates connections among economists with shared research interests, ultimately increasing productivity.

Table 2.11: Effect of Twitter Usage on Share of Co-authors from Twitter

	Dep. Var.: Share of Co-authors Twitter			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.108***	0.107***	0.108***	0.108***
	(0.006)	(0.005)	(0.006)	(0.006)
Observations	14,292	14,292	14,253	14,253
Mean of DV	0.19	0.19	0.19	0.19
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average yearly share of Twitter-using co-authors. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

As a last piece of evidence for the effect of Twitter on co-authorship, we study the impact of Twitter on the quality of co-authors. More specifically, we analyze whether Twitter changes the average number of Top 100, Top 50, Top General Interest, or Top 5 publications per co-author. The estimates from this exercise are reported in Figure 2.10. While we find insignificant and negative estimates for the number of Top 100 and Top 50 publications, we observe positive and statistically significant effects on the number of Top General Interest and Top 5 publications. Given the particular importance of publication in high-quality journals in economics, the results provide evidence for the fact that Twitter users, on average, are able to write papers with higher-quality co-authors.

We report the precise estimates from this exercise in Appendix Table A2.9. While the estimates appear small, the coefficient for the number of Top 5 publications implies that the average co-author of Twitter users has additional 0.03 Top 5 publications, an increase of around 6% relative to the mean (see Column (4)). Similarly, the estimate for publications in Top General Interest journals implies an increase of around 7% (see Column (3)).

Together, these results indicate that Twitter enhances the productivity of economists by expanding the number of co-authors, as well as improving the quality of co-authorship. We next turn our attention to the acquisition of new knowledge about papers through Twitter.

Figure 2.10: Effect on Quality of Co-Authors

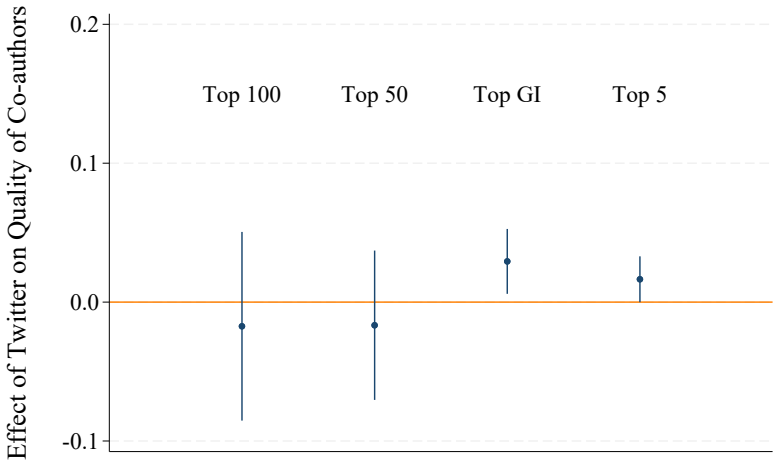

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications in different brackets of journal impact factor of researchers' average co-author. The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

5.2 Referencing

Another important channel through which Twitter could impact research productivity is by helping economists to learn about new papers they otherwise would have missed. We study this question in Table 2.12, which analyzes the average number of references economists provide in their papers. Note that for this analysis, we exclude self-citations and control non-parametrically for the number of publications of each author. As a result, the estimates will only pick up if, conditional on the number of papers, Twitter increases the number of references to other people's research. We find that Twitter significantly increases the average number of references authors cite in their work across all specifications. This is in line with the hypothesis that Twitter usage helps economists learn about new and relevant research for their own work.

As further evidence for this idea, we investigate the share of references that cite work by authors who are also on Twitter (see Table 2.13).[11] Similar to the results for co-authors, we find that Twitter usage significantly increases the share of references to work by other Twitter users. The estimates are large, positive, and significant, suggesting a more than doubling of the share of references to work by Twitter users.

[11] As before, a paper is considered to be authored by a Twitter user, if at least one of the authors uses Twitter.

Table 2.12: Effect of Twitter Usage on Number of References

	Dep. Var.: Number of References			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	11.032***	10.430***	10.351***	10.157***
	(0.676)	(0.586)	(0.648)	(0.649)
Observations	15,289	15,289	15,286	15,286
Mean of DV	43.40	43.40	43.40	43.40
Number of Publications FE	Yes	Yes	Yes	Yes
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of references provided per year (excluding references to own papers). All specifications include fixed effects for the average number of publications per year, and cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2.13: Effect of Twitter Usage on Share of References to Twitter Users

	Dep. Var.: Share of References Twitter			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.081***	0.086***	0.090***	0.090***
	(0.004)	(0.003)	(0.004)	(0.004)
Observations	15,029	15,029	15,015	15,015
Mean of DV	0.12	0.12	0.12	0.12
Number of Publications FE	Yes	Yes	Yes	Yes
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average yearly share of references to papers with at least one Twitter-using author (excluding references to own papers). All specifications include fixed effects for the average number of publications per year, and cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

As a last piece of evidence on the impact of Twitter on citation behaviors, we study the impact of Twitter on the number of references by quality. In Figure 2.11, we visualize the estimates for the number of references to papers published in the Top 100, Top 50, Top General Interest, and Top 5 journals. Detailed estimates are reported in Appendix Table A2.10. Note that all of these specifications include fixed effects for the total number of publications. We find that Twitter usage increases the number of references across the quality distribution of journals. In other words, Twitter usage does *not* disproportionately concentrate citations in higher- or lower-ranked journals. Overall, this finding also suggests that Twitter helps economists discover papers relevant to their research.

Figure 2.11: Effect on Quality of References

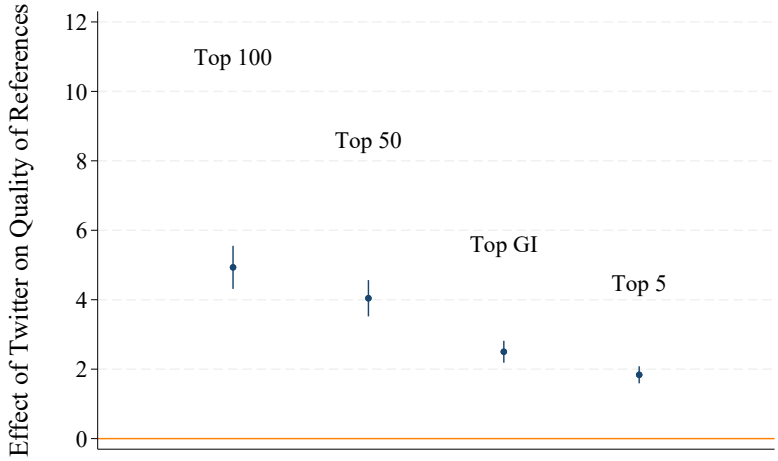

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly references to publications in different brackets of journal impact factor. The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

6 Effect of Twitter on Research Topics

Given the significant impact of Twitter on the co-authorship and citation networks we have documented in the previous part, it is only natural to ask to what extent Twitter also impacts the direction of research. In the last part of the paper, we therefore analyze to what extent Twitter influences the topics that researchers work on. First, we study the similarity of research done by other economists on and off Twitter. Second, we analyze whether Twitter increases the probability that economists work on "hot" topics such as Covid-19, cryptocurrencies, or Trump. Third, we investigate whether Twitter usage enhances the visibility and impact of economists beyond academia, as measured by Altmetric scores.

6.1 Topic Similarity to Twitter

We start by analyzing whether Twitter usage increases the similarity of one's own research to work done by other Twitter users. More specifically, we create embeddings for each paper in our data based on the words used in the abstract and calculate the cosine similarity to the 100 most similar papers published in the same year (excluding one's own papers). We repeat this similarity calculation separately for other papers that were authored by Twitter users and non-Twitter users. In Panel A in Table 2.14, we then report the estimates for the similarity to work done by Twitter users, while Panel B reports the estimates for the similarity to work done by non-Twitter users. To ease interpretation, we standardize the dependent variable to have a mean of 0 and a standard deviation of 1.

We find that Twitter usage indeed seems to lead economists to work on topics that are more similar to the work done by other economists on Twitter. On average, Twitter usage increases the cosine similarity to work done by other Twitter users by more than 1 standard deviation (see Column (4) Panel A). In stark contrast, Twitter usage appears to shift research away from the topics of economists who are not on Twitter, as we find significant decreases in the cosine similarity to the work of economists who are not using Twitter of around 0.2 standard deviations (see Panel B). These results indicate that Twitter adoption shifts economists' research away from more traditional topics in economics, toward topics that are more prominent among Twitter users. Further, they also suggest a divergence of research topics between economists who use and do not use Twitter.

6.2 Hot Topics

Motivated by the previous results, we investigate some of the topics that might underlie the shifts in similarity of research. More specifically, we analyze whether the increasing similarity of economists' research to that of other Twitter users is reflected by a shared focus on trending ("hot") topics. To do so, we select a set of keywords that represent prominent hot topics and flag papers on these topics if they use one of the keywords in their abstract. Note that the aim of this analysis is not to cover all hot topics, but rather investigate some topics with a high salience on Twitter.

The results from this analysis are presented in Table 2.15, in which we report the effect of Twitter usage on the number of papers covering the topics: 1) Covid-19, 2) Climate Change, 3) Cryptocurrencies, 4) Gender, and 5) Donald Trump. Except for climate change, we find that Twitter usage significantly increases the likelihood of writing a paper on one of these hot topics. Most notably, an additional year of Twitter raises the number of ab-

Table 2.14: Effect of Twitter Usage on Similarity of Research to Twitter

	Dep. Var.: Similarity of Research (std.) to			
	(1)	(2)	(3)	(4)
Panel A: Research of Economists on Twitter				
Years of Twitter Usage	0.668***	0.765***	0.852***	0.864***
	(0.037)	(0.037)	(0.045)	(0.047)
Observations	14,089	14,089	14,054	14,054
Mean of DV	0.01	0.01	0.01	0.01
Panel B: Research of Economists off Twitter				
Years of Twitter Usage	-0.132***	-0.125***	-0.102***	-0.108***
	(0.020)	(0.020)	(0.021)	(0.021)
Observations	14,215	14,215	14,180	14,180
Mean of DV	0.01	0.01	0.01	0.01
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average textual similarity of research to that of Twitter users and non-users, respectively. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

stracts mentioning Covid-19 by 3.8 percentage points, on average, from a baseline of 5 percent. Thus, roughly one out of every 14 economists on Twitter publishes a paper related to Covid-19 specifically due to joining the platform. We find similar, though smaller, increases in the number of papers on Cryptocurrencies, Gender, and Trump. The results are overall consistent with Twitter shifting economists' research towards topics that they see trending on Twitter.

6.3 Attention outside of Academia

As a final piece of evidence, we study the attention that is paid to the work of economists outside of academia. For this analysis, we investigate whether Twitter usage impacts the Altmetric scores of researchers, as a proxy for the broader attention research receives. To start, we analyze the overall Altmetric Attention Score – standardized to mean 0 and standard deviation 1 – of research published by Twitter and non-Twitter users. The results in Table 2.16 show a substantial increase in the Attention Score for Twitter users. Each additional year on Twitter leads to a 0.08 standard deviation increase in attention, translating into a 0.15 standard deviation increase for the average Twitter user relative to non-Twitter users in our data.

Table 2.15: Effect of Twitter Usage on Research on Hot Topics

	Dep. Var.: Number of Abstracts with Keyword				
	Covid	Climate	Crypto	Gender	Trump
	(1)	(2)	(3)	(4)	(5)
Years of Twitter Usage	0.038***	0.002	0.002***	0.006*	0.002***
	(0.004)	(0.004)	(0.000)	(0.003)	(0.001)
Observations	15,011	15,011	15,011	15,011	15,011
Mean of DV	0.05	0.04	0.00	0.07	0.00
Years Since PhD FE	Yes	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes	Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of papers mentioning specific keywords in abstract. All specification include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 2.16: Effect of Twitter Usage on Impact of Research

	Dep. Var.: Average Altmetric Score (std.)			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.086***	0.078***	0.079***	0.077***
	(0.013)	(0.013)	(0.015)	(0.016)
Observations	15,289	15,289	15,286	15,286
Mean of DV	0.00	0.00	-0.00	-0.00
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average yearly Altmetric Attention Score. The specifications cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

To get a better understanding of what underlies these increases in the Altmetric Attention Score, we study its individual components. The results in Table 2.17 indicate that the increase in the Attention Score is primarily driven by mentions in mainstream news media and, to a lesser extent, by mentions in Twitter posts. In addition, we also find significant increases in the mentions of research in blogs, policy documents, and on Wikipedia. The magnitudes imply that the average Twitter user receives approximately 0.084 more mentions from mainstream media and 0.013 more mentions from blogs. These increases are sizable given the low baseline rate of such mentions, representing more than 100 percent growth. Moreover, each individual mention in the news typically receives relatively high attention, which is reflected by their large weight in the Altmetric Attention Score.

