

Paywalls*

Desmond Ang[†]

Avinash Moorthy[‡]

Stephen Pettigrew[§]

Benjamin Schneer[¶]

July 10, 2026

Abstract

Over the past two decades, local newspapers have overwhelmingly adopted subscription paywalls, restricting access to those willing and able to pay. Combining hand-collected paywall data for the universe of U.S. daily newspapers with household-level web-browsing data, we document the distributional consequences of this transition for news consumption, political behavior, and civic knowledge. Paywalls reduce local readership and skew the remaining audience wealthier. Rather than seeking out alternative sources of area news, displaced readers substitute toward national outlets, which lean more liberal than local papers. This reallocation drives voters leftward and reduces local knowledge, with the largest effects among low-income households.

*Our research has been supported by the Ash Center for Democratic Governance and Innovation. We thank Matt Baum, Stefano DellaVigna, Matt Gentzkow, Jim Snyder, and seminar participants at the Harvard Kennedy School, Yale University, and the Elections, Public Opinion, and Voting Behavior conference for helpful comments. We are grateful to Rostom Rainauli and Abraham Sandoval for outstanding research assistance.

[†]Harvard Kennedy School, desmond_ang@hks.harvard.edu.

[‡]Harvard Kennedy School, amoorthy@g.harvard.edu.

[§]University of Pennsylvania, pettigr@sas.upenn.edu.

[¶]Harvard Kennedy School, benjamin_schneer@hks.harvard.edu.

1 Introduction

Local newspapers provide valuable information to communities, shaping public discourse and citizen engagement (Mondak 1995; Snyder and Strömberg 2010; Gentzkow et al. 2011; King et al. 2017). Yet in recent decades, the industry’s traditional economic model has been upended. As revenues from print circulation declined, newspapers around the country began charging for access to online content through subscription-based digital paywalls. Constructing a novel database tracking paywall implementation for the universe of U.S. daily newspapers, we find that the share of Americans who faced a paywall at their local newspaper rose from less than 2% in 2006 to over 90% in 2020 (Panel A of Figure 1). The shift from free to paid news content – during an era when the internet provides instant access to a near-infinite range of news and non-news substitutes – has fundamentally altered how Americans access information about their local communities.

[Figure 1 about here.]

Despite this sweeping transformation, little empirical research has examined the consequences of paywalls. Existing work focuses primarily on their direct effects on a newspaper’s own readership, typically using aggregate web traffic data from a limited set of outlets (Chiou and Tucker 2013; Pattabhiramaiah et al. 2019; Kim et al. 2020; Olsen et al. 2020; Streyczek 2025). While informative, these studies provide limited insight into broader civic and informational implications. When local news becomes costly, where do readers turn? And how do the resulting shifts affect public knowledge, perceptions, and engagement?

We address these questions through the first comprehensive examination of newspaper paywalls. Our analysis combines household-level web-browsing records for over a million panelists with the hand-collected information on paywall adoption and website-level measures of partisan slant from Robertson et al. (2018, 2023). These data reveal two important features of the local news landscape. First, local papers skew sharply right of other news sources: the median paper is roughly as conservative as the Wall Street Journal. Second, local newspaper

readership is disproportionately concentrated among higher-income households – a pattern that sharpened over the past two decades.

Both features are central to understanding what happens when local newspapers become gated. Because local papers provide unique coverage, any substitution is likely to expose displaced readers to meaningfully different online content. As displacement itself depends on households’ willingness and ability to pay for continued access to local news, the downstream consequences of paywalls may also diverge across the income distribution. Together, these dynamics recall longstanding theories of the costly acquisition of information (Downs 1957) and its unequal diffusion (Tichenor et al. 1970), suggesting that even modest price barriers can reshape not just who consumes the news, but also the knowledge and behavior that depend on it.

To trace the full arc of these effects, we exploit variation in the timing of paywall adoption across newspapers in a staggered difference-in-differences design. At the newspaper level, we compare readership before and after paywall implementation relative to newspapers without paywalls. At the household level, complete browsing records allow us to further identify where affected readers go and how their information diets change. Using local voting and survey data, we then examine how these consumption shifts spill over into political behavior and knowledge and perceptions of area conditions. Throughout, we explore differential effects by income.

We find that paywalls sharply reduce online readership. Total web traffic at paywalled newspaper websites declines, on average, by nearly 30% a year after adoption. Readership losses appear across multiple margins: browsing time, page views, and unique visitors all drop significantly in the wake of paywalls. These effects are robust across a range of specifications – including alternative sample definitions, controls for newspaper-specific time trends, and Callaway and Sant’Anna (2021) corrections for staggered adoption – and are validated by parallel pre-trends in the months leading up to adoption.

For displaced readers, the shift away from local newspapers represents a net decline in

local news consumption. We find no evidence that households substitute to alternative sources of area information – such as local television, municipal, or community websites – nor do we find evidence of a broader shift toward non-news content like streaming or social media. Instead, readers turn to national newspapers and other outlets that are, on average, less conservative than the local papers they left behind. The result is a significantly more liberal news diet.

This reallocation generates measurable spillovers on local communities. We find that paywalls significantly increase Democratic support in affected counties, mirroring the leftward shift in households’ news consumption. Turnout in federal elections rises as well, but only in presidential years, when top-of-ticket races attract outsized national attention. Finally, consistent with local papers’ distinct coverage of crime and civic affairs, paywalls bolster perceptions of community safety and reduce knowledge of local – but not national – economic conditions.

From the direct readership effects to the downstream civic consequences, we find consistent evidence of income-based stratification. Paywalls skew local reader bases away from the bottom of the income distribution toward the top, reducing the share of readership from low-income households by nearly 20%. Disproportionately filtered out of local news access, poorer households drive the leftward shift in news consumption as they seek out alternative outlets. These same households then experience the largest changes in political behavior and the sharpest declines in local knowledge.

