Review of Economic Studies (2024) **00**, 1–35 doi:10.1093/restud/rdae049

Advance access publication 6 May 2024

The Impact of Online Competition on Local Newspapers: Evidence from the Introduction of Craigslist

Milena Djourelova
Cornell University, USA and CESifo
Ruben Durante
National University of Singapore, Singapore, ICREA-UPF, Spain, BSE, Spain, CEPR, CESifo and IZA
and
Gregory J. Martin
Stanford Graduate School of Business, USA

First version received November 2021*; Editorial decision September* 2023*; Accepted May* 2024 *(Eds.)*

How does competition from online platforms affect the organization, performance, and editorial choices of newspapers? What are the implications of these changes for the information voters are exposed to and for their political choices? We study these questions using the staggered introduction of Craigslist (CL)—the world's largest online platform for classified advertising—across U.S. counties between 1995 and 2009. This setting allows us to separate the effect of competition for classified advertising from other changes brought about by the Internet, and to compare newspapers that relied more or less heavily on classified ads *ex ante*. We find that, following the entry of CL, local newspapers reliant on classified ads experienced a significant decline in the number of management and newsroom staff, including in the number of editors covering politics. These organizational changes led to a reduction in news coverage of politics and resulted in a decline in newspaper readership, particularly among readers with high political interest. Finally, we document that reduced exposure to local political news was associated with an increase in partisan voting and increased entry and success of ideologically extreme candidates in congressional elections. Taken together, our findings shed light on the determinants of the decline of print media and on its broader implications for democratic politics.

Key words: Newspapers, Internet, Advertising, Ideological polarization

JEL codes: L82, L86, D72

The editor in charge of this paper was Katrine L∅ken.

1. INTRODUCTION

The Internet has profoundly changed the environment in which newspapers operate. Competition from online platforms has contributed to a sharp decline in newspapers' revenues over the last two decades, forcing many news outlets to drastically rethink their business model and organization. These changes, some warn, may have detrimental consequences for the quality of news reporting and the provision of political information (McChesney and Nichols, 2011; Starkman, 2014; Peterson, 2021). Given the key role played by newspapers in informing citizens about their representatives (Besley and Burgess, 2002; Snyder and Strömberg, 2010), they may also have important political implications.

Despite the potentially grave consequences of these transformations for the future of journalism, rigorous evidence on the impact of online platforms on newspapers' organization and editorial choices is surprisingly scant. One reason for this is the challenge of separating the effect of online competition from other technological and socioeconomic changes brought about by the Internet, which may affect both the demand and the supply side of the newspaper market in other ways.

In this paper, we investigate the impact on U.S. newspapers of the introduction of Craigslist (henceforth CL), the world's largest online platform for classified ads. CL's entry disrupted the market for classified ads, formerly a lucrative niche for newspapers (Kroft and Pope, 2014; Seamans and Zhu, 2014), which accounted for 40% of newspapers' advertising revenues and about 30% of total revenues in the year 2000. Tracking the expansion of CL across U.S. counties between 1995 and 2009, we examine how the entry of a local CL website affected the organization, editorial decisions, and content production of local newspapers. We then trace the impact of these changes on local representation in the U.S. Congress.

The expansion of CL in the U.S. provides an attractive setting for several reasons. First, CL's staggered expansion over a period of 15 years, combined with the limited geographic scope of local CL websites, generates significant variation over time and across space in the degree of online competition for classified ads faced by local newspapers. Second, since CL websites do not feature news content or display advertising, CL's entry represents a specific shock to revenues from classified ads but leaves other market conditions unaffected. CL's narrow focus on classified ads provides an additional source of variation, since the entry of CL should disproportionately affect newspapers that relied more heavily on classified ads *ex ante*. Finally, with only a few exceptions in the biggest cities, ads on CL are free of charge, and most local websites do not generate profit for the company. The lack of a clear profit maximization strategy[1] alleviates concerns that the timing of CL's entry might have been driven by strategic considerations related to the conditions of local newspaper markets. We document that the timing of CL's entry into a local market is not correlated with the characteristics of local newspapers once population and the quality of the local Internet connection are controlled for.

To analyse the impact of CL on local newspapers, we collect data on the date of entry of each local CL website, and digitize comprehensive data on the organization and market outcomes of more than 1,500 newspapers (the universe of daily newspapers), covering the period from 1995 to 2010. Our empirical strategy uses two sources of variation. The first source of variation is CL's staggered introduction: we compare the evolution of outcomes of interest between areas with and without access to a local CL website, before and after the website is introduced. To make sure we separate the effect of CL from that of Internet penetration or other correlates of CL entry, we control for the quality of local broadband Internet, log population, and baseline

1. CL founder Craig Newmark was sued in 2010 by eBay, which held a minority stake in CL, for failing in his fiduciary duty to maximize shareholder returns.

county characteristics interacted with time fixed effects. While this variation identifies the average effects of CL entry, only newspapers in the classified ad market are plausibly exposed to competition from CL. Therefore, the second source of variation comes from differences across newspapers in the reliance on classified ads prior to the entry of CL, proxied by the presence of a dedicated classified ads manager in the newspaper's staff. We validate this proxy by documenting that it is strongly correlated with both the share of a newspaper's total pages devoted to classified ads and with the rates charged by newspapers for classified ads before the entry of CL.

We begin our analysis by investigating the "first-stage" impact of the local entry of a CL website on the use of the platform and on newspapers' classified ad business. We verify that local take-up of the CL platform increases after the entry of a local CL website, estimating a significant increase in local visits to the craigslist.org domain. We also find evidence that this take-up leads to a significant decline in the volume of classified ads in local newspapers. Following CL entry, the share of classified pages in local newspapers declines by about 3 percentage points on average, or 10% relative to the sample mean. This effect is driven entirely by newspapers that relied more heavily on classified ads at baseline, *i.e.* those that had a dedicated classified manager. The magnitude of the decline for such newspapers is 4.2 percentage points, or 15% of the mean.[2] Since 25% of newspapers experience the entry of CL and have a classified manager at baseline, this estimate implies that CL led to an overall reduction in the share of classified pages of about 1 percentage point, or 12% of the overall decline observed between 1995 and 2010 (9 percentage points).

Second, we turn to the impact of CL on newspapers' organization, finding substantial downsizing. After the entry of CL into a county, newspapers headquartered in that county cut 1.1 jobs on average, or about 5% relative to the sample mean. This effect is driven entirely by newspapers that relied more heavily on classified ads at baseline, for which we estimate a decline of about 3 jobs (14% of the sample mean). This accounts for 10% of the total decline in newspaper jobs observed between 1995 and 2010 (a 35% decline).

Staff cuts are not limited to advertising managers, but extend to managerial and editorial positions. For editorial staff, a detailed analysis of job titles indicates that the cuts affected political editors disproportionately, while editors covering other areas, such as sports and entertainment, were not significantly affected.