Overall, these results confirm that Twitter usage significantly increases the attention paid to research, not only inside academia but also by the media and policymakers. Thus, Twitter usage enhances the visibility of economists' work beyond academic circles.

Table 2.17: Effect of Twitter Usage on Altmetric Attention Score Components

	Dep. Var.: Yearly Number of Altmetric Mentions by				
	Mainstream Media	Blogs	Policy Documents	Wikipedia	Twitter Posts
	(1)	(2)	(3)	(4)	(5)
Years of Twitter Usage	0.044***	0.007***	0.011***	0.002**	0.498***
	(0.015)	(0.001)	(0.002)	(0.001)	(0.077)
Observations	15,286	15,286	15,286	15,286	15,286
Mean of DV	0.06	0.01	0.03	0.01	0.67
Years Since PhD FE	Yes	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes	Yes
Attention Score weight	8	5	3	3	0.25

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of mentions in different media, as recorded by Altmetric. The weight of each mention in the Altmetric Attention Score by type is indicated below. All specification include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

7 Conclusion

This paper provides the first evidence that Twitter has substantially influenced economic research productivity, collaboration networks, and topic selection among economists. Based on a novel database of economists, their PhD institutions, and their Twitter accounts, we construct an instrumental variables strategy exploiting peer effects in Twitter usage across PhD cohorts, which allows us to isolate the effect of Twitter usage. We find that Twitter usage increases economists' academic output and citations, leads to more co-authorships with other economists on Twitter, and a higher share of references to research of other Twitter users. Moreover, research focuses more on topics that are prominent on Twitter, leading to increased attention both on Twitter and from the broader public.

Overall, our research highlights that Twitter has transformed the research done by economists, shaping not only how research is conducted and disseminated but also what research is produced. The key open questions involve assessing whether the observed pivot toward socially engaging topics represents an overall beneficial evolution in economic research, enhancing its relevance and societal impact, or whether it signifies a movement toward short-term interests and potentially superficial engagement. These considerations remain vital for shaping policy and institutional practices around the adoption and use of social media within academia.

Bibliography

Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi, "Spread of (mis) information in social networks," *Games and Economic Behavior*, 2010, *70* (2), 194–227.

_ , Munther A Dahleh, Ilan Lobel, and Asuman Ozdaglar, "Bayesian learning in social networks," *The Review of Economic Studies*, 2011, *78* (4), 1201–1236.

Advani, Arun, Elliott Ash, Anton Boltachka, David Cai, and Imran Rasul, "Race-related research in economics," 2025.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow, "The welfare effects of social media," *American Economic Review*, 2020, *110* (3), 629–676.

Angrist, Joshua D, "The perils of peer effects," *Labour Economics*, 2014, *30*, 98–108.

Aridor, Guy, Rafael Jiménez-Durán, Ro'ee Levy, and Lena Song, "The economics of social media," *Journal of Economic Literature*, 2024, *62* (4), 1422–1474.

Banerjee, Abhijit V, "A simple model of herd behavior," *The quarterly journal of economics*, 1992, *107* (3), 797–817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of political Economy*, 1992, *100* (5), 992–1026.

Bond, Robert M, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler, "A 61-million-person experiment in social influence and political mobilization," *Nature*, 2012, *489* (7415), 295–298.

Braghieri, Luca, Ro'ee Levy, and Alexey Makarin, "Social media and mental health," *American Economic Review*, 2022, *112* (11), 3660–3693.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin, "Identification of peer effects through social networks," *Journal of econometrics*, 2009, *150* (1), 41–55.

_ , _ , and _ , "Peer effects in networks: A survey," *Annual Review of Economics*, 2020, *12* (1), 603–629.

Chaisemartin, Clément De and Xavier d'Haultfoeuille, "Two-Way Fixed Effects Estimators With Heterogeneous Treatment Effects," *American Economic Review*, 2020, *110* (9), 2964–96.

_ and Xavier D'Haultfoeuille, "Two-Way Fixed Effects and Differences-In-Differences With Heterogeneous Treatment Effects: A Survey," Technical Report, National Bureau of Economic Research 2022.

Eysenbach, Gunther et al., "Can tweets predict citations? Metrics of social impact based on Twitter and correlation with traditional metrics of scientific impact," *Journal of medical Internet research*, 2011, *13* (4), e2012.

Fujiwara, Thomas, Karsten Müller, and Carlo Schwarz, "The effect of social media on elections: Evidence from the United States," Technical Report, National Bureau of Economic Research 2021.

Galasso, Alberto and Mark Schankerman, "Patents and cumulative innovation: Causal evidence from the courts," *The Quarterly Journal of Economics*, 2015, *130* (1), 317–369.

Garg, Prashant and Thiemo Fetzer, "Political expression of academics on Twitter," *Nature Human Behaviour*, 2025.

Giorgi, Giacomo De, Michele Pellizzari, and Silvia Redaelli, "Identification of social interactions through partially overlapping peer groups," *American Economic Journal: Applied Economics*, 2010, *2* (2), 241–275.

Goldsmith-Pinkham, Paul and Guido W Imbens, "Social networks and the identification of peer effects," *Journal of Business & Economic Statistics*, 2013, *31* (3), 253–264.

Haustein, Stefanie, Isabella Peters, Cassidy R Sugimoto, Mike Thelwall, and Vincent Larivière, "Tweeting biomedicine: An analysis of tweets and citations in the biomedical literature," *Journal of the Association for Information Science and Technology*, 2014, *65* (4), 656–669.

Huang, Jen-Hung and Yi-Fen Chen, "Herding in online product choice," *Psychology & Marketing*, 2006, *23* (5), 413–428.

Jaffe, Adam B, Manuel Trajtenberg, and Rebecca Henderson, "Geographic localization of knowledge spillovers as evidenced by patent citations," *the Quarterly journal of Economics*, 1993, *108* (3), 577–598.

Jones, Benjamin F, Stefan Wuchty, and Brian Uzzi, "Multi-university research teams: Shifting impact, geography, and stratification in science," *science*, 2008, *322* (5905), 1259–1262.

Koffi, Marlène, Roland Pongou, and Leonard Wantchekon, "The color of ideas: Racial dynamics and citations in economics," Technical Report, National Bureau of Economic Research 2024.

Le, Quoc and Tomas Mikolov, "Distributed representations of sentences and documents," in "International conference on machine learning" PMLR 2014, pp. 1188–1196.

Lin, Xu, "Identifying peer effects in student academic achievement by spatial autoregressive models with group unobservables," *Journal of Labor Economics*, 2010, *28* (4), 825–860.

Manski, Charles F, "Identification of endogenous social effects: The reflection problem," *The review of economic studies*, 1993, *60* (3), 531–542.

Mongeon, Philippe, Timothy D Bowman, and Rodrigo Costas, "An open data set of scholars on Twitter," *Quantitative Science Studies*, 2023, *4* (2), 314–324.

Moser, Petra and Alessandra Voena, "Compulsory licensing: Evidence from the trading with the enemy act," *American Economic Review*, 2012, *102* (1), 396–427.

Müller, Karsten and Carlo Schwarz, "Fanning the flames of hate: Social media and hate crime," *Journal of the European Economic Association*, 2021, *19* (4), 2131–2167.

_ and _, "From Hashtag to Hate Crime: Twitter and Anti-Minority Sentiment," *Forthcoming American Economic Journal: Applied Economics*, 2022.

Özkent, Yasemin, "Social media usage to share information in communication journals: An analysis of social media activity and article citations," *Plos one*, 2022, *17* (2), e0263725.

Peri, Giovanni, "Determinants of knowledge flows and their effect on innovation," *Review of economics and Statistics*, 2005, *87* (2), 308–322.

Rosario, Ana Babić, Francesca Sotgiu, Kristine De Valck, and Tammo HA Bijmolt, "The effect of electronic word of mouth on sales: A meta-analytic review of platform, product, and metric factors," *Journal of marketing research*, 2016, *53* (3), 297–318.

Sathianathen, Niranjan Jude, Robert Lane III, Declan G Murphy, Stacy Loeb, Caitlin Bakker, Alastair D Lamb, and Christopher J Weight, "Social media coverage of scientific articles immediately after publication predicts subsequent citations-# SoME_Impact Score: observational analysis," *Journal of medical Internet research*, 2020, *22* (4), e12288.

Scharfstein, David S and Jeremy C Stein, "Herd behavior and investment," *The American economic review*, 1990, pp. 465–479.

Scotchmer, Suzanne, "Standing on the shoulders of giants: cumulative research and the patent law," *Journal of economic perspectives*, 1991, *5* (1), 29–41.

Sias, Richard W, "Institutional herding," *The Review of Financial Studies*, 2004, *17* (1), 165–206.

Smith, Zachary L, Austin L Chiang, Deborah Bowman, and Michael B Wallace, "Longitudinal relationship between social media activity and article citations in the journal Gastrointestinal Endoscopy," *Gastrointestinal endoscopy*, 2019, *90* (1), 77–83.

Tonia, Thomy, Herman Van Oyen, Anke Berger, Christian Schindler, and Nino Künzli, "If I tweet will you cite? The effect of social media exposure of articles on downloads and citations," *International journal of public health*, 2016, *61*, 513–520.

Tucker, Joshua A, Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan, "Social media, political polarization, and political disinformation: A review of the scientific literature," *Political polarization, and political disinformation: a review of the scientific literature (March 19, 2018)*, 2018.

Welch, Ivo, "Herding among security analysts," *Journal of Financial economics*, 2000, *58* (3), 369–396.

Wuchty, Stefan, Benjamin F Jones, and Brian Uzzi, "The increasing dominance of teams in production of knowledge," *Science*, 2007, *316* (5827), 1036–1039.

Zhuravskaya, Ekaterina, Maria Petrova, and Ruben Enikolopov, "Political effects of the internet and social media," *Annual review of economics*, 2020, *12*, 415–438.

Appendix

A.1 Additional Details: Data

Figure A2.1: Twitter share of active researchers over time, by quartiles of citations

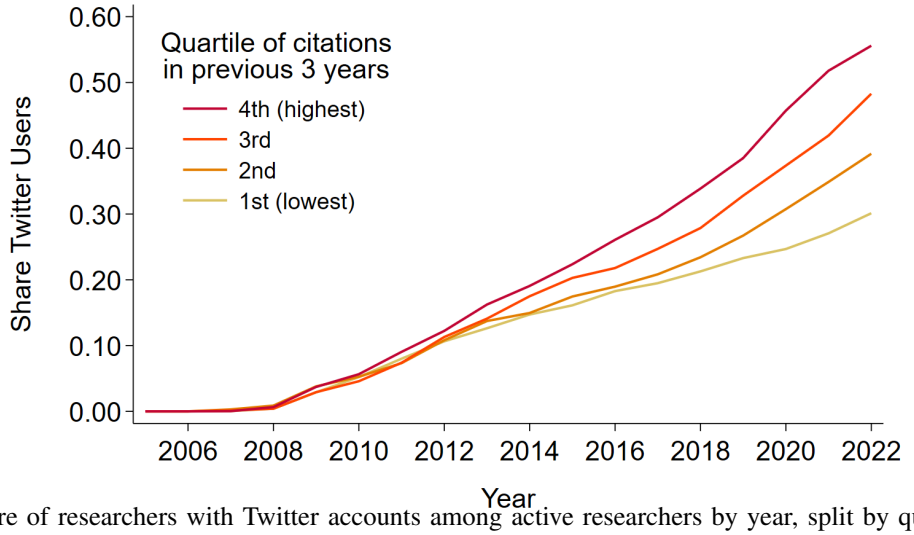

Notes: Share of researchers with Twitter accounts among active researchers by year, split by quartiles of number of citations in past three years.