These findings make two main contributions. First, we build upon a nascent literature examining newspaper paywalls, one of the key media innovations of the early 21st century. While prior work has documented sharp reductions in readership among limited samples of Scandinavian and American papers, the breadth and granularity of our data allow us to move beyond direct effects to document the substitution patterns and downstream spillovers that follow. In doing so, we complement a large body of research exploring the influence of media on voting (DellaVigna and Kaplan 2007; Gerber et al. 2009; Martin and Yurukoglu 2017),

social capital (Olken 2009; Bauernschuster et al. 2014; Campante et al. 2022; Ang 2023), and public discourse (King et al. 2017). Rather than the civic withdrawal typically associated with reduced access to news, we find reorientation, with increasing national engagement even as local knowledge erodes. During an era of online abundance, these results demonstrate how the consequences of media transformations depend not just on what people lose access to, but on what they consume instead (Gentzkow 2006).

Second, we contribute to research on the decline of local news (Hayes and Lawless 2015; Djourelova et al. 2025; Peterson 2021; Ewens et al. 2022). Unlike the newspaper closures and revenue shocks examined elsewhere, paywalls do not eliminate local coverage but instead impose a direct price on it. As such, they offer insight into fundamental but understudied questions: who actually demands local news, and how substitutable are those preferences? While information about area politics, business, and culture may be broadly relevant, we find that demand for local news is surprisingly elastic, particularly among low-income households. When access becomes costly, many readers abandon local news in favor of categorically different content, producing distinct information environments across the income distribution. Knowledge gaps thus arise not from differential absorption mediated by income (Tichenor et al. 1970), but from differential access purchased with it. In this light, paywalls may present a classic equity-efficiency tradeoff: a market-based solution to the local news crisis that nonetheless compounds existing information inequities.

The paper proceeds as follows. Section 2 describes the data and documents key patterns in the local news landscape. Section 3 establishes direct effects of paywalls on newspaper readership. Section 4 examines substitution patterns. Section 5 explores downstream civic consequences. Section 6 concludes.

2 Data and Descriptives

2.1 Data

Paywalls

We construct a comprehensive database tracking paywall implementation for the universe of U.S. daily newspapers. For newspapers with active websites, we identify whether and when they first implemented a paywall by hand-collecting information from newspaper announcements, media research organizations, the Internet Archive’s Wayback Machine, direct correspondence with editors, and prior work (Kim et al. 2020). The control group consists of 146 newspapers that had not implemented a paywall by the end of our Comscore sample period in 2020. The treatment group consists of 869 newspapers for which we were able to identify the month or quarter of paywall implementation.¹

Panel B of Figure 1 shows the timing of paywall adoption. The first major wave occurred from 2010 to 2014, followed by a smaller wave beginning in 2018. Around 80% of paywalls were “soft” or metered, allowing users access to a limited number of free articles each month (typically 5-15) before requiring payment (Appendix Figure A.1). The remainder were “hard” paywalls requiring subscription to access any article, or freemium models offering free headlines but restricted premium content. Monthly subscriptions for local paywalls averaged around \$10 at implementation, similar to the fees initially charged by national newspapers.²

Internet Browsing

Internet browsing data from 2006 to 2020 come from Comscore’s web behavior database. Each calendar year, Comscore enlists a panel of internet users to install a computer plug-in

¹Given the temporal granularity of our analysis, papers with indeterminate paywall timing are excluded from the sample. These outlets account for a small share of local news readership, averaging 80% less web traffic than newspapers in the analysis sample.

²Digital access to *New York Times* and *Washington Post* cost \$14.99/month and \$9.99/month when their paywalls were first implemented in 2011 and 2013, respectively.

tracking web browsing during that year.³ While mobile browsing is not captured, computers remained the dominant mode of online news consumption through 2017, after the bulk of paywalls had been adopted (Nielsen et al. 2023). Each observation records the date, time, number of pages viewed, and number of minutes spent on a website by a computer during a browsing session.⁴ Comscore also includes self-reported information on household zip code, income, and other demographics. In total, the data contain billions of site visits across over a million households.⁵

To identify each household’s relevant local newspaper, we calculate the non-national newspaper with the most online readership in a county from 2006 to 2009, prior to the first wave of paywalls.⁶ To characterize other online news consumption, we use the list of “hard news” websites compiled by Allen et al. (2020). Finally, to measure the political slant of news outlets, we draw upon scores generated by Robertson et al. (2018) from the differential social media sharing patterns of Democrats and Republicans. While highly correlated with Gentzkow and Shapiro’s (2010) text-based measures of newspaper slant (Appendix Figure A.2), the Robertson scores cover a wide range of newspaper and non-newspaper domains, including popular sites like CNN and Fox News.⁷

2.2 Local News Landscape

The web browsing data reveal several important facts about local newspapers and their online readership. First, local newspapers occupy a distinct position in the news landscape.

³The vast majority of computers are only observed in the panel for a single calendar year.

⁴Comscore only reports root domains and not the full URL, so we are unable to disentangle, for example, visits to google.com from news.google.com or google.com/travel/flights, or to identify the exact news article a user read on nytimes.com.

⁵In practice, the Comscore data identify observations by computer, with no way to determine if multiple computers belong to a single household or if multiple users use the same computer. We use the terms computers, users, and households interchangeably throughout.

⁶We exclude the four national newspapers (i.e., New York Times, Washington Post, Wall Street Journal and USA Today) except in their county headquarters. For instance, the Washington Post is considered “local” to the District of Columbia. Results are similar under alternative definitions, including requiring the local newspaper to be headquartered in the household’s state or to account for over half of county newspaper readership.