Third, we examine how these organizational changes affected newspapers' editorial priorities, with particular regard to the coverage of politics. To do so, we fit a topic model to a corpus of two million randomly drawn articles (covering ∼850 newspapers in our sample) and examine the distribution of newspapers' coverage across topics. Consistent with the notion that staff cuts are detrimental to the coverage of politics (Peterson, 2021), we find that CL entry is associated with an 8.3% decline in the prevalence of political topics, while "soft news" topics such as sports and entertainment are not affected. We complement this approach with full-text keyword searches for mentions of politicians' names in the universe of articles archived by NewsBank (covering ∼900 newspapers and more than 100M articles). We document that, following the entry of CL, news coverage of local congressional representatives and candidates for congressional offices declines by about 12%.[3] This decline is particularly pronounced for articles published prior to primary elections.

2. We detect no changes in classified ad rates, implying that the quantity loss should translate proportionately into revenue loss. We can hence roughly approximate the resulting revenue loss using typical levels of classified pages and prices in our dataset. Our estimates of page and price changes imply a median revenue loss among classified-manager papers of at least $137K–$298K per year.

3. We document similar declines in the number of mentions of national politicians, such as the president or party leaders in Congress, as well as the number of mentions of titles of local- and state-level politicians.

Fourth, we examine how readers respond to these changes in content. We document that, in the years after the entry of CL, affected newspapers experienced a decline in circulation per capita of about 4%, a finding we confirm using self-reported newspaper readership from two independent large-scale surveys of media consumption. The survey data further suggest that the drop in readership is not driven by readers that are likely to read the classified ads section of newspapers, but by readers interested in general news (among whom the decline in readership is as large as 15%). This finding is consistent with readers responding to changes in news content triggered by the entry of CL, rather than merely lower demand for print classified ads.[4] Evidence from both survey and browsing data suggests that the decline in newspaper readership is unlikely to be compensated by increased news consumption online.

Finally, we study how reduced news coverage of politics—and in particular, reduced coverage of local congressional candidates—affected local electoral outcomes, focusing on turnout, and ideological polarization in congressional elections. We find that the entry of CL has no significant effect on voter turnout. In contrast, we document that it increased the correlation between results in the local House or Senate race and the national presidential race, reflecting a decline in split-ticket voting. This finding is consistent with a greater tendency of voters to rely on national partisan cues when less information about local candidates for office is available.

We also find suggestive evidence that CL entry is associated with an increase in the entry and electoral performance of ideologically extreme candidates in House elections. This result is consistent with the role of newspapers in allowing voters to discriminate between extreme and moderate candidates, particularly in primary elections where voters cannot rely on party labels (Hall and Lim, 2018).[5]

Taken together, our results indicate that the impoverishment of local newspapers due to competition from online platforms can jeopardize their ability to inform citizens about politics, with the effect of reducing the centripetal pressure on candidate ideology that elections provide. This evidence supports the concerns expressed by some regulators that newspapers' financial distress, due to lower advertising revenues, may threaten quality reporting and pluralism (FCC, 2016).

Our paper contributes to several streams of literature. First, it relates to prior work by Seamans and Zhu (2014) on the impact of CL on newspapers' business strategies. Our analysis expands upon their findings by documenting the broader implications of the CL shock on newspapers' organization, editorial priorities, news content, and, ultimately, political outcomes.[6]

Second, our paper relates to studies on the effects of new media technologies on incumbent media. Bhuller *et al*. (2024) and Gavazza *et al*. (2019) study the roll-out of broadband Internet in Norway and in the U.K. respectively, and document large declines in newspapers' print circulation. Gentzkow (2006) and Angelucci *et al*. (2024) study the introduction of television in the U.S., documenting that it presented a significant negative shock to newspapers' readership and advertising. While these papers study broad technology shocks that may simultaneously

4. Lower readership is unlikely to be due to newspapers raising subscription prices, since we find if anything a negative effect of CL entry on that dimension.

5. The classic Calvert–Wittman model of electoral competition with policy-motivated candidates (Wittman, 1983; Calvert, 1985) generates the prediction that greater uncertainty over the preference of the median voter leads to greater divergence in equilibrium platforms. Matějka and Tabellini (2021) microfound this result in a model of rational inattention; when the cost to voters to acquire information rises, voters with more extreme preferences are more willing to keep paying, leading politicians to cater to them more. "That is, as the cost of attention rises, the equilibrium moves closer to the bliss point of the group with more extreme preferences (p. 1,920)."

6. Other studies have used the expansion of CL across the U.S., or particular design features of the platform, to investigate questions related to matching efficiency in labour and housing markets (Kroft and Pope, 2014), and the impact of online personal ads on sexually transmitted diseases and violence against women (Chan and Ghose, 2014; Cunningham *et al*., 2019).

affect readers' demand, advertising markets and the production of news, our setting allows us to separate the effect of a specific shock to the advertising market.

Our findings also dovetail with previous evidence of how shocks to advertising revenues affect news producers. In a historical perspective, Hamilton (2004) and Petrova (2011) argue that the growth of the print advertising market in the late 19th century was essential to the emergence of an independent (non-partisan) press. Our analysis differs in that we study the reverse phenomenon, *i.e.* the decline in newspapers' advertising business due to online competition and its negative implications for political coverage. This complements evidence from France by Angelucci and Cagé (2019), who find that the introduction of advertising on TV in the 1960s and 1970s reduced newspapers' production of journalistic-intensive content.[7]

More broadly, our analysis relates to previous work on the effect of the Internet on electoral politics (Falck *et al.*, 2014; Campante *et al.*, 2018; Larcinese and Miner, 2018; Gavazza *et al.*, 2019). While these studies assess the aggregate effect of the various changes brought about by the Internet, our analysis isolates the political impact of the Internet through the disruption of legacy media.

Finally, we contribute new causal evidence on the link between the media environment and the nationalization of politics (Trussler, 2021; Moskowitz, 2021) and success of extremist candidates (Hall and Lim, 2018). Our results demonstrate how news media support the functioning of elections as a selection mechanism. The magnitude of the electoral penalty that extreme candidates face, and voters' ability to punish extremists electorally (Hall, 2015; Hall and Thompson, 2018), hinges on the quality of the information environment to which voters have access. Our results show how the impoverishment of local newspapers and the resulting changes in organization, content, and readership can limit the electorate's ability to discriminate between candidates and thus weaken the ideologically moderating force of elections.

The remainder of the paper is organized as follows. Section 2 briefly describes our conceptual framework. Section 3 provides background information on CL and its expansion and describes the data used in the analysis. Section 4 discusses the empirical strategy. Section 5 presents the results and evidence on possible mechanisms. Section 6 concludes.

2. CONCEPTUAL FRAMEWORK

Our conceptual framework starts from the assumption that the coverage of politics is costlier than the coverage of soft news, such as sports and entertainment, since it requires greater investments in reporting staff with expertise and more time for investigative activities (Hamilton, 2016).[8] Prior to CL, newspapers' profits from classified advertising cross-subsidized the production of news, in general, and political news, in particular. The entry of a superior online competitor unbundled the two products and reduced newspapers' revenues from classified ads, forcing them to cut staff and limit more costly activities such as political reporting. The reduction in political coverage in turn reduced demand among readers interested in this type of content.

This framework is consistent with a model of a two-sided market, such as that proposed by Angelucci and Cagé (2019), in which advertisers pay for readers' attention, (some) readers

7. Another relevant contribution on the impact of competitive pressures on newspapers is George and Waldfogel (2006). The study documents how the diffusion of the New York Times to over 100 U.S. towns where it was previously not distributed affected the circulation and content of local newspapers.