A.2 General Definition of Instrument

We define each researcher's Twitter Pressure as the average share of Twitter users in the fields of their PhD colleagues (excluding the researcher's own fields). Formally,

$$Twitter\ Pressure_{\ i,t} = \frac{1}{|Cohort_i| - 1} \sum_{\substack{c \in Cohort_i, \\ c \neq i}} \sum_{\substack{f \in Fields_c, \\ f \notin Fields_i}} Field\ Intensity_{\ c}^{\ f} \times Sh.\ Twitter\ in\ Field_{\ t}^{\ f}\ . \tag{2.4}$$

$Twitter\ Pressure$ for researcher i in year t is a weighted average of Twitter adoption in fields f of PhD colleagues c. The summation runs over all colleagues in researcher i's cohort and over all fields except i's own ones. The weights $Field\ Intensity_{\ c}^{\ f}$ measure the relative focus of colleague c on field f. If colleague c is active in only one field, then $Field\ Intensity_{\ c}^{\ f}$ is an indicator equal to one for that field, and zero otherwise. Instead, if c is active in more than one field, then each field is weighted according to c's relative

focus on that field. In particular,

$$Field\ Intensity\ {}^{f}_{c} = \frac{1}{N\ Papers\ _{c}} \sum_{p_c \in Papers\ _{c}} \frac{N\ Field\ f\ in\ p_c}{N\ Fields\ in\ p_c}\ , \tag{2.5}$$

where $N\ Papers\ _{c}$ is the number of c's papers that include a JEL code, and $N\ Field\ f\ in\ p_c$ counts how often field f appears in a given paper p_c. If, for example, a paper mentions JEL codes A10, A11, and B10, then the two-character JEL code A1 would receive weight two-third, and B1 would receive one-third. Moreover, in papers that mention fewer JEL codes, each individual JEL code receives more weight.

Similarly, the share of twitter users in each field is the sum of Twitter users in that field, weighted by the importance of that field for each researcher:

$$Sh.\ Twitter\ in\ Field\ {}^{f}_{t} = \sum_{\ell \in {}^{Researchers}_{in\ Field\ f}} \frac{Field\ Intensity\ {}^{f}_{\ell}}{\sum_m Field\ Intensity\ {}^{f}_{m}} \times I(Twitter_{\ell,t})\ , \tag{2.6}$$

where $I(Twitter_{\ell,t})$ is an indicator for whether we observe a Twitter account for researcher ℓ by year t. This weighting scheme ensures that researchers whose work is more exclusively focused on field f contribute more strongly to that field's Twitter share compared to researchers active across multiple fields. When computing the instrument, we apply a leave-one-out measure by excluding the researcher in question from all calculations.

A.3 Additional Results: First Stage

Table A2.1: Robustness First Stage: Cohort Lengths

	Dep. Var.: Years Twitter Usage					
	1-year (1)	2-year (2)	3-year (3)	4-year (4)	5-year (5)	6-year (6)
Twitter Pressure 1-year (std.)	0.613***					
	(0.035)					
Twitter Pressure 2-year (std.)		0.742***				
		(0.033)				
Twitter Pressure 3-year (std.)			0.825***			
			(0.031)			
Twitter Pressure 4-year (std.)				0.871***		
				(0.031)		
Twitter Pressure 5-year (std.)					0.890***	
					(0.031)	
Twitter Pressure 6-year (std.)						0.899***
						(0.031)
Observations	15,417	15,417	15,417	15,417	15,417	15,400
Mean of DV	1.94	1.94	1.94	1.94	1.94	1.94
Years Since PhD FE	Yes	Yes	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Regressions of the number of years of Twitter usage on standardized Twitter Pressure, for different definitions of cohort lengths. Years represent the maximum difference between graduation years to be considered part of the same cohort. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2.2: Robustness First Stage: Research fields

	Dep. Var.: Years Twitter Usage				
	JEL1 (1)	JEL 1 (Top 1) (2)	JEL2 (3)	JEL2 (Top 3) (4)	JEL2 (Top 5) (5)
Twitter Pressure JEL1 (std.)	0.205***				
	(0.033)				
Twitter Pressure JEL1 Top1 (std.)		0.856***			
		(0.031)			
Twitter Pressure JEL2 (std.)			0.655***		
			(0.032)		
Twitter Pressure JEL2 Top3 (std.)				0.890***	
				(0.031)	
Twitter Pressure JEL2 Top5 (std.)					0.892***
					(0.031)
Observations	15,417	15,417	15,408	15,417	15,411
Mean of DV	1.94	1.94	1.94	1.94	1.94
Years Since PhD FE	Yes	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes	Yes

Notes: Regressions of the number of years of Twitter usage on standardized Twitter Pressure, for different definitions of research fields based on JEL codes. Column (1) uses all one-character codes (A-Z), and Column (2) uses each researcher's most frequent one-character code. Similarly, Column (3) uses all two-character codes (A1-Z3), while Columns (4) and (5) use the Top 3 and Top 5, respectively. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2.3: Robustness First Stage: Clustering

	Dep. Var.: Years Twitter Usage			
	(1) PhD Institution-by-Year	(2) PhD Institution	(3) PhD Year	(4) Individual
Twitter Pressure JEL2 Top3 (std.)	0.890***	0.890***	0.890***	0.890***
	(0.031)	(0.039)	(0.052)	(0.030)
Observations	15,417	15,417	15,417	15,417
Mean of DV	1.94	1.94	1.94	1.94
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes

Notes: Regressions of the number of years of Twitter usage on standardized Twitter Pressure, clustering standard errors at different levels. Column (1) clusters by unique combinations of PhD institution and graduation year, while Columns (2) and (3) cluster separately by institution and graduation year, respectively. Column (4) does not cluster standard errors, computing simple heterogeneity-robust standard errors. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2.4: Robustness First Stage: Functional Form

	Dep. Var.: Years Twitter Usage					
	(1)	(2)	(3)	(4)	(5)	(6)
Twitter Pressure (std.)	0.890***					
	(0.031)					
Twitter Pressure (pct.)		0.031***				
		(0.001)				
Twitter Pressure (asinh)			1.155***			
			(0.040)			
I[Twitter Pressure≥50th pct]				1.151***		
				(0.080)		
I[Twitter Pressure≥75th pct]					1.226***	
					(0.090)	
I[Twitter Pressure≥90th pct]						0.714***
						(0.091)
Observations	15,417	15,417	15,417	15,417	15,417	15,417
Mean of DV	1.94	1.94	1.94	1.94	1.94	1.94
Years Since PhD FE	Yes	Yes	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Regressions of the number of years of Twitter usage on different transformations of Twitter Pressure. Column (1) uses the z-standardized value. Column (2) transforms into percentiles, and Column (3) applies the inverse hyperbolic sine. Columns (4) to (6) use indicators for Twitter Pressure being above the 50th, 75th, and 90th percentile. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2.5: Robustness First Stage: Restrict to pre-2006 PhD Graduates

	Dep. Var.: Years of Twitter Usage			
	(1)	(2)	(3)	(4)
Twitter Pressure (std.)	1.018***	1.409***	1.395***	1.391***
	(0.064)	(0.069)	(0.078)	(0.078)
Observations	7,088	7,088	6,963	6,963
Mean of DV	1.77	1.77	1.77	1.77
R^2	0.02	0.03	0.02	0.02
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: Regressions of the number of years of Twitter usage on standardized Twitter Pressure among researchers with at least one observed paper before 2006. The specifications include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A2.2: Leave-One-PhD-Institution-Out Estimation

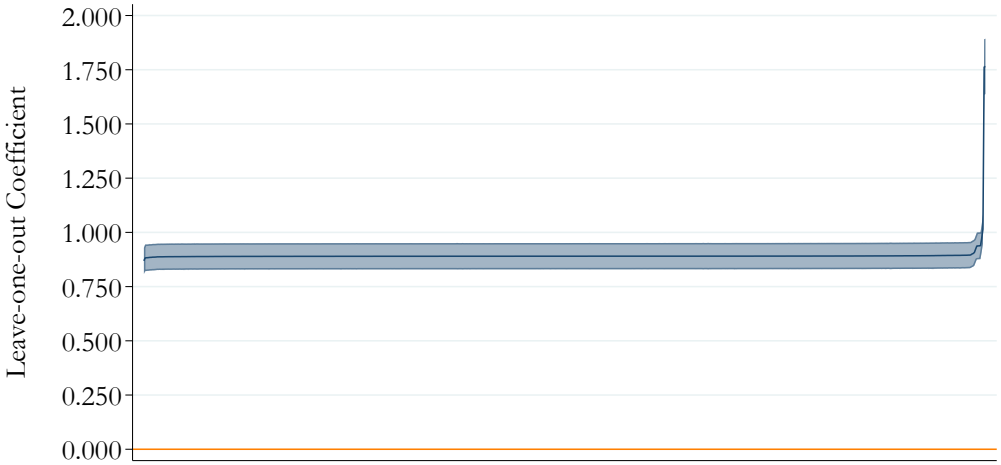

Notes: Coefficients and 95 percent confidence intervals from regressions of the number of years of Twitter usage on Twitter Pressure, leaving out one PhD institution at a time. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

A.4 Additional Results: Output

Table A2.6: Effect of Twitter Usage on Number of Publications by Journal Quality

	Dep. Var.: Nr. of Publications per Year				
	All	Top 100	Top 50	Top GI	Top 5
	(1)	(2)	(3)	(4)	(5)
Years of Twitter Usage	0.239***	0.091***	0.066***	0.025***	0.013***
	(0.015)	(0.006)	(0.005)	(0.002)	(0.002)
Observations	15417	15417	15417	15417	15417
Mean of DV	1.07	0.41	0.27	0.07	0.04
F-Stat (KP)	824.55	824.55	824.55	824.55	824.55
Years Since PhD FE	Yes	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes	Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of publications in journals of different quality brackets. Column (1) includes all publications, while Columns (2) and (3) count publications in the Top 100 and Top 50 journals by impact factor, respectively, while Columns (4) and (5) refer to the 12 Top General Interest and Top 5 journals, respectively. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2.7: Effect of Twitter Usage on Number of Publications (Controlling for Pre-Period)

	Dep. Var.: Nr. of Publications per Year			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.098***	0.196***	0.194***	0.193***
	(0.023)	(0.019)	(0.020)	(0.021)
Observations	7,045	7,045	6,918	6,918
Mean of DV	1.26	1.26	1.26	1.26
Pre-Twitter Publication FE	Yes	Yes	Yes	Yes
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of publications per year, among researchers with at least one observed paper before 2006. All specifications include fixed effects for the number of publications before 2006, and cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A2.3: Heterogeneity by Twitter Usage Intensity (Tweets)

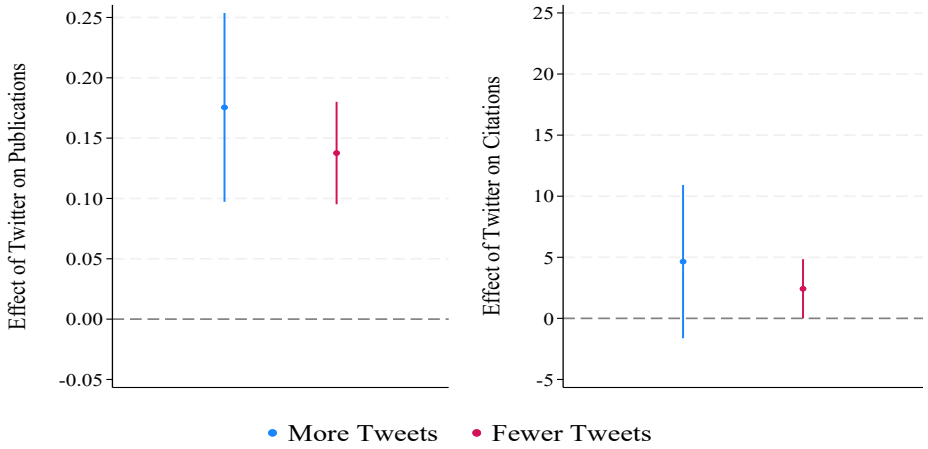

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications and citations among Twitter users. Blue and red represent results from separate regressions in the subgroups of users with high and low intensity of Twitter usage, respectively, defined by whether the number of tweets exceeds the 75$^{\text{th}}$ percentile ($\sim 1,600$). The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

Figure A2.4: Heterogeneity by Twitter Usage Intensity (Following)

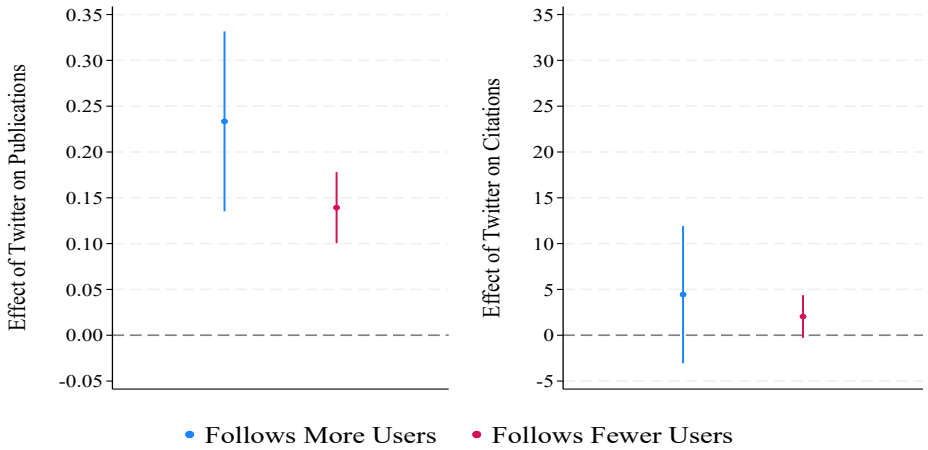

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of yearly publications and citations among Twitter users. Blue and red represent results from separate regressions in the subgroups of users with high and low intensity of Twitter usage, respectively, defined by whether the number of followed accounts exceeds the 75$^{\text{th}}$ percentile (~ 800). The specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors are clustered by PhD institution-by-year.