⁷In aggregate, Robertson-scored domains capture nearly 20 times the web browsing traffic of outlets covered by Gentzkow and Shapiro (2010).

As shown in Panel C of Figure 1, local newspapers skew more conservative than other sources of news. Local newspapers are, on average, 0.57 standard deviations more conservative than national newspapers and 0.21 standard deviations more conservative than all other online news websites ($p < 0.001$). Put differently, the median local newspaper is roughly as conservative as the nation’s most right-leaning “newspaper of record,” the Wall Street Journal.

Second, local newspaper readers comprise a small and demographically distinct segment of households. Fewer than 15% of Comscore users spent at least 10 minutes on their local newspaper’s website over the course of a year. Those who did were whiter (82% versus 63%), older (60% versus 56% over age 44), and less likely to be low-income (20% versus 23% below \$25,000 a year) than typical Comscore users (Appendix Table A.1).

Third, local newspaper readership evolved markedly over the sample period. In 2006, households spent nearly 70% more time at their local newspaper’s website than at the websites of all four national newspapers combined (76 minutes versus 45 minutes). By 2020, the pattern had reversed, with households spending roughly 25% less time at local sites (15 minutes versus 19 minutes). Local news consumption also became increasingly selected on income over this interval. As shown in Panel D of Figure 1, the income distribution of local news readers mirrored that of the Comscore user base early in the sample period. However, beginning around 2010, the share of local readers from the upper-end of the income distribution steadily rose, while the share earning less than \$25,000 a year fell. By 2020, local news readers were 12 percentage points – or roughly 50% – less likely than the full sample to come from the bottom of the income distribution.

These descriptive patterns highlight the unique content and evolving readership composition of local newspapers, and motivate the causal analysis in the following sections.

3 Direct Readership Effects

We begin by examining the direct effect of paywalls on newspaper readership. To do so, we construct a newspaper-month panel aggregating web traffic across all Comscore households, and estimate the following event-study model:

$$y_{pt} = \sum_{k \neq -1} \beta_k \times D_{pt}^k + \gamma_p + \delta_t + \epsilon_{pt}. \quad (1)$$

Outcome y_{pt} is the log of total time spent at newspaper p 's website during month t . D_{pt}^k denote month-to-paywall indicators for treatment papers, omitting the month before a paywall is introduced.⁸ Periods more than 12 months from paywall adoption are top- and bottom-coded following Schmidheiny and Siegloch (2023). γ_p and δ_t are newspaper and month fixed effects, respectively, and standard errors are clustered by newspaper, the level of treatment.

[Figure 2 about here.]

Results are shown in Panel A of Figure 2. While local news consumption was falling nationwide over the sample period, we find little evidence of differential pre-trends between treatment and control newspapers in the months leading up to paywall implementation. However, after implementation, online readership of paywalled papers drops significantly. Effects emerge immediately and persist over time, with average web traffic falling by roughly 30% one year after adoption.

Table 1 demonstrates robustness to a host of other specifications. Column 1 summarizes our baseline effects on log duration using a single pooled post-treatment indicator. Columns 2 and 3 instead examine log page views and log unique visitors, respectively, providing evidence of both intensive- and extensive-margin reductions. Column 4 weights observations by pre-treatment web traffic. Column 5 allows for newspaper-specific linear time trends. Column 6 drops papers owned by media conglomerates to ensure results are not driven by a

⁸For newspapers where only the quarter of implementation is known, we impute the treatment month as the middle month of the quarter.

small number of large companies. Finally, Column 7 applies Callaway and Sant’Anna (2021) adjustments to account for potential treatment heterogeneity.⁹ In all cases, we find that paywalls significantly reduce newspaper readership by 15 to 40%, broadly consistent with estimates from Streyczek (2025) for a sample of major U.S. newspapers.

Beyond aggregate declines in readership, paywalls also alter who consumes local news. Panel B of Figure 2 presents difference-in-differences estimates of paywall effects on the income composition of readership. After implementation, the share of local web traffic coming from households earning less than \$25,000 a year falls, on average, by 3 percentage points. The effect represents a nearly 20% decline from baseline and is offset by a similarly sized, though noisy, increase among households earning over \$100,000 a year. These compositional shifts demonstrate how paywalls disproportionately filter out those least able to pay, exacerbating the widening income divides in local news access documented in Section 2.

4 Household Substitution Patterns

4.1 Empirical Strategy

While the paper-level analysis shows that paywalls alter both the aggregate level and composition of newspaper readership, understanding the broader consequences requires tracing where displaced readers redirect their attention. The granular browsing data allow us to zoom in on affected households and identify the types of content they choose to consume instead. To do so, we employ the following event study design comparing trends in browsing behavior among households in paywalled counties with those elsewhere:

$$y_{it} = \sum_{k \neq -1} \beta_k \times D_{it}^k + \gamma_i + \delta_t + \epsilon_{it}. \quad (2)$$

The model is analogous to the newspaper-level approach described in Equation 1 with

⁹Full event study estimates for the page views and Callaway-Sant’Anna specifications are shown in Appendix Figure A.3.

a few distinctions. Here, y_{it} may represent the log of time spent by a household i during month t across a range of website categories, not just at the paywalled newspaper. D_{it}^k denote month-to-paywall indicators for household i 's local paper. γ_i and δ_t are household and month fixed effects, respectively, and standard errors are clustered by county.

Because Comscore panels are constructed annually, we restrict the treatment group to households observed in the same calendar year that their local newspaper adopted a paywall and pool time-to-treatment indicators beyond ± 6 months. The control group includes all households assigned to local newspapers that had not implemented a paywall by the end of the sample period. Finally, to ensure a balanced and consistent panel, we follow Athey et al. (2021) and restrict the sample to active internet users, defined as households that visited at least one newspaper website and spent at least 50 hours online over the panel year.