8. Hamilton argues that this cost differential arises from the combination of two factors. First, the ease of discovering primary information: "press agents are eager to promote stars," while public affairs stories can "require determination to pursue information that government actively tries to keep hidden." Second, the certainty that reporting effort will produce a good story: in sports coverage, "a game's unfolding and outcomes are guaranteed to provide a story," while "accountability coverage is akin to drilling for oil, since tips and suspicions may not pan out." (p. 16)

value political coverage, and producing political coverage involves paying higher fixed costs as it requires a larger newsroom. In such a model, a reduction in advertisers' willingness to pay induces newspapers to reduce the size of the newsroom and produce less political coverage.[9] Intuitively, newspapers have weaker incentives to attract readers interested in political news through quality political coverage when advertisers are less willing to pay for readers' attention. This framework predicts that newspapers affected by CL should: (1) reduce the size of their newsroom (particularly in the area of politics); (2) reduce their coverage of politics; (3) experience a decline in readership (particularly among readers interested in political news).

Given the role of local newspapers as the primary mainstream source of information on local political candidates,[10] a reduction in the volume of newspaper coverage in this domain can have important implications for electoral politics. First, voters may be less likely to participate in elections.[11] Second, with less specific information available on local candidates, voters may rely more on partisan cues when casting their vote, leading to declines in split-ticket voting in general elections.[12] Third, a rise in the cost of accessing information about candidates in mainstream sources should advantage more ideologically extreme candidates. The reason is that politicians cater their platforms to more informed voters (Strömberg, 2001). When political coverage in mainstream sources falls, moderate voters become differentially less informed compared to more extreme voters, because they are on average less willing to pay the costs of seeking out substitutes from less accessible sources (Matějka and Tabellini, 2021). Moderates will thus have a harder time differentiating ideologically extreme candidates from moderate ones. Together with the second implication above, this effect makes it more likely that relatively extreme candidates will be willing to enter electoral contests, and achieve better electoral returns when they do.

3. BACKGROUND AND DATA

3.1. *Background*

Craiglist.org (CL) is the world's largest online platform for classified ads and is consistently ranked among the top 20 websites by traffic in the U.S. The company was created in San Francisco in 1995 and served only the Bay Area until 2000, when it began expanding to other U.S. locations. The expansion started with large cities such as Boston, New York, and Chicago and continued to smaller markets over time, covering 115 locations in 2005, 331 in 2008, and 416 today.

The CL platform is characterized by a simple layout which has remained largely unchanged over time (Figure B1). CL websites only host classified ads and do not include any display ads or news content. The most popular ad categories are housing, jobs, and items for sale.

Posting, browsing and responding to CL ads in the vast majority of locations is completely free of charge. The only exceptions are brokered apartment rentals in New York City and job posts for employers in some major cities, with fees ranging from $3 to $75.[13] Fees for these few

9. This model's prediction for subscription prices, on the other hand, is ambiguous.

10. *E.g.* Mahone *et al*. (2019) show that newspapers produce a disproportionate share of what the FCC calls "Critical Information Needs" (CIN) coverage. This category includes "information on local, regional, and county candidates at all levels of governance (p. 26)." In contrast, local television news tends to have a more national focus—with political coverage heavily tilted towards the presidential level—and provide less substantive information about politicians' policy stances (Fowler, 2020).

11. Evidence from a variety of settings suggests that substitution into media with less political coverage leads to declines in participation; see Gentzkow (2006), Ellingsen and Hernæs (2018) and Campante *et al*. (2018).

12. Other work that examines the relationship of the media environment to split-ticket voting includes Moskowitz (2021), Angelucci *et al*. (2024) and Darr *et al*. (2018).

13. A full list of the exceptions as of 2010 is available at: https://web.archive.org/web/20100706030043/https://www.craigslist.org/about/help/posting_fees.

special categories constitute CL's only source of revenue. This unconventional business model reflects the views of CL's founder, Craig Newmark, who originally founded CL as a non-profit and prioritized providing a useful service to local communities over profit maximization even after the company was incorporated in 1999.[14]

The period of CL's expansion and growth in popularity coincides with a collapse in newspapers' lucrative classified ad business. According to data from the Newspaper Association of America, in the year 2000 U.S. newspapers' classified advertising revenues amounted to $20 billion, relative to $49 billion in total advertising revenues and $11 billion in circulation revenues. By 2012, classified advertising revenues had fallen to $4.6 billion—a 77% decline. Panel (a) of Figure B2 shows that this decline was steeper and occurred earlier than that in other revenue sources. Similarly, the aggregate number of newsroom workers has been on a steep downward trend since the mid-2000 s (panel b).

3.2. *Data*

Our analysis combines data on: (i) CL's expansion across the U.S.; (ii) characteristics, organization and market outcomes of daily newspapers; (iii) newspapers' content; (iv) survey data on media consumption; (v) political outcomes including turnout, voting choices, and candidates' ideology, and (vi) additional covariates, including a proxy for local Internet penetration.

Craigslist's expansion. To construct a measure of the availability of CL in each county, we first collect information on the timing of the entry of each of CL's local websites. For a subset of 197 websites, this information is directly available on CL's "about" webpage (https://www.craigslist.org/about/expansion). For the remaining 219 websites, we assign the date of the first snapshot recorded by the Internet Archive Wayback Machine (https://archive.org), cross-checking this date against past snapshots of CL's list of websites. The maps in Figure 1 show the geographic distribution of local CL websites in 2000, 2005, and 2010, respectively.

To define the area served by a given CL website, we consider the county (or counties) included in the location reported in the website's URL. This is usually a single city, but can also be a combination of nearby cities, a region, or, in a few cases, an entire state. For example, we assume that the website chicago.craigslist.org (Chicago) serves Cook and DuPage counties, and that the website westernmass.craigslist.org (Western Massachusetts) serves Hampden, Hampshire, Worcester, Berkshire, and Franklin counties.[15]

Figure 2 illustrates the timeline of CL's expansion and take-up. Panel (a) plots the evolution of the share of U.S. counties served by a CL website, along with the evolution of a proxy for Internet quality: the average number of Internet service providers by zip code. Panel (b) plots the evolution of CL's traffic, measured by the share of visits to the domain craigslist.org recorded by Comscore, a dataset that tracks browsing behaviour for a representative sample of U.S. online users (more information below). For comparison, we also plot the equivalent measure of traffic for CL's major competitors in the three main categories of classified ads: monster.com for job ads, realtor.com for housing ads, and ebay.com for items for sale. The figure shows that the use of CL evolved differently from that of its competitors and from the general trend in Internet penetration.

14. For a profile of Craig Newmark and his business strategy see https://www.theguardian.com/technology/2006/feb/19/news.theobserver1 and https://www.wired.com/2009/08/ff-craigslist/.

15. As an alternative approach, we consider a broader set of counties: any counties that account for a non-negligible share of the ads posted on the website in year two after its entry (retrieved from the Wayback Machine). We discuss this approach in Supplementary Material, Appendix B.1.1.