Table A2.8: Effect of Twitter Usage on Number of Citations (Controlling for Pre-Period)

	Dep. Var.: Nr. of Citations per Year			
	(1)	(2)	(3)	(4)
Years of Twitter Usage	2.751***	2.878***	3.346***	2.855***
	(0.907)	(0.745)	(1.002)	(0.996)
Observations	5,491	5,490	5,356	5,356
Mean of DV	25.84	25.84	26.01	26.01
Pre-Period Citation FE	Yes	Yes	Yes	Yes
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE		Yes	Yes	Yes
PhD Institution FE			Yes	Yes
JEL Code Control				Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average number of citations received per year, among researchers with at least one observed paper before 2006. All specifications include fixed effects for the number of citations before 2006, and cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.5 Additional Results: Production of Papers

Table A2.9: Effect of Twitter Usage on Quality of Co-authors

	Dep. Var.: Nr. of Publications of co-authors			
	Top 100 (1)	Top 50 (2)	Top GI (3)	Top 5 (4)
Years of Twitter Usage	-0.017	-0.017	0.029**	0.016*
	(0.035)	(0.027)	(0.012)	(0.008)
Observations	15216	15216	15216	15216
Mean of DV	1.98	1.34	0.40	0.22
F-Stat (KP)	453.95	453.95	453.95	453.95
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average co-author's number of publications in journals of different quality brackets. Columns (1) and (2) count publications in the Top 100 and Top 50 journals by impact factor, respectively, while Columns (3) and (4) refer to the 12 Top General Interest and Top 5 journals, respectively. All specifications include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A2.10: Effect of Twitter Usage on Quality of References

	Dep. Var.: Nr. of References to Publication in:			
	Top 100 (2)	Top 50 (3)	Top GI (4)	Top 5 (5)
Years of Twitter Usage	4.932***	4.043***	2.502***	1.838***
	(0.317)	(0.267)	(0.161)	(0.125)
Observations	15286	15286	15286	15286
Mean of DV	21.24	17.64	9.05	7.17
F-Stat (KP)	504.79	504.79	504.79	504.79
Number of Publications FE	Yes	Yes	Yes	Yes
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the average yearly number of references provided to publications in journals of different quality brackets. Columns (1) and (2) count publications in the Top 100 and Top 50 journals by impact factor, respectively, while Columns (3) and (4) refer to the 12 Top General Interest and Top 5 journals, respectively. All specifications include fixed effects for the average number of publications per year, and cumulatively include fixed effects for the number of years since PhD graduation (Column 1), size of the PhD cohort (2), PhD-granting institution (3), and the most frequent one-character JEL code across all years (4). Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.6 Additional Results: Research Topics

Table A2.11: Effect of Twitter Usage on Twitter mentions

	Dep. Var.: Yearly Altmetric Twitter Mentions by			
	Public	Scientist	Clinician	Communicator
	(1)	(2)	(3)	(4)
Years of Twitter Usage	0.363***	0.120***	0.010***	0.005***
	(0.061)	(0.017)	(0.002)	(0.002)
Observations	15,286	15,286	15,286	15,286
Mean of DV	0.50	0.14	0.02	0.01
Years Since PhD FE	Yes	Yes	Yes	Yes
PhD Cohort Size FE	Yes	Yes	Yes	Yes
PhD Institution FE	Yes	Yes	Yes	Yes
JEL Code Control	Yes	Yes	Yes	Yes

Notes: 2SLS estimates for the effect of years of Twitter usage, instrumented by cohort-specific Twitter Pressure, on the number of Twitter mentions by different user types, as recorded by Altmetric. Column (1) refers to member of the public, while Columns (2) to (4) indicate mentions by users who indicate connections to academic research through their biography or tweets, separately classified as scientist, clinical scientist, or science communicator. All specification include fixed effects for the number of years since PhD graduation, size of the PhD cohort, PhD-granting institution, and the most frequent one-character JEL code across all years. Standard errors (in parentheses) are clustered by PhD institution-by-year. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Chapter 3

Disruptive Activism and Environmental Policy Support

Abstract

Social movements may choose disruptive forms of protest to attract controversy, increasing public attention at the cost of negative sentiment toward itself and ideologically similar groups. This paper analyzes the short-term effects of such protests carried out by the *Last Generation*, a German environmental activist group known for blocking roads and covering public structures in orange paint. Combining detailed data on protests with large-scale survey responses, I exploit the timing and location of protests to identify causal effects on environmental attitudes and support for the Green Party. The results indicate no direct impact of the protests themselves. Instead, short-lived increases in Green Party support and environmental concerns emerge immediately before the protests, concentrated among politically informed respondents and accompanied by increased urban-rural polarization. These findings suggest that external news events likely shape both local attitudes and local propensities to protest.

Keywords: Protests, Environment, Climate Change, News, Information, Politics
JEL Codes: D72, D74, L82, Q58, Z13

1 Introduction

Attention is critical for social movements to shape the public debate, but difficult to attract even for large movements. Gethin and Pons (2024) document that most major protests in recent US history have failed to shift public opinion, pointing to a lack of media coverage as a likely explanation. One possible solution for movements seeking greater exposure is to disrupt everyday life, thereby forcing public attention and sparking controversy. Prominent recent examples include the yellow vests protests in France (2018–2020), the convoy protests in Canada (2020), and the *Last Generation* protests in Germany (2022–2023), which all attracted significant media coverage by, among other activities, repeatedly blocking major roads and disrupting traffic.

The effectiveness of these disruptive tactics in building broader support for environmental policies hinges on a delicate trade-off. On one hand, disruptive protests generate negative sentiment, potentially increasing visibility of climate issues and raising support for climate action. On the other hand, if the public perceives these protests as too extreme, these actions may backfire, undermining support not only for the movement itself but also for ideologically aligned groups. Despite the prevalence and public debate, empirical evidence on the actual impacts of disruptive protests remains scarce.

This paper addresses this gap by analyzing the short-term effects of protests carried out by the *Last Generation*, a German environmental activist group known for street blockades and covering structures of public interest with orange paint (e.g., Schuetze, 2023). I combine detailed data on protest timing and location with responses in a large-scale survey to identify causal effects on attitudes toward environmental policies and support for the Green Party.

I combine two main datasets. First, I use protest data from the *Armed Conflict Location & Event Data (ACLED)* project, which systematically records details of protests based on news reporting (Raleigh et al., 2023). The dataset provides information on date and city-level location for all protests attributed to the *Last Generation* in Germany from early 2022 through the beginning of 2024, when the group announced it would cease street blockades (Last Generation, 2024). Second, the *German Socio-Economic Panel (SOEP)* is a nationally representative longitudinal survey that covers approximately 15,000 private households per year, with interviews spread throughout most of the calendar year. My analysis focuses on two types of questions: Respondents' preferred political party, and their degree of worry about various topics, notably climate conservation and environmental protection, measured on a four-point scale.

The empirical analysis leverages variation in protest location and timing, as well as the timing of survey interviews around protests. I implement a stacked event-study design that compares survey respondents interviewed shortly before and after protests in treated city commuting zones, relative to respondents in commuting zones without protests within the same time window. I restrict the analysis to isolated protest events, ensuring that treated localities experience no other protests in the preceding or subsequent 14 days. City–stack fixed effects absorb attitudes specific to localities in a given time window, while day–stack fixed effects control for nationwide high-frequency shocks that may be correlated with protests. Identification of the average causal effect of protests on environmental attitudes relies on two key assumptions. First, survey interview timing must be as good as random within the 28-day windows around protests, conditional on observable controls. Second, protest occurrence and timing must be uncorrelated with other local events affecting attitudes. I verify the former via balance checks on observable individual characteristics, while leveraging the high frequency of observations to assess the latter.

I find no statistically significant effect of protests themselves on the expressed preference for the Green Party or worry about climate conservation and environmental protection. This result is robust to different choices of fixed effects, geographical aggregation, and the length of time windows around protests. However, the estimates are rather imprecise due to low overlap between the timing of protests and survey responses.

Additionally, I find evidence for shifts in environmental attitudes *prior to* protests. The event studies reveal increases in Green Party support and environmental concerns around one to two days before protests take place, which decrease over the following three days. Moreover, heterogeneity analyses show that this pre-protest increase is entirely driven by individuals who consume news daily and have previously expressed a party preference. Further, this phenomenon coincides with increased polarization: Urban respondents become more supportive of environmental positions, while rural respondents become less so.

The observed shifts in attitudes prior to protests suggest that external events likely influence both local attitudes toward environmental issues and local propensities to engage in protest. The concentration of these associations among politically informed individuals points toward media coverage as a potential driver. Nationwide news stories or policy debates around environmental issues may resonate more strongly or receive heightened attention in certain locations, such as student or politically left-leaning cities, thereby triggering simultaneous shifts in both opinions and activism. Moreover, the observed polarization between urban and rural respondents may reflect differences in news exposure or interpretation, shaped by regional media coverage or underlying beliefs and circumstances.

To explore these hypotheses more deeply, further research could benefit from extending the analysis using the 2023 wave of the *SOEP*, once available. This data might provide more precise estimates due to better temporal overlap with *Last Generation* protests. Additionally, to better understand the underlying mechanisms, one could instrument key variables of protest effectiveness, such as media attention using news pressure, or disruptiveness through weather-driven variation in traffic conditions.

This project contributes to a mature literature across the social sciences on the effects of political protests (e.g., Elsbach and Sutton, 1992; Meyer, 2004; Robinson and Acemoglu, 2012; Aidt and Franck, 2015; Battaglini et al., 2020; Tarrow, 2022; Cantoni et al., 2023). Scholars have studied the incentives of social movements for choosing violent versus non-violent tactics (Stephan and Chenoweth, 2008; Cornell and Grimes, 2015; Wang and Piazza, 2016; Simpson et al., 2018; Orazani and Leidner, 2019). While recent survey experiments have found mixed effects of disruptive protests on attitudes towards protesters and their demands (Feinberg et al., 2020; Saldivia Gonzatti et al., 2023; Menzies et al., 2023), Kountouris and Williams (2023) find no effects of street blockades by environmental activists in England. Instead, Fabel et al. (2025) find that participation rates in peaceful but controversial *Fridays-for-Future* protests in Germany were associated with higher vote shares for the Green Party. As one of the key mechanisms, they point toward increased media presence of the Green Party and environmental topics. I contribute to this literature by providing quasi-experimental, causal evidence on the effects of non-violent, disruptive protests. By doing so, I test whether this disruptive strategy optimized for media attention proves more effective in persuading the public.

2 Data

I combine two main data sources that allow tracing protests by the *Last Generation* as well as environmental attitudes at high geographical and temporal resolution.

2.1 Protests

I use publicly available information on protests from the *Armed Conflict Location & Event Data (ACLED)* project (Raleigh et al., 2023), which records events of armed and non-armed political conflicts worldwide. I restrict this sample to events in Germany throughout 2022 and 2023 that name the *Last Generation* as the main actor. The result are 619 protest events, defined as unique combinations of date, city, and (where available) sub-location. Because several protests may occur in different spots of the same city on a single day, these events collapse to 494 distinct city–date pairs.