4.2 Destination Sites

Panels C and D of Appendix Figure A.3 replicate the direct effects of paywalls on a paper's own readership using our household-level approach, finding consistent support for parallel pre-trends followed by significant post-implementation declines.¹⁰

To explore whether households mitigate potential information losses by substituting to other sources of local news, we identify the five closest alternatives to each newspaper using data on organic search competitors from Adobe Semrush, a digital marketing and search engine optimization platform. Panel A of Figure 3 shows little evidence that households substitute toward these outlets. While many are local television, radio, or municipal websites providing overlapping content with local newspapers, the null effects suggest either that readers do not view them as adequate replacements or that the friction costs of switching exceed the perceived value of local information. Both interpretations imply elastic demand for local news.

¹⁰The household-level approach estimates effects only for households in a newspaper's county, weighting households equally, whereas the paper-level approach assesses readership regardless of household location and implicitly weights toward households that consume more news.

[Figure 3 about here.]

Instead, as shown in Panel B of Figure 3, we find that barriers to local news access lead displaced readers to turn to national newspapers.¹¹ While all national papers except USA Today were paywalled for much of the sample period, our results suggest that many displaced households were nonetheless willing to pay for access to those outlets – consistent with findings by Gentzkow (2006) showing that locally-oriented news content was readily abandoned at the advent of television. A reallocation from local to national news is corroborated by Panel B of Table 1, which presents difference-in-differences estimates across a range of website categories. Null effects for local community platforms and total hard news consumption confirm that households do not seek out other local sources of information and instead switch to national outlets. Meanwhile, positive but statistically insignificant estimates for search, streaming, social media, and total online consumption indicate limited substitution to non-news websites.¹²

4.3 News Slant

The decline of local newspapers is typically understood as a geographic phenomenon – communities losing coverage of local institutions, elections, and events. Yet given their conservative lean, paywalls may have ideological consequences as well, shifting the partisan composition of households’ news diets as readers migrate to other online sources. To test this, we categorize all news domains by quantiles of partisan bias scores (Robertson et al. 2018), allowing us to examine how paywalls shift the distribution of partisan consumption across a household’s entire news portfolio.

Difference-in-differences estimates for time spent in each slant category are displayed in Panel C of Figure 3. We find that paywalls significantly increase the amount of time local households spend at center-left news websites, with suggestive evidence that this increase

¹¹Appendix Figure A.4 demonstrates consistent effects across each of the four national newspapers.

¹²Search engines, streaming and video platforms, and social media sites are identified using Shallalist, a blacklist service frequently used to categorize URLs (Paredes et al. 2018; Laohaprapanon and Sood 2021).

comes at the expense of more right-leaning outlets. The pattern is consistent with the null effect on total news consumption and indicates that paywalls induce a net leftward shift in local households’ news diets.

Given that paywalls disproportionately reduce readership among low-income households, the effects on partisan consumption may also differ across the income distribution. We thus allow treatment estimates to vary by income group and examine changes in the average partisan slant of a household’s news diet.¹³ As shown in Panel D of Figure 3, the leftward shift in news consumption is concentrated among the lowest-income households – precisely those whose access to local news was most disrupted by paywalls. In contrast, we find near-zero estimates for higher-income households.

Together, the substitution patterns highlight the distinct informational and distributional consequences of paywalls. Poor households, in particular, experience net reductions in local news consumption and a meaningful shift toward ideologically and topically distinct content. Given the well-documented role of news media in shaping community beliefs and behavior, this reshuffling may cascade into a range of downstream consequences, which we explore next.

5 Downstream Civic Consequences

5.1 Voter Behavior

We first examine the political consequences of paywalls using county-level voting data. To do so, we estimate an event-study model analogous to Equation 2 and compare trends in election outcomes before and after paywall implementation. Treatment is again determined by a county’s most read local paper from 2006 to 2009.

Panel A of Figure 4 displays effects on Democratic vote share in congressional elections.

¹³To measure average news partisanship, websites are scored by slant category from -3 (“extreme left”) to +3 (“extreme right”). We then calculate the average score across all websites visited by a household in a month, weighting by the time spent at each site.

Consistent with parallel trends, we find little evidence of differential voting patterns across treatment and control counties before paywall adoption. Democratic vote share then increases steadily after implementation. These results align with household-level substitution away from local papers toward more liberal national news and are the mirror image of the rightward shifts caused by exposure to conservative cable outlets documented in prior work (DellaVigna and Kaplan 2007; Martin and Yurukoglu 2017).

[Figure 4 about here.]

In contrast to vote choice, the effects of paywalls on voter turnout depend on election timing. As shown in Panel B of Figure 4, we find significant increases in turnout in presidential elections, but near-zero estimates in midterm elections. This asymmetry is consistent with greater national news consumption raising the salience of presidential races more than that of midterm contests, where participation often depends on localized coverage of congressional candidates (Gentzkow et al. 2011; Hayes and Lawless 2015).

Panel C of Figure 4 examines how the political consequences of paywalls vary with county income. We find that the leftward shift in Democratic vote share and the increase in presidential turnout are concentrated in lower-income communities. The heterogeneity analysis further reveals that the aggregate null effects on midterm turnout mask significant reductions in participation in these areas, consistent with declining access to local news.