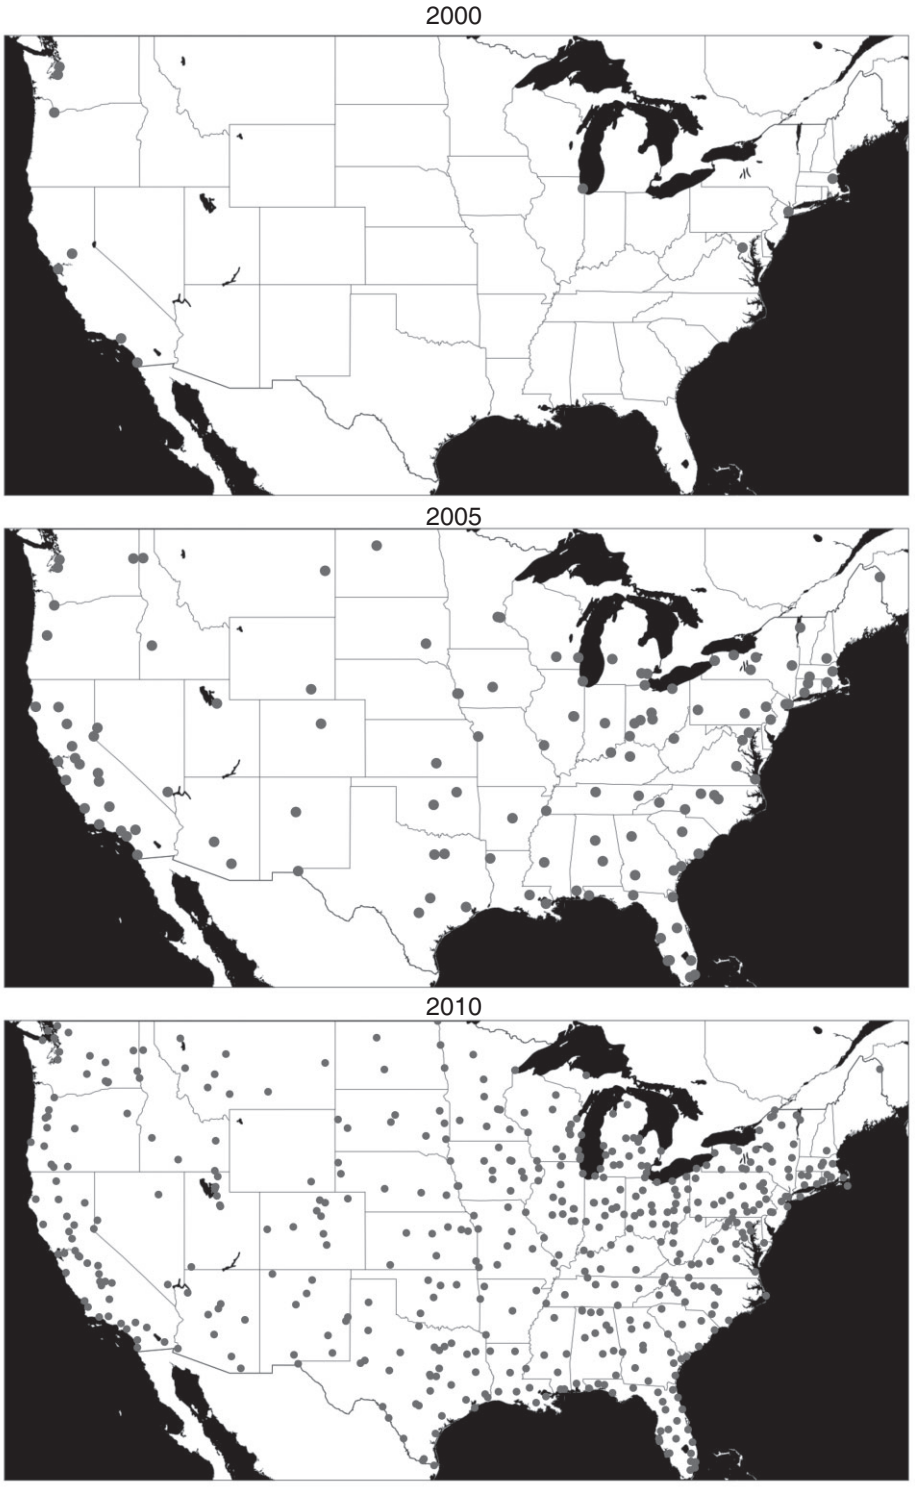

FIGURE 1
Craigslist locations over time

Notes: Geographic distribution of local craigslist.org websites at three points in time