Figure 3.1: Timing of protests

(a) Number of protests by month

(b) Share of protests by day of week

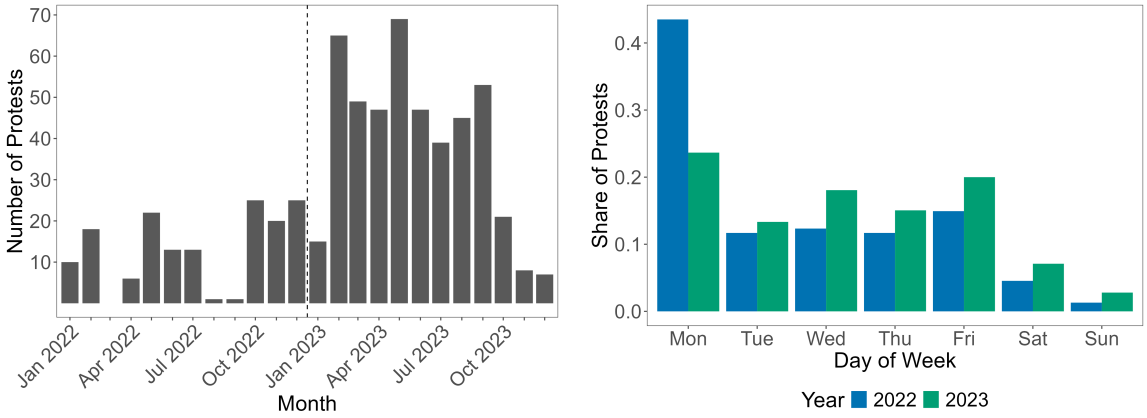

Notes: Panel (a): Monthly number of protest events in Germany involving the *Last Generation* as recorded by *Armed Conflict Location and Events Data (ACLED)*. Panel (b): Share of *Last Generation* protests falling on specific day of week, separately for 2022 and 2023.

Protest activity by the *Last Generation* started in 2022 and intensified over time. Figure 3.1a shows the number of events per month – 154 in 2022 and 465 in 2023 – with most of them between February and September 2023. These protests usually took place on weekdays. Figure 3.1b displays the share of events by day of week, revealing that in both 2022 and 2023, more than eighty percent of protests happened between Monday and Friday. This pattern reflects the unconventional strategy of the Last Generation, which targets rush our traffic to cause maximal disruption.

While the *Last Generation* was by far most active in Berlin, many other cities also experienced considerable protest activity. Figure A3.1 details the geographical location of protests throughout 2022 and 2023, which spanned over 78 cities. Figure A3.2 reveals that Berlin alone accounts for 53 protest days in 2022 and 94 in 2023. Munich, Dresden, Hamburg, Köln, and Leipzig each faced at least ten protest days in 2023.

Table 3.1 reports summary statistics, revealing several additional details. First, most protests outside of Berlin occurred in West Germany. Second, the median recorded number of protesters is 10, reflecting the strategy of creating large (media) attention with few participants. Third, keywords appearing in *ACLED*'s standardized event descriptions explicitly refer to most events as a *blockade* or *disruptive*. In particular, I label an event a *blockade* if the description mentions both (i) a verb related to interruption ("block", "interrupt", "glue", etc.) and (ii) a noun related to transport infrastructure ("highway", "bridge", "intersection", etc.). I define an event as *disruptive* if it is a blockade *or* if it refers to the common tactic of vandalizing surfaces with paint ("spray", "paint", "orange", etc.). Depending on the year, between 74% and 90% of all events meet the *disruptive* criterion.

Table 3.1: Summary statistics for protests of Last Generation

	Mean	Median	SD	Min	Max	N
2022						
Location: Berlin	0.50	0	0.50	0	1	154
Location: West Germany (except Berlin)	0.42	0	0.50	0	1	154
Location: East Germany (except Berlin)	0.08	0	0.27	0	1	154
Keywords: Disruptive	0.90	1	0.30	0	1	154
Keywords: Blockade	0.84	1	0.36	0	1	154
Number of Protesters	17.32	10	18.26	3	100	93
2023						
Location: Berlin	0.38	0	0.49	0	1	465
Location: West Germany (except Berlin)	0.49	0	0.50	0	1	465
Location: East Germany (except Berlin)	0.12	0	0.33	0	1	465
Keywords: Disruptive	0.74	1	0.44	0	1	465
Keywords: Blockade	0.69	1	0.46	0	1	465
Number of Protesters	50.06	10	128.74	3	1000	263

Notes: Summary statistics for protest events in Germany in 2022 and 2023 involving the *Last Generation*, as recorded by *Armed Conflict Location and Events Data (ACLED)*.

Media outlets typically covered the protest on the day they unfolded. Traffic disruptions often triggered immediate news reports, which were followed by segments in evening television, next day's print editions, and more extensive analytical pieces in weekend sections. The *Last Generation* usually did not advertise protests in advance, ensuring that police would take longer to disperse the blockades. Only large, multi-day campaigns were announced beforehand and received occasional ex-ante coverage.

In January 2024, the *Last Generation* announced that it would step away from blockading traffic, and instead focus on other forms of civil unrest (Last Generation, 2024).

2.2 German Socio-Economic Panel

The *German Socio-Economic Panel (SOEP)*[1] is a nationally representative longitudinal survey of private households conducted by the *German Institute for Economic Research (DIW)*. I use wave 39, the most recent published version at the time of writing, which includes responses up until the 2022 round of interviews. As the *Last Generation* protests started only that year, I restrict the data to interviews conducted in 2022. The resulting sample contains 28,606 adult respondents from 21,046 households.

[1] Socio-Economic Panel (SOEP).

Figure 3.2: Number of SOEP interviews per month in 2022

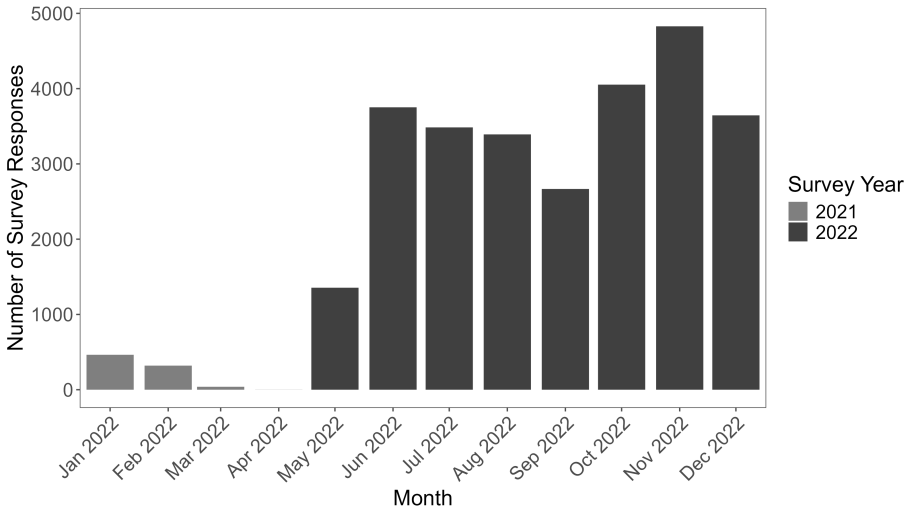

Notes: Number of individual survey responses collected per month of 2022 by the *German Socio-Economic Panel (SOEP)*. Bars in light grey refer to the survey wave of 2021, while dark grey denotes the 2022 wave.

Interviews are conducted on a continuous basis throughout most of the year. Figure 3.2 shows that in 2022, a small number of interviews were recorded between January and March as part of the 2021 survey wave. Instead, interviews for the 2022 wave start in May, with more than 3,000 responses per month from June onward.

The survey offers a rich set of variables on demographics and attitudes. Table 3.2 reports summary statistics for individual characteristics as well as my two main outcomes: First, respondents name their preferred political party, where around half express no preference at all. Interestingly, the right-populist *Alternative for Germany (AfD)* seems to be underrepresented: Merely 2 percent prefer the *AfD*, compared to their 10.4 percent vote share in the 2021 federal election, potentially reflecting sampling bias or social desirability concerns. Second, respondents rate how much they worry about several political issues, including climate change and environmental protection. On a four-point scale, around half report the highest degree of concern for both topics.

I link survey observations to protest exposure using respondents' locations, which are in principle available on street block level.[2] I aggregate protests and respondents by 96 *Raumordnungsregionen*, which correspond to commuting zones around economic centers. In the following, I will refer to these commuting zones as "cities" for brevity. Alternatively, I aggregate to 38 NUTS-2 administrative units.

[2]Restricted micro data on geographical locations of survey respondents is available for on-site access at the *DIW*.

Table 3.2: Summary statistics for German Socioeconomic Panel, 2022

	Mean	SD	Min	Max	N
Demographics					
Female	0.51	0.50	0	1	28,022
Birth Year	1971	16.95	1921	2004	28,042
College Degree	0.29	0.48	0	1	24,595
Vocational Degree	0.54	0.50	0	1	24,595
Employment: Full-Time	0.41	0.50	0	1	27,940
Employment: Part-Time	0.21	0.41	0	1	27,940
Employment: None	0.37	0.46	0	1	27,940
Household Income per Month (€)	3,650.19	2,397.89	0	15,000	27,877
Reads Newspaper Daily	0.50	0.50	0	1	14,813
Party preference					
Has Party Preference	0.49	0.50	0	1	25,770
CDU/CSU	0.14	0.32	0	1	25,708
Green	0.13	0.34	0	1	25,708
SPD	0.11	0.29	0	1	25,708
FDP	0.03	0.16	0	1	25,708
AfD	0.02	0.15	0	1	25,708
Left	0.02	0.16	0	1	25,708
Worry "much" about topic					
Climate Conservation	0.51	0.50	0	1	25,734
Environmental Protection	0.46	0.50	0	1	25,734

Notes: Summary statistics of selected variables in the *German Socio-Economic Panel (SOEP)* among interviews conducted in 2022. Household income is winsorized at the 99^{th} percentile. Worry about topics is rated on a four-point scale, with "much" representing the highest value. Survey weights are included.

3 Empirical strategy

I estimate the short-term causal effect of protests by exploiting variation in the exact timing and location of protests. The empirical design is a stacked event study that compares respondents interviewed shortly before a protest in a given city to those interviewed shortly after, relative to respondents in cities without protest during the same time window.

3.1 Design

To isolate the effect of individual protests, I restrict the analysis to protest-city pairs with no other recorded protest in the same city during the 14 days before or after the event. I further retain only those cities whose $\pm$14-day window contains at least one pre- and one post-interview in both treated and untreated locations. Applying these restrictions yields 29 protest–city pairs in 2022 and 1,068 unique *SOEP* observations from respondents in treated cities.[3]

The stacked event study estimates a difference-in-difference regression for each protest event, and averages ("stacks") the resulting estimates. Each difference-in-difference regression is estimated in a separate subset of the data in the $\pm$14-day window around the protest ("stack"), consisting of i) survey respondents in the treated locality, and ii) survey respondents in control localities that did not experience any protest in the time window.

I employ the following specification:

$$y_{\,i,c,t,s} = \alpha_{c,s} + \gamma_{t,s} + \sum_{\tau=-14,\tau\neq-1}^{14} \beta_\tau \; I(Days\ since\ Protest_{\,c,t} = \tau) + \delta' \, X_{i,c,t} + \varepsilon_{i,c,t,s} \tag{3.1}$$

$y_{\,i,c,t,s}$ is a survey outcome for participant i residing in city c, surveyed on date t, in stack s. The survey outcome is an indicator, either for supporting the Green Party or for worrying "much" about climate conservation and environmental protection, respectively. The city–stack fixed effect $\alpha_{c,s}$ absorbs time-invariant heterogeneity across cities *within each protest event*. In other words, it ensures that the treatment effect is obtained by comparing the change of respondents' attitudes within treated cities to the change of attitudes within control cities, all within the 28-day window of each protest. Therefore, this fixed effect

[3] Unfortunately, the temporal overlap between protests and survey responses is modest in 2022. Roughly half of all protests precede May, which is when most *SOEP* interviews begin. Moreover, protest activity largely pause in August and September. The amount of protests as well as the temporal overlap are higher in 2023.

controls for macro trends in attitudes within city, for example due to protest history. The date–stack fixed effect $\gamma_{t,s}$ captures daily shocks to attitudes that are common across all localities, for example due to common news shocks. Allowing this fixed effect to vary across stacks absorbs differential shocks based on the specific subset of localities used for estimating the stack-specific effect of a given protest. $\mathbf{X_{i,c,t}}$ is a vector of individual demographic control variables in the form of indicators for bins of age (6 groups), sex (2), education (7), household income (6), and employment status (4). These variables may capture different propensities across individuals to prefer a certain party or worry about the environment. Finally, $Days\ since\ Protest_{c,t,s}$ denotes the number of days since the protest of stack s took place in city c, if applicable. The coefficients β_τ represent event study estimates for the average effect of the protest on exposed localities, across all protests in the sample. I estimate Equation (3.1) by ordinary least squares (OLS), apply *SOEP* sampling weights, and cluster standard errors by city.

3.2 Identification

Identification of β_τ rests on two main assumptions. First, conditional on covariates, the *SOEP* interview date needs to be as-good-as random with respect to attitudes. Table A3.1 confirms that respondents surveyed in the 14 days before protests largely resemble those surveyed in the 14 days after protests along observable characteristics. I control for these variables in all specifications. Second, the exact protest day within the 28-day window needs to be unrelated with concurrent shocks to attitudes that are specific to the the treated city. While not directly testable, the high frequency of the design helps rule out confounders by assessing whether the timing of observed effects aligns plausibly with protest events.