5.2 Local Knowledge and Perceptions

As local papers provide unique coverage of economic and social conditions that many residents cannot observe firsthand, paywalls may alter a range of beliefs about community life (Lippmann 1922). To explore these potential effects, we turn to data from Gallup Daily, a repeated cross-section survey of roughly 1,000 Americans each day. The survey has two unique features critical for our setting. First, the data are identified by date and county, allowing us to map responses to the timing and location of paywalls. Second, from 2008

to 2017, the survey included consistently asked questions regarding beliefs about both local and national conditions.¹⁴

Given the extensive crime coverage found in local news (Weitzer and Kubrin 2004), we first explore effects on perceptions of safety. Panel D of Figure 4 presents event study estimates using the survey microdata, including county and date fixed effects and controlling for respondent age, gender, race, marital status, education, and income.¹⁵ Validating a parallel trends assumption, perceived safety evolved similarly across groups in the months leading up to paywalls, before rising steadily among respondents in paywalled counties. This pattern is consistent with existing work documenting how crime reporting shapes fear (Mastrorocco and Minale 2018; Mastrorocco and Ornaghi 2025), and suggests that paywalls increased citizens’ sense of safety through reduced exposure to local crime coverage.

To examine effects on knowledge, we construct measures comparing respondents’ assessments of local and national economic conditions with objective benchmarks. For national conditions, responses to a question asking whether economic conditions in “this country, as a whole, are getting better or worse” are benchmarked against national GDP growth. For local conditions, responses to a question asking whether local businesses are “expanding” or “reducing” their workforces are benchmarked against analogous measures of county economic growth.¹⁶ Given that the exact reference frame any given respondent was assessing is unknown, neither measure is perfect. However, the Appendix provides validating evidence that both measures capture meaningful variation in economic knowledge.¹⁷

Event study estimates are presented in Panel E of Figure 4. Local economic knowledge

¹⁴Gallup Daily was restructured in 2018, sharply reducing the number of responses to these questions from 150,000 per year to as few as 3,000. We therefore restrict our sample accordingly.

¹⁵The outcome is an indicator set to one if a respondent answered “Yes” to “*Do you feel safe walking alone at night in the city or area where you live*” (polled from 2008 to 2013) or “Agree” or “Strongly Agree” to “*You always feel safe and secure*” (polled from 2014 to 2017).

¹⁶For each geographic level, we define a knowledge indicator equal to one if a respondent reported that conditions were improving and actual GDP growth was in the top two quartiles of the sample distribution, reported that conditions were staying the same and actual growth was in the middle two quartiles, or reported that conditions were worsening and actual growth was in the bottom two quartiles.

¹⁷Panel A of Figure A.5 shows that raw assessments of economic conditions are positively and significantly correlated with actual GDP growth at each geographic level, while Panel B demonstrates that more educated respondents hold more accurate economic beliefs according to our knowledge measures.

falls significantly in the months after local newspaper paywalls are implemented. In contrast, we find little change in national economic knowledge. These patterns highlight the asymmetric availability of different news content: a wide variety of sources cover national news, but only local newspapers report on local conditions. While prior work has established the importance of local newspapers for political knowledge (Mondak 1995; Hayes and Lawless 2015; Gentzkow et al. 2011; Streycek 2025), our results further demonstrate how these institutions shape citizens’ fundamental understanding of the communities they live in.

Panel F of Figure 4 examines how effects vary with household income. After paywalls, self-reported feelings of safety increase significantly among low-income households, while accurate knowledge of local economic conditions falls. In contrast, higher-income households exhibit little change on either margin, echoing the sharp divergence observed for readership, substitution, and voting.

6 Conclusion

This paper provides the first comprehensive account of how newspaper paywalls have reshaped access to local news and the communities it serves. We show that paywalls do not simply reduce a newspaper’s own readership; they redirect attention away from local news entirely. Instead of seeking out alternative sources of similar information, households substitute toward more nationally oriented and liberal-leaning content. The shift from relatively conservative local news to relatively liberal national news generates corresponding spillovers across a range of outcomes, from political preferences and participation to civic knowledge and beliefs.

Given the unique role of local newspapers in the wider information environment, these effects likely understate the full impact of paywalls. Local papers are the primary vehicles through which many residents learn about community gatherings, business openings, and municipal affairs. To the extent that paywalls distort “the pictures in our heads” (Lippmann

1922) and erode a shared portrait of civic life, their ramifications may extend well beyond the ballot box to the very fabric of American society (Putnam 2000; Bauernschuster et al. 2014; Campante et al. 2022).

Compounding these potential consequences is the fact that newspaper paywalls stratify information access within communities, with the largest shifts in news consumption, political behavior, and local knowledge all accruing to the lowest-income households. This pattern takes on added significance against the backdrop of a broader digital transformation. As social media platforms, streaming providers, and AI services increasingly adopt gated models, the deeper consequence may not be reduced access in the aggregate, but rather an increasingly fractured society in which what people know, believe, and engage with is determined by their ability to pay.

References

- Allen, Jennifer, Baird Howland, Markus Mobius, David Rothschild, and Duncan J. Watts, “Evaluating the Fake News Problem at the Scale of the Information Ecosystem,” *Science Advances*, 2020, 6 (14), eaay3539.
- Ang, Desmond, “The birth of a nation: Media and racial hate,” *American Economic Review*, 2023, 113 (6), 1424–1460.
- Athey, Susan, Markus Mobius, and Jeno Pal, “The Impact of Aggregators on Internet News Consumption,” *NBER Working Paper 28746*, 2021.
- Bauernschuster, Stefan, Oliver Falck, and Ludger Woessmann, “Surfing Alone? The Internet and Social Capital: Evidence from an Unforeseeable Technological Mistake,” *Journal of Public Economics*, 2014, 117, 73–89.
- Callaway, Brantly and Pedro H.C. Sant’Anna, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Campante, Filipe, Ruben Durante, and Andrea Tesei, “Media and Social Capital,” *Annual Review of Economics*, 2022, 14, 69–91.
- Chiou, Lesley and Catherine Tucker, “Paywalls and the demand for news,” *Information economics and policy*, 2013, 25 (2), 61–69.
- DellaVigna, Stefano and Ethan Kaplan, “The Fox News Effect: Media Bias and Voting,” *The Quarterly Journal of Economics*, 2007, 122 (3), 1187–1234.
- Djourelouva, 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*, 2025, 92 (3), 1738–1772.
- Downs, Anthony, *An Economic Theory of Democracy*, New York: Harper and Row, 1957.
- Ewens, Michael, Arpit Gupta, and Sabrina T. Howell, “Local Journalism Under Private Equity Ownership,” *NBER Working Paper 29743*, 2022.
- Gentzkow, Matthew, “Television and voter turnout,” *The Quarterly Journal of Economics*, 2006, 121 (3), 931–972.
- and Jesse M. Shapiro, “What Drives Media Slant? Evidence from U.S. Daily Newspapers,” *Econometrica*, 2010, 78 (1), 35–71.
- , —, and Michael Sinkinson, “The Effect of Newspaper Entry and Exit on Electoral Politics,” *American Economic Review*, 2011, 101 (7), 2980–3018.
- 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.