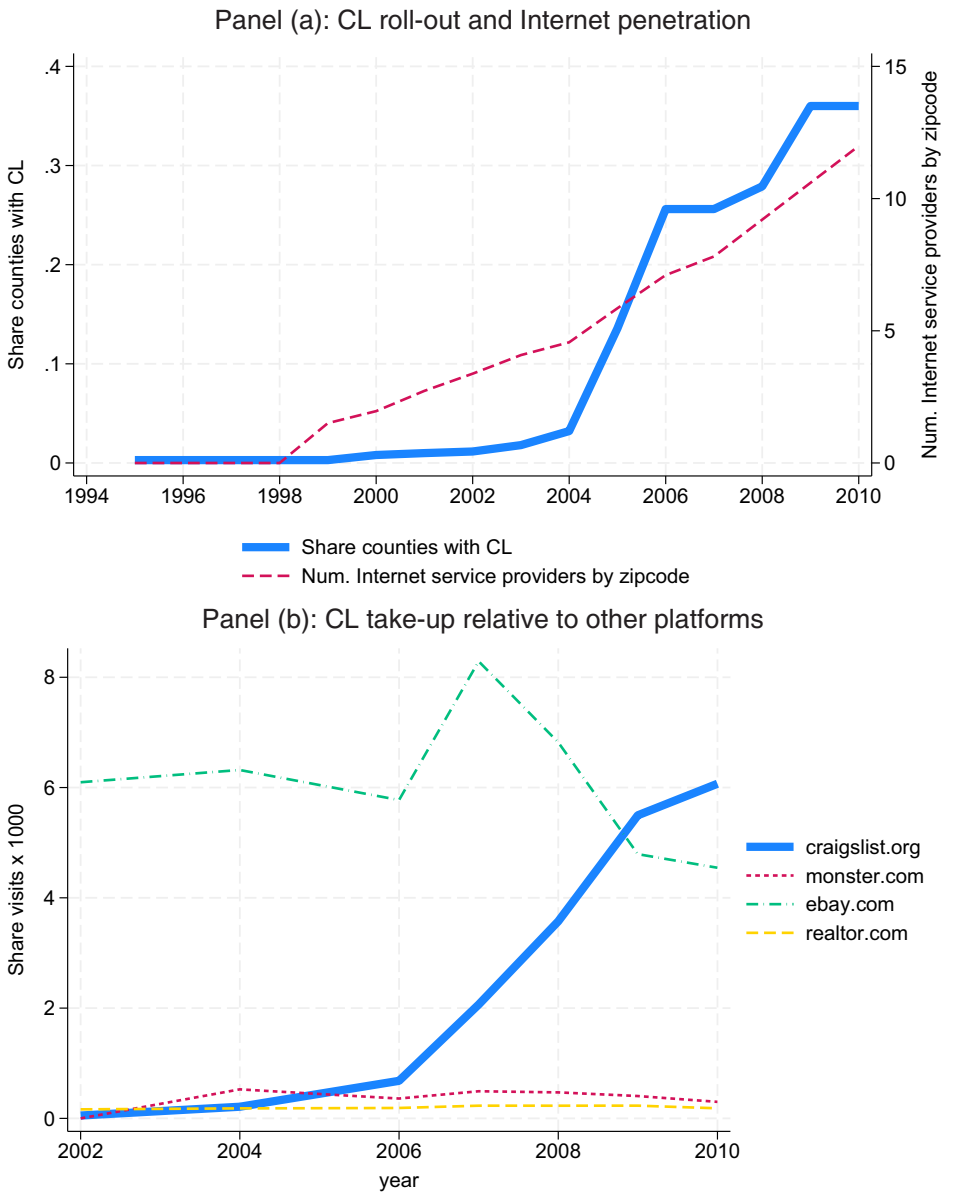

FIGURE 2
Craigslist's roll-out and take-up over time

Notes: Panel (a): Share of U.S. counties with access to a local CL website (left axis) and average number of Internet service providers by zipcode (right axis). Panel (b): Number of visits to craigslist.org, monster.com, ebay.com, and realtor.com as a share of total visits by Comscore panelists

Newspaper characteristics and outcomes. We collect comprehensive data on relevant newspaper characteristics and outcomes from industry yearbooks published by the company Editor & Publisher (E&P, 2010). We accessed print copies of the yearbooks for the period 1995–2010 and

digitized them using OCR software.[16] The yearbooks contain detailed information for over 1,500 U.S. daily newspapers, including: address of the headquarters (HQ), ownership group, circulation, subscription price, as well as the complete list of editorial and managerial staff members with names, broad job categories, and job titles. Figure B3 presents examples of the raw E&P data.

To define newspaper markets, we assume that newspapers' readership is concentrated in the county where the newspaper HQ is located. This approach is common in the literature (Gentzkow and Shapiro, 2010; Seamans and Zhu, 2014) and a good approximation for most newspapers: disaggregated circulation data suggests that the HQ county accounts for 85% of the readership of the median newspaper.[17]

The extent to which a newspaper is affected by the entry of CL depends on how heavily it initially relied on revenues from classified ads. To proxy for pre-CL reliance on classified ads revenue, we consider whether, in the year 2000 (prior to CL's major expansion), the newspaper had a dedicated classified ads manager in its staff. To validate this measure and to estimate the size of the revenue shock to affected newspapers, we collect information from newspapers.com—an online archive of historical newspapers—on the share of pages per issue devoted to classified ads (available for a subset of about 250 newspapers for the period 1995–2010). We also collect information from Standard Rate and Data Service (SRDS , 2006) on classified prices (available for all newspapers for the period 1995–2006).

We correlate these measures of classified reliance with the presence of a classified ads manager prior to CL's entry. The results, presented graphically in Figure 3 and in tabular form in Table B1, indicate that, controlling for a broad set of characteristics of the newspapers and of the county in which they are headquartered, newspapers with a classified ad manager devoted about 7 percentage points more pages to classified ads and had classified ad prices about 9% higher than newspapers without a classified ad manager.

Newspapers' content. We obtain data on newspapers' content from NewsBank (2010), which contains the full text and metadata of more than 100 million articles published in about 900 newspapers in our sample. We use the data in two ways. First, we perform keyword searches on the full text of all articles looking for names of specific politicians (*e.g.* "Rep. Paul Ryan," "Senator Dianne Feinstein," etc.), and use the number of mentions in a given newspaper-year as a measure of the amount of coverage it devotes to specific groups of politicians. Second, we extract the text of the lead paragraphs for a random sample of 2 million articles and estimate a topic model on the resulting corpus. Specifically, we estimate a Correlation Explanation (CorEx) model (Gallagher *et al*., 2017), which has the advantage of producing coherent topics for corpora consisting of short texts. We obtain a distribution of topic weights for each article in the corpus, which we then aggregate up to the newspaper-year level. Further details about the procedures used to construct these variables are reported in Supplementary Material, Appendix B.1.4.

Internet penetration and other county controls. To control for local Internet penetration, following Larcinese and Miner (2018), Seamans and Zhu (2014) and Lelkes *et al*. (2017), we use the number of Internet service providers (ISPs) registered by zip code and year. These data are available for the period 1998–2008 from the Federal Communication Commission (FCC, 2008) and cover all providers with more than 250 high-speed lines in a state and transfer speed greater than 200 kilobits per second. We assign zero ISPs to all zip codes for the years before 1998, and

16. To alleviate transcription errors, we flagged observations for which our key variables deviate substantially from past/ future values for the same newspaper, and interpolated flagged or missing values.

17. In Supplementary Material, Appendix B.1.2, we describe an alternative approach based on identifying the broad set of counties where a given newspaper circulates using geographically disaggregated circulation data. These data, however, are only available for a subset of the newspapers in the E&P sample.