On a technical note, Wing et al. (2024) shows that the OLS estimator imposes specific weights when aggregating stack-specific estimates for the treatment effect. In result, the vector of pooled coefficients β_τ does not necessarily recover any meaningful aggregation of the stack-specific estimates $\{\beta_\tau^{\mathbf{s}}\}_{\mathbf{s}\in\text{Stacks}}$. In a robustness check, I verify that in this case the results are virtually identical when using the proposed bias-corrected weights.

4 Results

4.1 Main result

I find no statistically significant effect of protests on attitudes towards the Green Party as well as concerns about climate conservation and environmental protection. Table 3.3 reports aggregate effects over the two weeks following protests, obtained via the static version of Equation (3.1). For all three outcomes, the coefficients are positive, but insignificant with p-values below 0.50.

Table 3.3: Effect of Protests on Green Party Support and Topic Concerns

Dependent Variable:	Support Party:	Worry Much About:	
	Green	Climate	Environment
	(1)	(2)	(3)
Protest x Post	0.013	0.032	0.027
	(0.031)	(0.046)	(0.047)
Dep. Var. Mean	0.13	0.51	0.46
Individual Controls	✓	✓	✓
Protest-ID x City FE	✓	✓	✓
Protest-ID x Days since Protest FE	✓	✓	✓
Day of Week FE	✓	✓	✓
Protests	29	29	29
Treated Observations	1,068	1,068	1,068
Control Observations	38,740	38,740	38,740

Notes: Coefficients for static (pre-post) version of Equation (3.1): Stacked event study around city-by-day protest events, comparing cities with exactly one protest in 28-day window to cities without any protests. Outcomes are indicators for preferring the Green Party (Column 1), and worrying about climate conservation and environmental protection to highest degree on a four-point scale, respectively (Columns 2–3). All specifications include individual controls for age, sex, education, household income, and employment status, as well as fixed effects for protest-by-city, protest-by-days-since-protest, and day of week. Standard errors in parentheses are clustered by city. Survey weights are included. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The event studies confirm that the protests themselves do not affect these environmental attitudes, even over shorter time frames. Figure 3.3 plots event study estimates for Green Party support obtained from estimating Equation (3.1). To facilitate interpretation, I choose the second week before the protest as the reference period, instead of the more commonly used day before. In other words, I omit days -14 to -8 relative to the protest rather than -1. This choice yields a more precise baseline by averaging more estimates, and reports the standard error for the -1 coefficient, allowing to better interpret the point estimate. The estimates reveal no clear effect. Both pre- and post-treatment coefficients largely fluctuate around zero and show neither a consistent trend nor an increase after the protest.[4] The same

[4]Because the error bars report pointwise rather than uniform confidence intervals, isolated significant

is true for for concerns about climate change (Figure A3.3) and environmental protection (Figure A3.4), which display qualitatively similar patterns.

However, all three figures show an increase in environmental attitudes on the day *before* the protest. In each case, a single (marginally significant) spike stands out, vanishing over the following two to three days. Thus, while the protests studied in this analysis by themselves seem not to affect attitudes, they are appear with higher likelihood in cities that experience a *prior* increase in Green Party support as well as concerns about climate and environment.

Figure 3.3: Effect of protest on probability of preferring Green Party

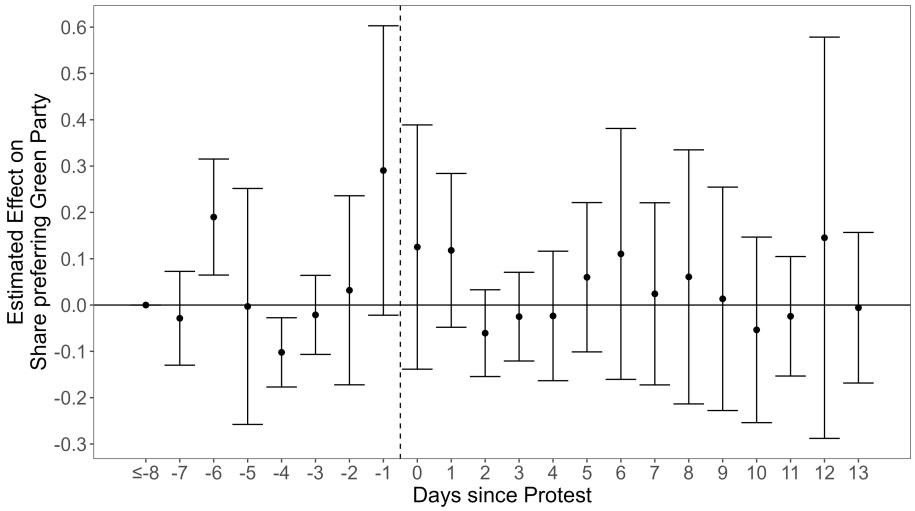

Notes: Daily estimates for Green Party support around *Last Generation* protests in 2022, based on Equation (3.1). Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents preferring the Green Party. The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

4.2 Robustness checks

The findings are robust to varying the design of the empirical analysis in several ways. Figure A3.5 reports the respective event studies for Green Party support. Panel (a) varies the degree of saturation with fixed effects. Estimates in orange represent a basic specification that includes only treatment status and event time. Green additionally controls for day of week, and blue replicates Equation (3.1), including city–stack and date–stack fixed effects. Point estimates and inference remain virtually unchanged. Panel (b) aggregates protests

coefficients do not necessarily imply robust effects.

and survey respondents on 38 administrative (NUTS-2) regions, instead of 98 city commuting zones, and Panel (d) chooses 14-day instead of 28-day windows around protests. In both cases, the main results remain unchanged. Panel (c) varies weighting. Estimates in orange report results without any sample weights, showing considerably less variation while retaining the spike on day -1. Estimates in green report baseline results obtained using only *SOEP* sample weights, while those in blue apply weights that multiply *SOEP* sample weights with Wing et al. (2024) bias-correcting weights. The correction makes little difference, suggesting that in this case, the aggregation performed by OLS produces close-to-correct inference.

4.3 Mechanism

Heterogeneity checks help clarify the results and suggest mechanisms driving these findings. I split the sample by individual characteristics and estimate Equation (3.1) in each subsample.

First, I consider variables that indicate to what degree respondents are politically informed. Figure 3.4 reports results for Green Party support, dividing respondents by whether they report reading news daily. It reveals that the positive pre-treatment coefficient is driven entirely by daily readers, starting two days before the protest. Instead, those who follow the news less frequently experience similar trends before, but show no discernible difference in Green Party support around protests. Figure 3.5 separates respondents by whether they have ever declared any party preference in previous years. Again, the uptick appears only among politically engaged individuals, and two days before the protest. Respondents without a past party preference show no change.

Figure 3.6 displays outcomes among urban and rural residents, respectively. Interestingly, while pre-trends are very similar, attitudes diverge considerably on the day before the protest: Respondents in urban areas show significantly increased preference for the Green Party, while those in rural areas show significantly lower preference, compared to their respective pre-period levels. Equivalent splits for the outcomes on environmental concerns (Figures A3.6–A3.8) largely confirm the results, but with less clarity.

Together, these findings reveal three insights. First, the transient shift in Green Party preference *precedes* the protests included in my sample. Second, these increases are driven by urban, highly informed individuals. Third, the protests are associated with short-term political polarization across residents of urban and rural regions.

Figure 3.4: Effect of protest on probability of preferring Green Party - by news consumption

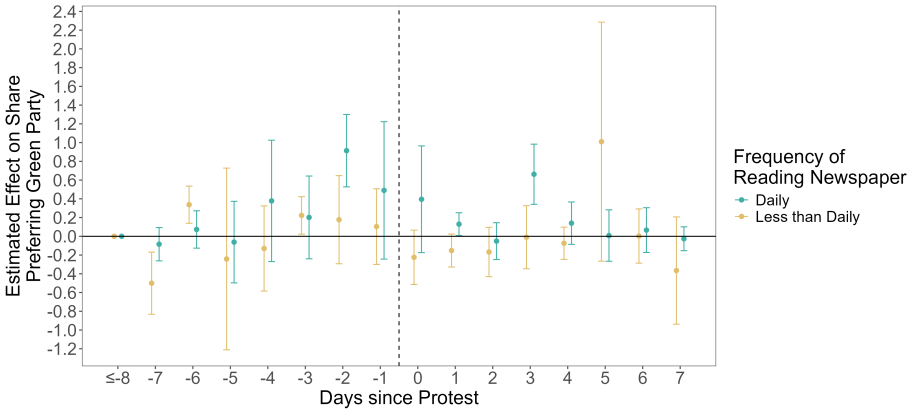

Notes: Daily estimates for Green Party support around *Last Generation* protests in 2022, based on Equation (3.1). Separate regressions for individuals reading news daily and less than daily, respectively. Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents preferring the Green Party. The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Figure 3.5: Effect of protest on probability of preferring Green Party - by having party preference in the past

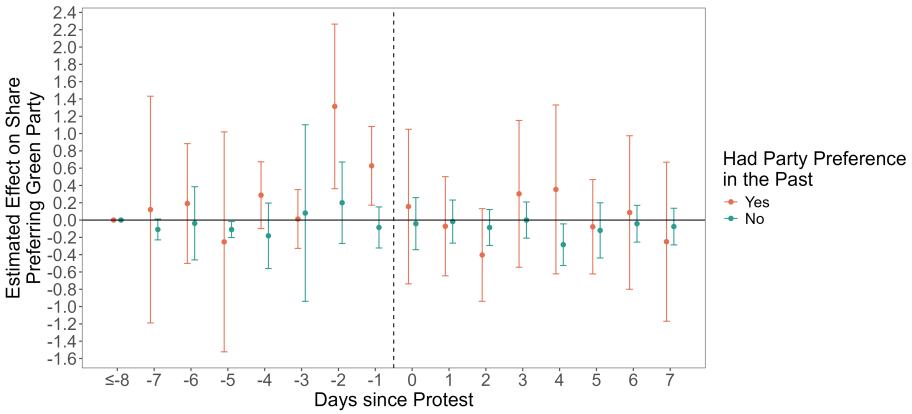

Notes: Daily estimates for Green Party support around *Last Generation* protests in 2022, based on Equation (3.1). Separate regressions by whether individuals report any party preference in the past. Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents preferring the Green Party. The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Figure 3.6: Effect of protest on probability of preferring Green Party - by urban/rural

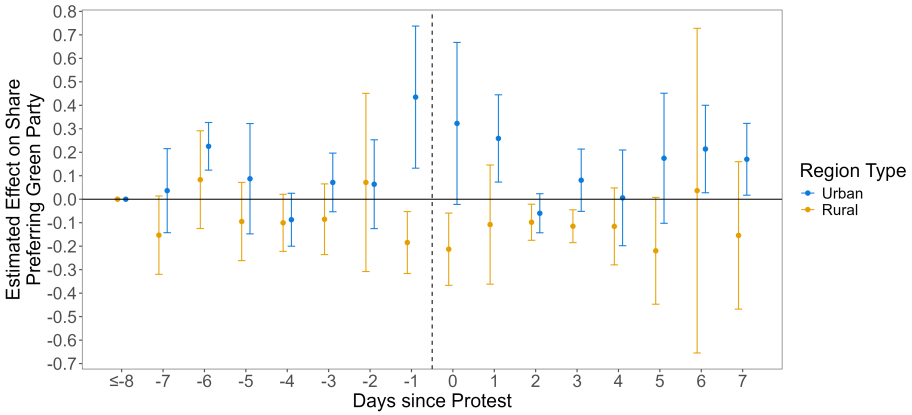

Notes: Daily estimates for Green Party support around *Last Generation* protests in 2022, based on Equation (3.1). Separate regressions for individuals residing in urban and rural areas, respectively. Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents preferring the Green Party. The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

4.4 Interpretation

The observation that attitudes shift before a protest and only among the most informed respondents may suggest that external news shocks related to environmental topics jointly drive environmental attitudes and increase the likelihood of protest. Two types of news shocks, such as stories or policy debates, are plausible. First, local news shocks might receive attention only in their respective city. Second, nation-wide shocks could receive disproportionate coverage, or resonate more strongly, in places with a higher latent propensity to mobilize, such as in student or otherwise left-leaning towns.[5]

Either scenario violates the exogeneity of protest timing: The empirical design no longer identifies the causal effect of protest on attitudes, but instead captures the correlation induced by the omitted news shock. Accordingly, the findings should not be interpreted as evidence for a treatment effect of the protest itself, but may hint at the existence of an information shock that is either received or interpreted differently among different populations.