- Hayes, Danny and Jennifer L. Lawless**, “As Local News Goes, So Goes Citizen Engagement: Media, Knowledge, and Participation in U.S. House Elections,” *The Journal of Politics*, 2015, 77 (2), 447–462.
- Kim, Ho, Reo Song, and Youngsoo Kim**, “Newspapers’ Content Policy and the Effect of Paywalls on Pageviews,” *Journal of Interactive Marketing*, 2020, 49 (1), 54–69.
- King, Gary, Benjamin Schneer, and Ariel White**, “How the News Media Activate Public Expression and Influence National Agendas,” *Science*, 2017, 358 (6364), 776–780.
- Laohaprapanon, Suriyan and Gaurav Sood**, “Domain Knowledge: Predicting the Kind of Content Hosted by a Domain,” in “13th International Conference on Computational Intelligence in Security for Information Systems (CISIS 2020) 12” Springer 2021, pp. 152–162.
- Lippmann, Walter**, *Public Opinion*, New York: Harcourt, Brace and Company, 1922.
- Martin, Gregory J. and Ali Yurukoglu**, “Bias in Cable News: Persuasion and Polarization,” *American Economic Review*, 2017, 107 (9), 2565–2599.
- Mastorocco, Nicola and Arianna Ornaghi**, “Who watches the watchmen? Local news and police behavior in the United States,” *American Economic Journal: Economic Policy*, 2025, 17 (2), 285–318.
- **and Luigi Minale**, “News media and crime perceptions: Evidence from a natural experiment,” *Journal of Public Economics*, 2018, 165, 230–255.
- Mondak, Jeffery J.**, “Newspapers and Political Awareness,” *American Journal of Political Science*, 1995, pp. 513–527.
- Nielsen, Rasmus Kleis, Nic Newman, Rasmus Fletcher, and Alexios Kalogeropoulos**, “Reuters institute digital news report 2020,” *Reuters Institute*, 2023.
- Olken, Benjamin A.**, “Do Television and Radio Destroy Social Capital? Evidence from Indonesian Villages,” *American Economic Journal: Applied Economics*, 2009, 1 (4), 1–33.
- Olsen, Ragnhild Kristine, Aske Kammer, and Mona Kristin Solvoll**, “Paywalls’ Impact on Local News Websites’ Traffic and their Civic and Business Implications,” *Journalism Studies*, 2020, 21 (2), 197–216.
- Paredes, Rolysent K., Ariel M. Sison, and Ruji P. Medina**, “Bandwidth Allocation Mechanism Based on Users’ Web Usage Patterns for Campus Networks,” *International Journal of Communication Networks and Information Security*, 2018, 10 (2), 380–388.
- Pattabhiramaiah, Adithya, Shreya Sriram, and Puneet Manchanda**, “Paywalls: Monetizing online content,” *Journal of Marketing*, 2019, 83 (2), 19–36.
- Peterson, Erik**, “Paper Cuts: How Reporting Resources Affect Political News Coverage,” *American Journal of Political Science*, 2021, 65 (2), 443–459.

- Putnam, Robert D**, *Bowling alone: The collapse and revival of American community*, Simon and Schuster, 2000.
- Robertson, Ronald E, Jon Green, Damian J Ruck, Katherine Ognyanova, Christo Wilson, and David Lazer**, “Users choose to engage with more partisan news than they are exposed to on Google Search,” *Nature*, 2023, *618* (7964), 342–348.
- Robertson, Ronald E., Shan Jiang, Kenneth Joseph, Lisa Friedland, David Lazer, and Christo Wilson**, “Auditing Partisan Audience Bias Within Google Search,” *Proceedings of the ACM on Human-Computer Interaction*, 2018, *2* (CSCW), 148.
- Schmidheiny, Kurt and Sebastian Siegloch**, “On event studies and distributed-lags in two-way fixed effects models: Identification, equivalence, and generalization,” *Journal of Applied Econometrics*, 2023, *38* (5), 695–713.
- Snyder, James M. and David Strömberg**, “Press Coverage and Political Accountability,” *Journal of Political Economy*, 2010, *118* (2), 355–408.
- Streyczek, Julian**, “Political Effects of Newspaper Paywalls,” *Available at SSRN 5124655*, 2025.
- Tichenor, Phillip J., George A. Donohue, and Clarice N. Olien**, “Mass Media Flow and Differential Growth in Knowledge,” *Public Opinion Quarterly*, 1970, *34* (2), 159–170.
- Weitzer, Ronald and Charis E. Kubrin**, “Breaking News: How Local TV News and Real-World Conditions Affect Fear of Crime,” *Justice Quarterly*, 2004, *21* (3), 497–520.