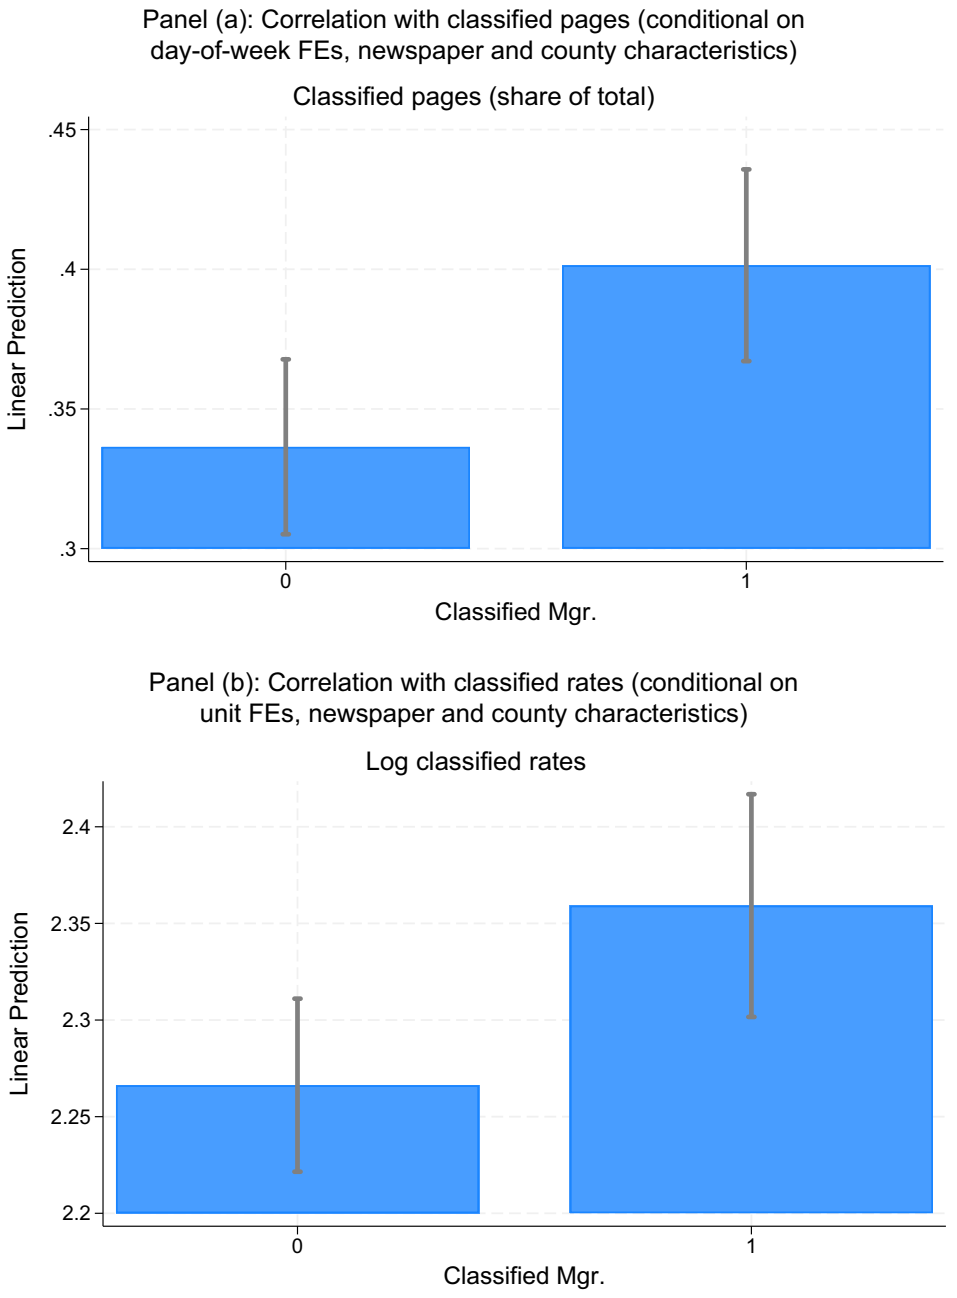

FIGURE 3

Classified manager as a proxy for reliance on classified ads

Notes: Linear predictions for the share of classified pages (Panel a) and for log classified rates (Panel b) in newspapers with and without a classified manager at baseline, conditional on covariates. The covariates include circulation, jobs-count, log population, number of Internet service providers, share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout, and share White/Black/Hispanic. Classified pages are pre-2000 newspaper × day-of-week specific averages, scaled by total pages per issue (recorded by E&P). Classified rates are pre-2000 newspaper × unit specific averages. Panel (a) additionally controls for day-of-week FEs and is based on regressions weighted by the number of issues sampled. Panel (b) additionally controls for unit FEs

use linear interpolation to fill missing data for years after 2008. We then aggregate the number of ISPs at the county level by taking the population-weighted average across all zip codes in a county.

While disaggregated data on the number of Internet subscribers is unavailable for the period of interest, prior studies show that the number of ISPs is a strong predictor of the number of subscribers at the state level, as well as at the county level in later periods. As further validation, we examine the correlation of the number of ISPs with self-reported measures of access to the Internet from two large-scale electoral and marketing surveys (described in greater detail below). The results, reported in Figure B4, show a strong positive relationship between the number of ISPs and the share of local respondents who report accessing the Internet regularly.

Throughout our analysis, we also use data on population by county and year (from the National Center for Health Statistics), and on a series of county-level variables in the baseline census year 2000 including: income per capita, share of the population in urban areas, share of the population with college education, share of the population who rent housing, racial composition and median age (all from the 2000 Census) unemployment rate (from the Bureau of Labor Statistics), and presidential turnout (from the David Leip Election Atlas).

Surveys and browsing data. We use individual survey data on self-reported media consumption from two large-scale surveys. The second source is the Survey of the American Consumer conducted by GfK Mediamark Research & Intelligence (GfK-MRI, 2010), a nationally representative rolling cross-sectional marketing survey conducted every year in the period 1999–2010. The second source is the National Annenberg Electoral Survey (NAES, 2008), a nationally representative rolling cross-sectional electoral survey conducted in the lead-up to the 2000, 2004, and 2008 presidential elections.[18] In addition to questions on the frequency and type of media the respondent consumes, GfK-MRI contains information on the newspaper sections read by the respondent, *e.g.* general news or classified advertising. This allows us to study the characteristics of news- versus classified-section readers.

We use additional data on online media consumption from WRDS Comscore (Comscore, 2010)—a dataset that tracks the browsing behaviour of a large sample of U.S. Internet users representative of online buyers (los Santos *et al.*, 2012). The data cover the years 2002, 2004, and 2006–2010. We aggregate the number of visits to domains of interest as well as total visits recorded by Comscore by county and year.

Political outcomes. We obtain county-level electoral data for House, Senate and presidential elections from the David Leip Atlas of American Elections (Leip, 2020). Following Darr *et al.* (2018), we use these data to construct a measure of split-ticket voting, defined as the absolute value of the difference between the Republican candidate's vote share in the presidential election and the Republican candidate's vote share in House and Senate elections in the same county and year.[19]

We also examine a set of outcomes related to the entry and electoral performance of ideologically extreme candidates. Following Hall (2015) and Autor *et al.* (2020), we classify candidates using information on contributors to their campaigns. Specifically, we use candidates' position in the distribution of campaign-finance-based ideological scores (CFScores) from the Database on

18. The 2000 and 2004 waves of the NAES include questions on general media exposure, while the 2008 wave only asks about exposure to campaign-specific information. We therefore focus the analysis of media consumption on the first two waves.

19. Darr *et al.* (2018) use the Senate-President difference, but since only a third of Senate seats are contested in each election cycle, also using House races expands the number of observations available. We measure vote shares for each office at the county level; House (Senate) shares are computed by aggregating across all House (Senate) votes cast by voters in the county.

Ideology, Money in Politics, and Elections (Bonica, 2016). Bonica's method is an unsupervised learning technique that recovers candidates' positions in a one-dimensional latent ideological scale under the assumption that donors give "spatially;" *e.g.* that a contribution from donor i to candidate j is more likely, the closer are i and j in the latent ideological space. Importantly, CFScores are fixed within candidate, meaning that change in the distribution of CFScores of candidates in a given district over time can come only from the entry of new candidates. We use the 25th and 75th percentiles of the distribution of scores for all House candidates in 2000 as the thresholds separating "extremists" from "moderates." The main outcomes of interest are the presence of an extremist in a primary election, the presence of an extremist in a general election (which implies that at least one primary election in that district was won by an extremist), and the probability that a general election is won by an extremist.

4. EMPIRICAL STRATEGY

4.1. *Determinants of CL entry*

To implement our empirical strategy, it is necessary first to understand what factors drove the timing of CL's staggered rollout. According to CL founder Craig Newmark, CL's entry into new locations was determined by market demographics and the quality of local broadband Internet, but apart from that was largely idiosyncratic.[20]

Importantly for our purposes, CL is not in the news business, and is thus unlikely to have considered demand-side factors in the news market. Furthermore, various accounts call into question the extent to which CL attempted to maximize profits.[21] The lack of a profit-maximization motive arguably allowed more flexibility for the idiosyncracies of CL's founder and its early user base to determine the timing of the roll-out, rather than systematic economic factors in local news markets.