[5] An additional potential source of endogeneity is non-response bias: Since survey response rates are usually low (around 9% in this case), results may be biased if the news or protest events under consideration change the set of individuals who participate in the survey. To produce the observed results, exposure to these events would need to increase the likelihood of Green Party voters to participate in the survey (in particular, among politically informed and urban participants), for example due to an increased sense of social responsibility.

5 Conclusion

This paper explores whether disruptive climate protests conducted 2022 by the *Last Generation* in Germany influenced environmental attitudes. I find little evidence that protests themselves had an effect. Instead, the timing of effects suggests that protests respond to external news shocks, which coincided with increased support for the Green Party as well as heightened worry about climate conservation and environmental protection. These effects emerge over a span of approximately three days and primarily affect politically informed individuals. Additionally, these events are associated with increased polarization, with urban and rural respondents exhibiting opposite attitudinal shifts.

A key limitation of the analysis are relatively imprecise estimates due to sparse temporal overlap between protest activity and survey interviews. Future research could benefit from extending the analysis to 2023, when protest timing aligns more closely with the *SOEP* fieldwork.

Such follow-up research should accomplish two key goals. First, it should pinpoint the types of news events that simultaneously trigger protests and attitude changes. Second, it could exploit additional variation to clarify mechanisms. For instance, researchers could instrument media attention using news pressure, for example using the *Global Database of Events, Language, and Tone (GDELT) Frontpage Graph*, which collects front-page content from news websites at high frequency. Alternatively, disruptiveness of protests could be instrumented by weather-induced traffic shocks: Rainy weather leads to increased car use, causing protests to generate larger disruptions, likely without significantly affecting protester turnout given the already substantial personal risks involved in protest participation.

Bibliography

Aidt, Toke S and Raphaël Franck, "Democratization under the threat of revolution: Evidence from the Great Reform Act of 1832," *Econometrica*, 2015, *83* (2), 505–547.

Battaglini, Marco, Rebecca B Morton, and Eleonora Patacchini, "Social groups and the effectiveness of protests," Technical Report, National Bureau of Economic Research 2020.

Cantoni, Davide, Andrew Kao, David Y Yang, and Noam Yuchtman, "Protests," Technical Report, National Bureau of Economic Research 2023.

Cornell, Agnes and Marcia Grimes, "Institutions as incentives for civic action: Bureaucratic structures, civil society, and disruptive protests," *The Journal of Politics*, 2015, *77* (3), 664–678.

Elsbach, Kimberly D and Robert I Sutton, "Acquiring organizational legitimacy through illegitimate actions: A marriage of institutional and impression management theories," *Academy of management Journal*, 1992, *35* (4), 699–738.

Fabel, Marc, Matthias Flückiger, Markus Ludwig, Helmut Rainer, Maria Waldinger, and Sebastian Wichert, "The relationship between the youth-led Fridays for Future climate movement and voting, politician and media behaviour in Germany," *Nature Human Behaviour*, 2025, *9* (3), 481–495.

Feinberg, Matthew, Robb Willer, and Chloe Kovacheff, "The activist's dilemma: Extreme protest actions reduce popular support for social movements.," *Journal of personality and social psychology*, 2020, *119* (5), 1086.

Gethin, Amory and Vincent Pons, "Social movements and public opinion in the united states," Technical Report, National Bureau of Economic Research 2024.

Gonzatti, Daniel Saldivia, Sophia Hunger, and Swen Hutter, "Analysebericht zur Studie "Environmental Protest Effects on Public Opinion: Experimental Evidence from Germany"," *OSF Preprints*, June 2023.

Kountouris, Yiannis and Eleri Williams, "Do protests influence environmental attitudes? Evidence from Extinction Rebellion," *Environmental Research Communications*, 2023, *5* (1), 011003.

Last Generation, "Strategie für 2024," 2024. Press release, January 24, 2024, last accessed June 24, 2025.

Menzies, Rachel E, Matthew B Ruby, and Ilan Dar-Nimrod, "The vegan dilemma: Do peaceful protests worsen attitudes to veganism?," *Appetite*, 2023, *186*, 106555.

Meyer, David S, "Protest and political opportunities," *Annu. Rev. Sociol.*, 2004, *30*, 125–145.

Orazani, S Nima and Bernhard Leidner, "The power of nonviolence: Confirming and explaining the success of nonviolent (rather than violent) political movements," *European Journal of Social Psychology*, 2019, *49* (4), 688–704.

Raleigh, Clionadh, Roudabeh Kishi, and Andrew Linke, "Political instability patterns are obscured by conflict dataset scope conditions, sources, and coding choices," *Humanities and Social Sciences Communications*, 2023, *10*, 74.

Robinson, James A and Daron Acemoglu, *Why nations fail: The origins of power, prosperity and poverty*, Profile London, 2012.

Schuetze, Christopher F., "Climate Activists Find a Way to Get Germany's Attention: Stop Traffic," *The New York Times*, 1 2023.

Simpson, Brent, Robb Willer, and Matthew Feinberg, "Does violent protest backfire? Testing a theory of public reactions to activist violence," *Socius*, 2018, *4*, 2378023118803189.

Socio-Economic Panel (SOEP), "Socio-Economic Panel (SOEP), Version 39, Data for Years 1984–2022 (SOEP-Core v39, EU Edition)," 2024.

Stephan, Maria J and Erica Chenoweth, "Why civil resistance works: The strategic logic of nonviolent conflict," *International security*, 2008, *33* (1), 7–44.

Tarrow, Sidney, *Power in movement*, Cambridge university press, 2022.

Wang, Dan J and Alessandro Piazza, "The use of disruptive tactics in protest as a trade-off: the role of social movement claims," *Social forces*, 2016, *94* (4), 1675–1710.

Wing, Coady, Seth M Freedman, and Alex Hollingsworth, "Stacked difference-in-differences," Technical Report, National Bureau of Economic Research 2024.

Appendix

Figure A3.1: Number of days with protests by city, 2022–2023

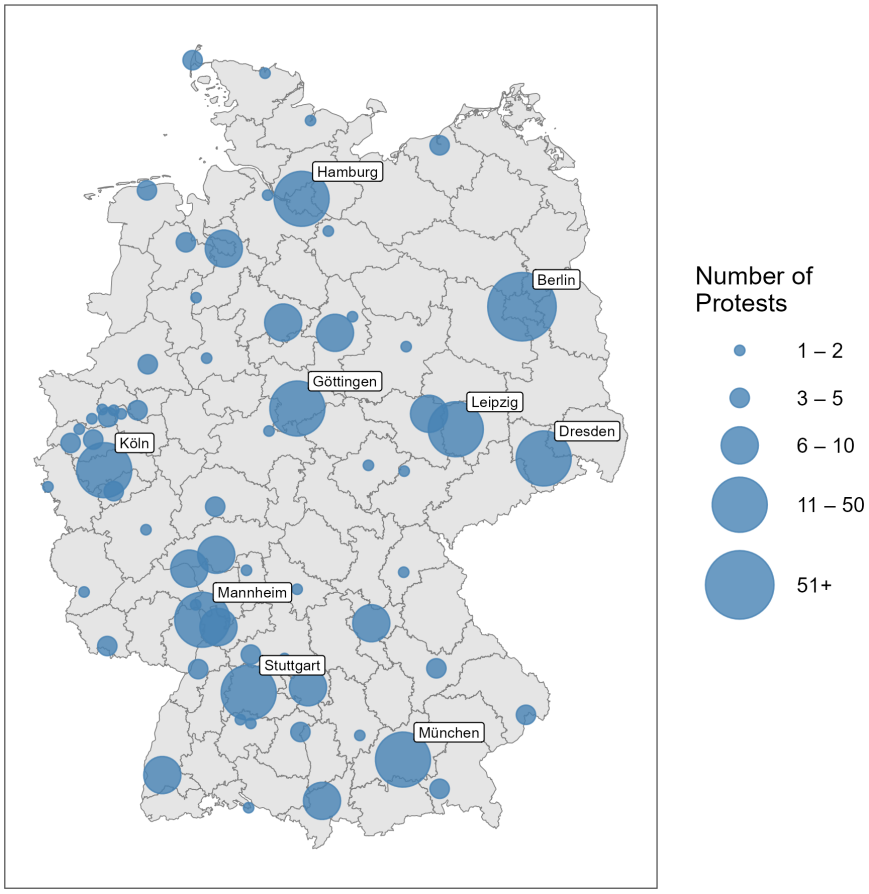

Notes: Number of days with *Last Generation* protest per city throughout 2022 and 2023. Areas indicate the 96 city commuting zones (*Raumordungsregionen*).

Figure A3.2: Number of days with protest for top 10 cities

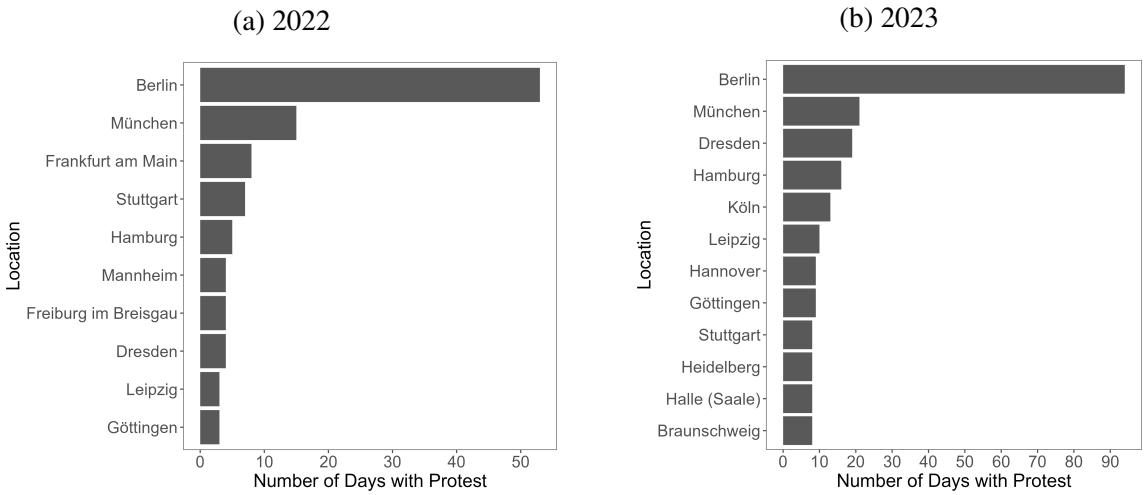

Notes: Number of days with *Last Generation* protest per city, separately for 2022 and 2023, among 10 cities with most protest days.

Figure A3.3: Effect of protest on worrying about climate conservation

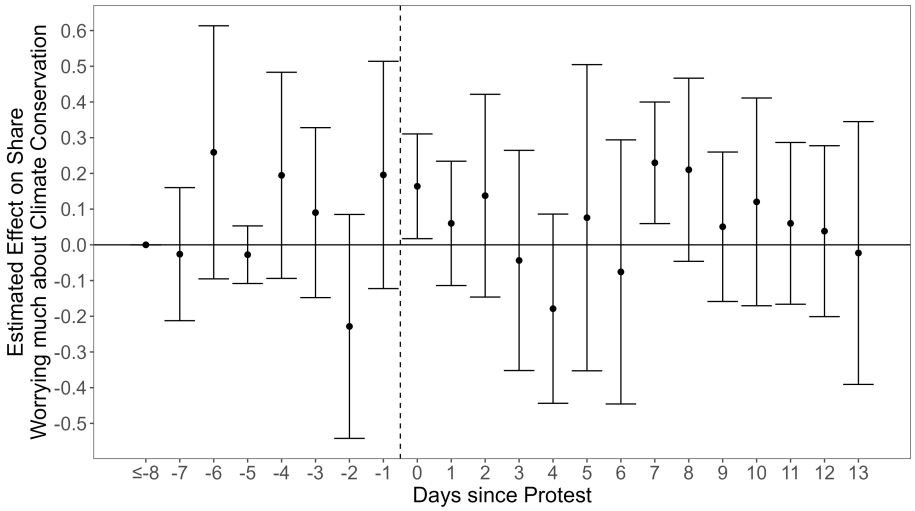

Notes: Daily estimates for reported worry about climate conservation around *Last Generation* protests in 2022, based on Equation (3.1). Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents reporting maximum worrying about climate conservation on a four-point scale. The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Table A3.1: Covariate balance: Stacked event studies