Figure 1: Local Newspaper Paywalls

Paywalls

(a) Paywall Coverage over Time




(b) Paywall Adoption Dates




Local Newspapers

(c) Distribution of Partisan Slant




(d) Income Composition of Readers



Note: Panel A plots the share of Comscore households with a paywall at their local newspaper over time. We define a household’s local newspaper as the non-national newspaper with the most online readership from 2006 to 2009 in the household’s county. Panel B plots the distribution of when daily newspapers introduced paywalls. Panel C plots the distribution of partisan slant among local newspapers, national newspapers (i.e., New York Times, Washington Post, USA Today, and Wall Street Journal), and other “Hard News” sites from Allen et al. (2020). Partisan slant scores are from Robertson et al. (2018), with lower values reflecting more left-leaning content and higher values reflecting more right-leaning content. Panel D plots the difference between the share of local newspaper readers in each income bracket and the share of all Comscore households in each income bracket, where local readers are defined as households who spent at least 10 minutes at their local newspaper website over the course of the year. The income brackets shown (<\$25k, \$25k-\$75k, \$75k-\$100k, \$100k+) are the only four categories that can be consistently observed throughout the Comscore sample.


Figure 2: Direct Effects of Paywalls




Note: Panel A plots event-study coefficients for the effect of paywalls on newspaper readership. Outcome is the log of total time (in minutes) spent by all Comscore households – regardless of location – at a newspaper in a given month. Months beyond +/- 12 are top- and bottom-coded following Schmidheiny and Siegloch (2023). Panel B plots difference-in-differences estimates (i.e. for a single post-treatment indicator) for the effect of paywalls on the income composition of local readers. Outcomes are the share of all local browsing time at a paper-month that came from households in each of the four consistently-observed Comscore income categories (<\$25k, \$25k-\$75k, \$75k-\$100k, \$100k+), weighting observations by the total browsing time at each newspaper (averaged from 2006 to 2009). We include newspaper and year-month fixed effects, and cluster standard errors at the newspaper-level.

Figure 3: Substitution Effects of Paywalls

News Consumption




News Partisanship




Note: Panels A and B display event study estimates for the effect of local newspaper paywalls on the log of total time spent by a household at local news competitors and national newspapers, respectively. Competitors include each paper’s five closest organic search competitors identified by Adobe Semrush, a digital marketing platform. National newspapers include the New York Times, Washington Post, USA Today, and Wall Street Journal. Event study results for each national paper are shown in Appendix Figure A.4. Months beyond ± 6 are top- and bottom-coded following Schmidheiny and Siegloch (2023). Panel C plots difference-in-differences estimates (i.e. for a single post-treatment indicator) from regressions of consumption in each category of partisan slant. Outcomes are the log of total time spent by a household at all web domains in a given slant quantile. Panel D presents difference-in-differences estimates of the effect of paywalls on average partisan slant by household income category. Average partisan slant for a household-month is calculated by averaging across time spent at websites in each partisan slant category – ranging from -3 (“extreme left”) to +3 (“extreme right”). Because each household is observed for only a single panel year, income categories are defined by quartiles of the full set of Comscore income categories available in that year (seven categories prior to 2014, eight thereafter). We include household and year-month fixed effects, and cluster standard errors at the county-level.

Figure 4: Downstream Effects of Paywalls

Voter Behavior



Knowledge and Perceptions



Note: Panels A through C examine downstream effects on elections. Panel A plots event study estimates for county Democratic vote share in congressional elections from 2006 to 2020, indexed in two-year election cycles. Panel B plots estimates for county presidential and midterm turnout, defined as votes cast divided by citizen voting age population, indexed in four-year election cycles. Citizen voting age population in a county-year comes from 5-year ACS estimates with the corresponding terminal year, available beginning in 2010. Panel C presents difference-in-differences estimates of paywall effects on Democratic vote share, presidential turnout, and midterm turnout for counties in the bottom quartile of median household income in 2010 and all other counties. Panels D through F examine effects on knowledge and perceptions using Gallup Daily survey data from 2008 to 2017 and include controls for respondent race, age, gender, marital status, education and income. Panel D plots estimates for perceived safety, defined as an indicator equal to one if a respondent reports feeling safe in their community. Panel E presents estimates for local and national economic knowledge, defined as an indicator set equal to one if a respondent correctly assesses county and national economic conditions. Panel F presents difference-in-differences estimates of paywall effects on safety and knowledge for households with annual income below and above \$24,000, the Gallup income category nearest to Comscore's <\$25,000 bracket. Panels A through C include county and year fixed effects. Panels D through F include county and date fixed effects. All specifications cluster standard errors at the county level. For event studies, endpoints are top- and bottom-coded following Schmidheiny and Siegloch (2023).

Table 1: Direct and Substitution Effects of Paywalls

<i>Panel A: Direct Effects on Newspaper Traffic (Paper-Level)</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Paywall x Post	-0.288*** (0.051)	-0.210*** (0.050)	-0.149*** (0.025)	-0.367* (0.188)	-0.201*** (0.045)	-0.183** (0.073)	-0.483*** (0.076)
Specification	Base	log(Pages)	log(Users)	Weighted	Paper trends	Drop conglom.	Callaway- Sant’Anna
Mean (level)	959	496	40.2	959	959	959	959
Observations	168,888	168,888	168,888	162,696	168,888	94,476	168,888

<i>Panel B: Substitution Effects by Website Category (Household-Level)</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Paywall x Post	-0.005 (0.013)	0.008 (0.008)	0.001 (0.001)	0.010 (0.013)	0.020 (0.013)	0.019 (0.015)	0.022 (0.015)
Category	Hard news	Soft news	Community	Search engines	Streaming	Social media	All
Mean (level)	464	34.6	.24	637	427	536	3,923
Observations	687,552	687,552	687,552	687,552	687,552	687,552	687,552


Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: Panel A presents difference-in-differences estimates (i.e. for a single post-treatment indicator) for the direct effects of paywalls on a newspaper’s own website traffic. Column 1 presents our baseline specification (Equation 1), which includes newspaper and year-month fixed effects and examines log of total time (in minutes) as the outcome. Columns 2 and 3 replace log duration with log page views and log unique visitors, respectively. Column 4 weights observations by the total browsing time at each newspaper (averaged from 2006 to 2009). Column 5 adds newspaper-specific linear time trends. Column 6 excludes newspapers owned by large conglomerates (i.e., Gannett/GateHouse, Tribune/Digital First, Lee Enterprises, CNHI, and Adams Publishing). Column 7 uses Callaway-Sant’Anna 2021 corrections to account for heterogeneous treatment effects under staggered adoption. Panel B presents difference-in-differences estimates for the effects of paywalls on household consumption across website categories, using our baseline household-level specification with household and year-month fixed effects (Equation 2). The outcome is the log of total time spent in each category per household-month. Columns 1 and 2 examine hard and soft news websites as classified by Allen et al. (2020). Column 3 examines community platforms *Patch.com* and *Nextdoor.com*. Columns 4 through 6 examine sites classified as search engines, streaming and video, and social media by Shallalist. Column 7 examines total web consumption.


A Additional Figures

Figure A.1: Paywall Subscription Models and Pricing




Note: Panel A plots the share of daily newspapers adopting freemium, hard, and soft (metered) paywalls. Panel B plots the distribution of free articles allowed under soft paywalls. Panels C and D plot the distribution of monthly subscription prices at time of paywall introduction and as of 2023, respectively.

Figure A.2: Correlation of Partisan Slant Measures




Note: Binned scatter plot of the correlation between Robertson et al. (2018) partisan bias scores and slant scores from Gentzkow and Shapiro (2010).

Figure A.3: Robustness of Direct Effects on Newspaper Readership



Note: Panels A and B present robustness checks for the newspaper-level effects of paywalls on readership. Panel A replaces the outcome with log page views. Panel B examines log duration using Callaway and Sant'Anna (2021) corrections to account for potential treatment effect heterogeneity under staggered adoption. We include newspaper and year-month fixed effects, and cluster standard errors at the newspaper-level. Panels C and D present estimates for the household-level effect of paywalls on local newspaper readership. Panel C presents our baseline specification. Panel D uses Callaway and Sant'Anna (2021) corrections. Because Comscore households are observed for a single calendar year, the treatment group is restricted to households observed in the same year as their local paywall adoption. We include household and year-month fixed effects, and cluster standard errors at the county-level. Coefficient magnitudes are not directly comparable across the two sets of results. The paper-level analyses examine readership regardless of household location, implicitly weighting toward households that consume more news. The household-level analyses examine only households in a newspaper's county, weighted equally. In Panels A and C, months beyond the event window are top- and bottom-coded following Schmidheiny and Siegloch (2023). As Callaway and Sant'Anna (2021) corrections do not allow for top- or bottom-coding, Panels B and D instead report coefficients for each individual month within the same window, without accumulating endpoints.


Figure A.4: Effect of Paywalls on National Newspaper Readership




Note: Figure presents event study estimates for the effect of local newspaper paywalls on readership of the New York Times, Washington Post, USA Today, and Wall Street Journal. Outcome is the log of time spent by a household at a given newspaper website. For users with paywalled newspapers, we restrict to the year of the paywall. For users with control newspapers, we include all years. Months beyond +6 and before -6 are top- and bottom-coded following Schmidheiny and Siegloch (2023). We include household and year-month fixed effects, and cluster standard errors at the county-level.

Figure A.5: Validation of Knowledge Measures

(a) Perceived vs. Actual Economic Conditions



(b) Economic Knowledge by Education



Note: Panel A presents binned scatter plots of households’ reported perceptions of economic conditions from the Gallup Daily surveys against actual GDP growth, separately for national and local measures. National perceptions are based on responses to a question about whether economic conditions in the country are “getting worse” (-1), “staying the same” (0), or “getting better” (+1). Local perceptions are based on responses to the question about whether local businesses are “reducing” (-1), “not changing” (0), or “expanding” (+1) their workforces. GDP growth is defined using the most frequent time horizon at which each measure is available – quarterly for national GDP and annual for county GDP. Perceived and actual measures are standardized to mean zero and unit variance within the local and national series prior to plotting. Panel B presents binned scatter plots of the constructed economic knowledge measures described in Section 5.2 against respondent educational attainment. For both panels, slopes are estimated via OLS regression with county-clustered standard errors in parentheses.

Table A.1: Demographics of Internet News Readers

	Full Sample	Local Reader	Not Local Reader
<i>Race</i>			
White	0.65	0.82	0.63
Black	0.15	0.10	0.16
Asian	0.04	0.02	0.05
Hispanic	0.16	0.13	0.17
<i>Income</i>			
Under \$25,000	0.23	0.20	0.23
\$25,000–\$75,000	0.45	0.47	0.45
\$75,000–\$100,000	0.12	0.14	0.12
\$100,000+	0.20	0.19	0.20
<i>Age</i>			
18–29	0.13	0.09	0.14
30–44	0.30	0.31	0.29
45–64	0.44	0.45	0.43
65+	0.14	0.15	0.13
Number of Households	1,064,043	124,076	939,967

Note: Table presents demographic characteristics of Comscore households from 2006 to 2020. Column 1 reports averages for the full sample. Column 2 reports averages for local newspaper readers, defined as households who spent at least 10 minutes at their local newspaper’s website during their panel year. Column 3 reports averages for households who spent less than 10 minutes at their local newspaper’s website. Demographics are self-reported. Race and age refer to the head of household, income refers to household income.