To test this conjecture, in Table 1 we examine the set of counties where at least one newspaper was headquartered in 2000, and regress the year of CL's entry on county and newspaper characteristics, expressed in levels in panel (a) and in changes between 1995 and 2000 in panel (b). The results confirm that population and the quality of the local Internet connection were important factors in CL's entry decision. The magnitudes are sizable: a one-standard deviation higher log population (number of ISPs) is associated with CL entering a county about 15 months (16 months) earlier. Together, these two variables account for 37% of the variation in year of entry. Given this strong relationship, we will control for time-varying log population and number of ISPs throughout our analysis. A few other covariates, such as the rental share of housing, median age, and the county's racial composition, display weaker correlations with CL's time of entry (see Table A1). To control flexibly for possible secular trends associated with these factors, in

20. According to Newmark: "We put up a city based on how many people are asking us to do so. It's also based on Jim's [Jim Buckmaster, CL's CEO] perception of a city's demographics and the city's broadband penetration and intuition. We use word of mouth to get the word out, though sometimes the local press is kind··· We just wait for things to happen." See "CRAIGSLIST/On the record: Craig Newmark." San Francisco Chronicle, 15 August 2004 (https://www.sfgate.com/business/ontherecord/article/CRAIGSLIST-On-the-record-Craig-Newmark-2733312.php).

21. In the ruling of a 2010 civil action against Newmark by eBay, the judge concluded that "Craigslist does not expend any great effort seeking to maximize its profits or to monitor its competition or its market share." (eBay Domestic Holdings Inc. v. Newmark, Delaware Court of Chancery Civil Action No. 3,705-CC, decision dated 9 September 2010, https://h2o.law.harvard.edu/cases/3472). The absence of profit maximization motivated the suit, as eBay (which held a substantial stake in CL) argued that founder Craig Newmark had failed his fiduciary duty to maximize returns to shareholders.

TABLE 1

Correlates of year of CL entry

	Dependent variable: Year of CL entry					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel (a): Correlates in levels						
Newspaper jobs	−0.004	−0.003				
	(0.003)	(0.003)				
Newspaper circulation per capita			−0.224	0.205		
			(0.512)	(0.687)		
Newspaper classified manager					−0.041	−0.027
					(0.122)	(0.121)
Log population	−0.443***	−0.466**	−0.534***	−0.510**	−0.508***	−0.523***
	(0.150)	(0.184)	(0.176)	(0.247)	(0.150)	(0.190)
Internet service providers	−0.647**	−0.389**	−0.659**	−0.403**	−0.670**	−0.401**
	(0.289)	(0.174)	(0.299)	(0.173)	(0.278)	(0.168)
Other county characteristics	No	Yes	No	Yes	No	Yes
Observations	616	616	616	616	616	616
R^2	0.38	0.45	0.38	0.45	0.38	0.45
Panel (b): Correlates in changes						
Δ Newspaper jobs	−0.003	−0.000				
	(0.005)	(0.005)				
Δ Newspaper circulation per capita			0.395	0.488		
			(0.844)	(0.727)		
Δ Newspaper classified manager					−0.017	−0.032
					(0.114)	(0.122)
Δ Log population	−4.176***	−2.759	−4.163**	−2.725	−4.184**	−2.741
	(1.180)	(2.500)	(1.852)	(2.519)	(1.847)	(2.514)
Δ Internet service providers	−0.953***	−0.587***	−0.953***	−0.587***	−0.953***	−0.587***
	(0.089)	(0.119)	(0.210)	(0.120)	(0.209)	(0.120)
Other county characteristics	No	Yes	No	Yes	No	Yes
Observations	615	615	615	615	615	615
R^2	0.35	0.42	0.35	0.42	0.35	0.42

Notes: Regressions of year of CL entry on county characteristics in the year 2000, or changes in county characteristics between 1995 and 2000. Standard errors clustered by CL−area. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

our baseline specification we interact a broad set of baseline county characteristics with year dummies.

Importantly for our analysis, we find no relationship between the timing of CL's entry and the state of local newspapers at baseline as measured by circulation, number of jobs, or the presence of a dedicated classified manager, once log population and the number of ISPs are controlled for. These results corroborate the view that CL did not prioritize areas where newspapers were under-performing or had greater dependence on classified ads.

4.2. *Main specifications*

To estimate the effect of CL entry, we employ a difference-in-differences strategy exploiting CL's staggered introduction across U.S. counties, combined with differences across newspapers in ex-ante reliance on classified advertising. The sample includes the universe of daily newspapers covered by E&P (and respectively, all counties with newspaper HQs), excluding newspapers

with national circulation—the New York Times, U.S. Today, and the Wall Street Journal—that do not serve a specific local market.[22]

The following equations summarize our econometric strategy:

$$Outcome_{nct} = \alpha_n + \beta_t + \gamma\, PostCL_{ct} + \theta Controls_{ct} + \epsilon_{nct}, \quad (1)$$

$$Outcome_{nct} = \alpha_n + \beta_t + \delta PostCL_{ct} + \lambda PostCL_{ct} \times ClassifMgr_n + \theta Controls_{ct} + \epsilon_{nct} \quad (2)$$

$Outcome_{nct}$ is an outcome of interest for newspaper n, headquartered in county c, at time t; $PostCL_{ct}$ is a an indicator equal to one for years after the entry of CL in county c; $ClassifMgr_n$ is an indicator for the presence of a dedicated classified ad manager at baseline (*i.e.* in the year 2000). The vector $Controls_{ct}$ includes time-varying log population and number of ISPs, as well as county-level controls from the 2000 census interacted with year fixed effects, including share of urban population, share with college degree, rental share of housing, income per capita, unemployment rate, median age, turnout and racial composition. α_n and δ_t are newspaper and year FEs, and ϵ_{nct} is the error term. We cluster standard errors by the area affected by the entry of a given CL website (*i.e.* a single county or group of counties), or, for newspapers never affected by CL, by county.

The identifying assumption of our estimation strategy is that, conditional on controls, the timing of CL entry is uncorrelated with pre-existing trends in the outcomes (specification (1)), or with pre-existing differential trends between newspapers with or without a classified manager (specification (2)). Table 1 lends support for this assumption by showing that the timing of CL entry is unrelated to levels or trends in newspaper outcomes or to the presence of a classified manager once we condition on Internet penetration and log population.

To further examine the timing of the effects, we estimate dynamic specifications following recent methodological advances in the difference-in-differences literature (Roth *et al.*, 2023). A concern raised by this literature is that, in settings with treatment effect heterogeneity over time or across treated units, dynamic models estimated using the standard two-way fixed effects estimator can estimate non-convex combinations of treatment effects. This concern applies to our setting since network effects in CL's local adoption predict heterogeneity over time, and newspapers' differential ex-ante reliance on classified ads predicts heterogeneity across units. Therefore, we estimate dynamic effects and first-difference pre-treatment placebos based on the DID_M estimator proposed by de Chaisemartin and D'Haultfoeuille (2020) which is robust to treatment effect heterogeneity. We consider two versions: one that defines the treatment as CL entry (akin to specification (1)) and one that defines the treatment as CL entry into the market of a newspaper with a classified manager (akin to specification (2)). In other words, in the latter version we consider newspapers without a classified manager as untreated.