	Before Protest		After Protest			P-Value	
Variable	**Mean**	**SD**	**Mean**	**SD**	**Mean Diff.**	**Conventional**	**Bonferroni**
Sex							
Male	0.50	0.50	0.50	0.50	-0.00	0.923	1
Birth year							
1950–1959	0.14	0.35	0.14	0.35	-0.00	0.912	1
1960–1969	0.21	0.40	0.19	0.39	-0.01	0.006^{***}	0.138
1970–1979	0.15	0.36	0.16	0.36	0.01	0.199	1
1980–1989	0.17	0.38	0.18	0.38	0.00	0.734	1
1990+	0.23	0.42	0.23	0.42	-0.00	0.573	1
Education							
None	0.02	0.14	0.02	0.15	0.00	0.435	1
Hauptschule	0.04	0.18	0.04	0.19	0.00	0.836	1
Realschule	0.02	0.13	0.02	0.14	0.00	0.253	1
Vocational	0.43	0.50	0.41	0.49	-0.02	0.003^{***}	0.069^{*}
Abitur	0.05	0.21	0.04	0.20	-0.00	0.224	1
College	0.23	0.42	0.22	0.42	-0.01	0.194	1
Other	0.02	0.15	0.02	0.15	0.00	0.47	1
Employment							
None	0.31	0.46	0.31	0.46	0.00	0.612	1
Training	0.02	0.15	0.02	0.15	0.00	0.208	1
Part-time	0.23	0.42	0.21	0.41	-0.01	0.011^{**}	0.253
Full-time	0.43	0.50	0.45	0.50	0.01	0.103	1
Household income							
< 1500	0.11	0.32	0.12	0.33	0.01	0.044^{**}	1
1500–2500	0.20	0.40	0.20	0.40	0.01	0.304	1
2500–3500	0.20	0.40	0.20	0.40	-0.00	0.345	1
3500–4500	0.20	0.40	0.18	0.38	-0.02	0.001^{***}	0.023^{**}
4500–5500	0.14	0.34	0.14	0.34	0.00	0.835	1
> 5500	0.16	0.36	0.16	0.37	0.01	0.177	1

Notes: Mean difference of individual characteristics among survey respondents interviewed in two weeks before versus two weeks after protests, within the stacked event study estimation sample. Conventional as well as Bonferroni-adjusted p-values are reported. Survey weights are included. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Figure A3.4: Effect of protest on worrying about environmental protection

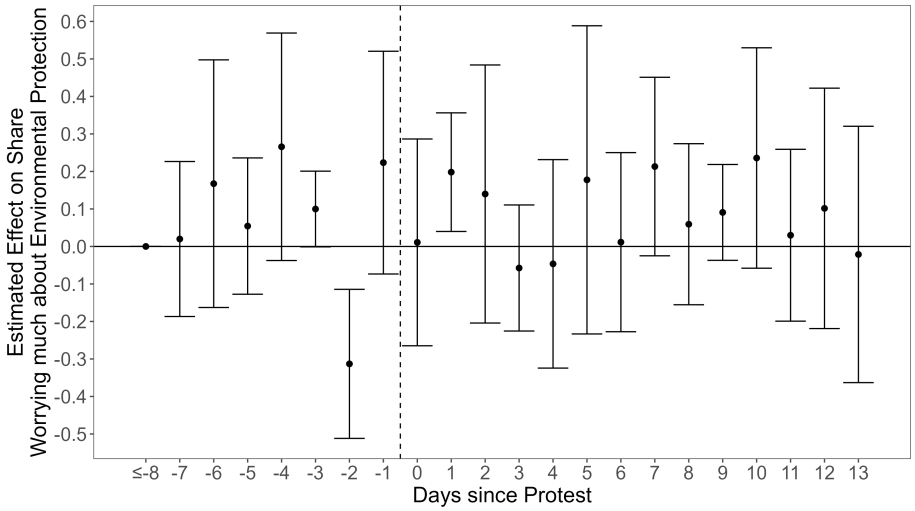

Notes: Daily estimates for reported worry about environmental protection around *Last Generation* protests in 2022, based on Equation (3.1). Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents reporting maximum worrying about environmental protection on a four-point scale. The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Figure A3.5: Effect of protest on probability of worrying - by urban/rural

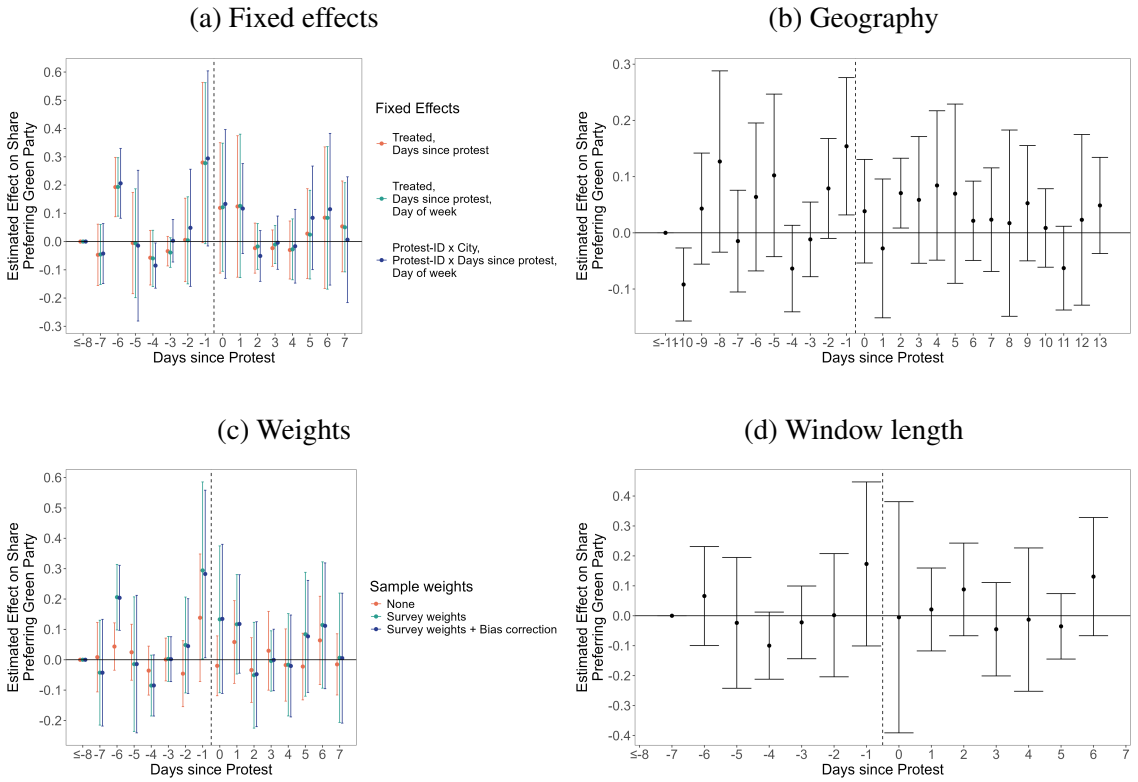

Notes: Daily estimates for Green Party support around *Last Generation* protests in 2022, based on stacked event studies following Equation (3.1). Panel (a) varies fixed effects. Panel (b) aggregates protests and survey respondents on the 38 NUTS-2 administrative regions instead of 96 city commuting zones. Panel (c) shows sensitivity to survey weights as well as Wing et al. (2024) bias-correcting weights for stacked event studies. Panel (d) focuses on unique protests within 14-day instead of 28-day windows.
The baseline specification reports coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for survey respondents preferring the Green Party. The treatment variable measures days relative to protest. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Figure A3.6: Effect of protest on worrying - by news consumption

(a) Climate conservation

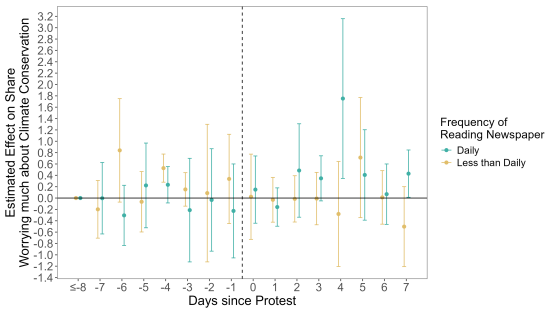

(b) Environmental protection

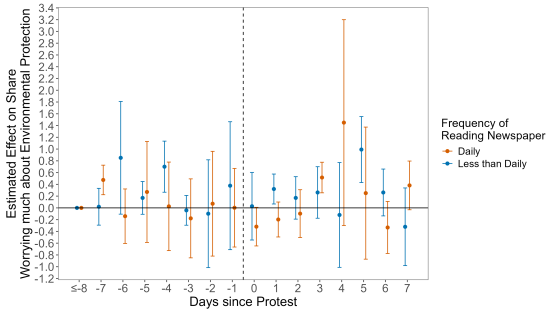

Notes: Daily estimates for worrying about environmental topics around *Last Generation* protests in 2022, based on Equation (3.1). Separate regressions for individuals reading news daily and less than daily, respectively. Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for maximal worrying on a four-point scale, about climate conservation in Panel (a) and environmental protection in Panel (b). The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Figure A3.7: Effect of protest on worrying - by having party preference in the past

(a) Climate conservation

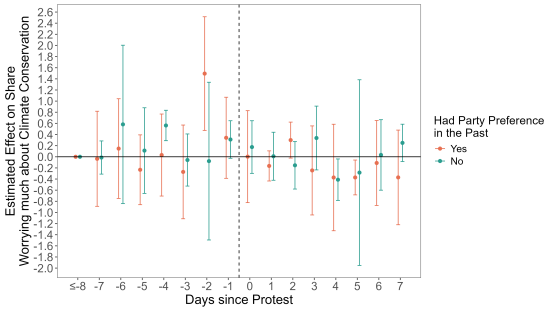

(b) Environmental protection

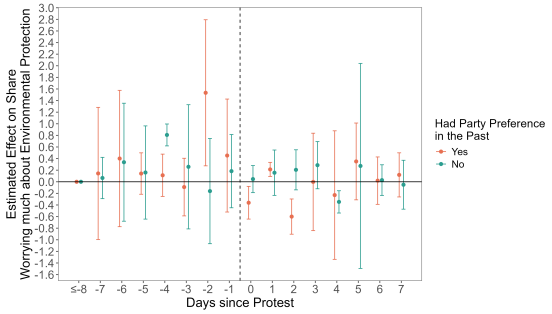

Notes: Daily estimates for worrying about environmental topics around *Last Generation* protests in 2022, based on Equation (3.1). Separate regressions by whether individuals report any party preference in the past. Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for maximal worrying on a four-point scale, about climate conservation in Panel (a) and environmental protection in Panel (b). The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.

Figure A3.8: Effect of protest on probability of worrying - by urban/rural

(a) Climate conservation

(b) Environmental protection

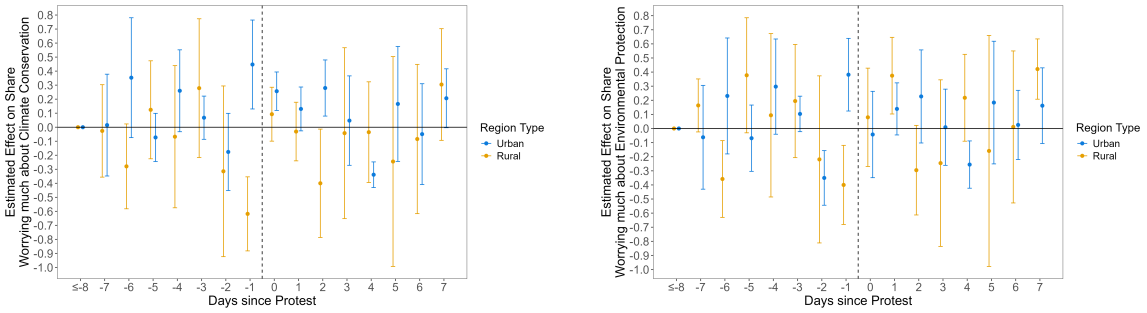

Notes: Daily estimates for worrying about environmental topics around *Last Generation* protests in 2022, based on Equation (3.1). Separate regressions for individuals residing in urban and rural areas, respectively. Coefficients from a stacked event study that compares city commuting zones experiencing exactly one protest within a 28-day window to those without any protests in the same period. The outcome is an indicator for maximal worrying on a four-point scale, about climate conservation in Panel (a) and environmental protection in Panel (b). The treatment variable measures days relative to protest, omitting days -14 to -8 as the reference period. Specification includes fixed effects for city-by-stack and day-by-stack, along with individual controls for age, sex, education, household income, and employment status. Original estimation sample includes 29 city-protest pairs, yielding 54,797 observations including 1,068 unique respondents in treated cities. Error bars denote pointwise 95 percent confidence intervals, clustered at the city level.