An important difference between specifications (1) and (2) is that they identify the effect of CL entry for different groups of newspapers. In specification (1), the coefficient γ captures the average impact across all newspapers in the sample. In specification (2), the coefficient λ captures a more local effect—the differential impact on newspapers that are reliant of classified ads ex ante, and hence more vulnerable to shocks to classified ad revenue. Since we are interested in the effect of exposure to the revenue shock rather than the effect of CL per se, the latter

22. The sample includes newspapers in counties where CL never entered, which are about 45% of the total and serve as a control group. The existence of a stable group of "never treated" units is a useful feature for the de Chaisemartin and D'Haultfoeuille (2020) method which we use to estimate dynamic effects robust to treatment-effect heterogeneity. That said, we also present results excluding never-treated newspapers and exploiting only variation in the timing of CL's entry across newspapers that were eventually treated.

approach comes closer to our estimand of interest. Additionally, this approach allows us to conduct robustness checks that relax the assumption on the conditional exogeneity of CL's entry by restricting the comparison to newspapers in the same location—*i.e.* subject to the same location-specific trends—but with different ex-ante reliance on classified revenues. Yet, this requires the additional assumption that differences in CL's effect between newspapers with and without a classified manager are indeed attributable to different classified reliance, and not to other characteristics (*e.g.* newspapers' size). We discuss robustness checks probing the validity of this assumption in Section 5.

For outcomes measured at the county level, we estimate specifications equivalent to (1) and (2), but replace newspaper FEs with county FEs and aggregate $ClassifMgr_n$ up to the county level by taking the circulation-weighted average across newspapers headquartered in the county, using circulation as of the year 2000. Note that this aggregation only applies to counties where newspapers are not local monopolies—which is the case for 14% of the counties in our sample.

For outcomes measured at the congressional district level, we adapt the above specifications to the level of county × congressional district cells, weighting observations by the share of the cell's voting-age population relative to the respective district and clustering standard errors by district.[23] It is worth noting that since outcomes in this case are measured at a higher level of aggregation than the treatment (our sample covers about 1,200 counties in 440 electoral districts), estimates in this specification may be downward biased. Indeed, the misalignment between units of observation inevitably reduces precision as districts contain a mixture of treated and untreated newspapers (*i.e.* ones that experience CL entry and ones that do not, and ones that are more and less reliant on classified ads ex ante).

5. RESULTS

5.1. *First stage: CL take-up and effects on the classified ad market*

The premise of our conceptual framework is that local CL entry triggers take-up of the platform by local consumers and thereby causes a decline in local newspapers' classified ad revenues.

While in theory CL websites have narrow geographic scope, in practice there are no restrictions on the locations of ads or of users. It is therefore an empirical question whether CL's local entry produces meaningful variation in local take-up. To verify that it does, we use data on web browsing behaviour from Comscore (described in Section 3.2) and estimate the effect of CL's local entry on the (ihs-transformed) visits to the craigslist.org domain by county and year. Figure 4 shows an event-study for this effect, estimated following the method proposed by de Chaisemartin and D'Haultfoeuille (2020). The figure presents first-difference placebos for the pre-entry period, and cumulative dynamic effects (corresponding to long differences relative to $t = -1$) for the post-entry period. It suggests a significant increase in visits after the entry of a local website, of about 50% relative to the last pre-entry period. Local take-up increases over time, consistent with network effects in the adoption of the website: a larger number of local users (and hence, a higher volume of local ads) makes the platform more useful and attracts yet more users.

23. To absorb variation due to changing congressional district boundaries, we further include district by redistricting regime fixed effects in all regressions. The major redistricting event in our sample period occurs following the decennial census in 2000, after which all states redrew district boundaries. A handful of states (North Carolina and Virginia in 1997, Texas in 2003, and Georgia in 2005) had additional significant district boundary changes, which we include as well. An example district-redistricting regime fixed effect would be GA-04-2005, which is treated as distinct from GA-04-2000. These fixed effects thus ensure that comparisons in the regressions are within invariant district boundaries.

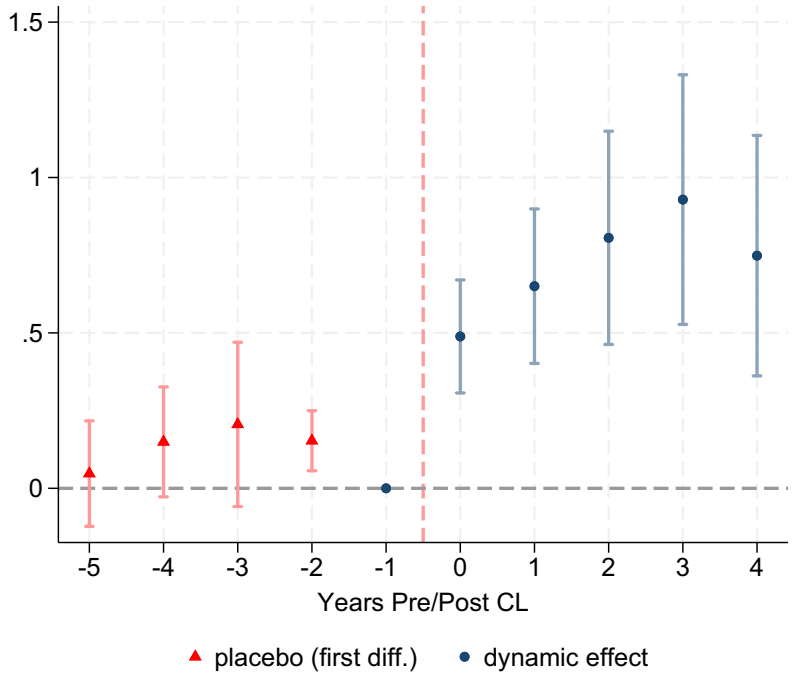

FIGURE 4
Visits to craigslist.org—Event Study

Notes: Pre-treatment placebos and dynamic effects of the entry of a local CL website on the (ihs-transformed) number of visits to craigslist.org. The graph presents coefficients and 95% confidence intervals based on the DID_M estimator proposed in de Chaisemartin and D'Haultfoeuille (2020). Pre-treatment coefficients correspond to first-differences between consecutive periods. Post-treatment coefficients correspond to cumulative effects for the respective length of exposure to CL. Controls include total Comscore visits (ihs-transformed), log population, and number of Internet service providers. Standard errors clustered by CL-area

Next, we examine the effect on the newspaper classified business. In Table 2, we estimate effects of CL entry on the quantity (measured in share of pages devoted to classifieds) and prices of classified ads. We find on average a 3 percentage point decline in the share of pages to classifieds (relative to a mean of 28 percent), and no effect on classified prices. This effect is driven entirely by papers with a classified manager at baseline, where the decline is 4.2 percentage points (corresponding to 15% of the mean). We find no effect on the prices of classified ads,[24] implying a proportional impact on revenue. Together with Figure 3, which shows that these papers were more reliant on classified ads at baseline, Table 2 confirms that the entry of CL was a revenue shock with heterogeneous effects depending on baseline classified reliance.[25]

24. This result differs from that found by Seamans and Zhu (2014), who document a significant decline in classified rates for newspapers affected by CL. Several factors could explain the discrepancy between our results and theirs. First, our sample covers a larger number of newspapers (about 1,400 per year versus less than 600 per year in Seamans and Zhu (2014)) and a longer time period (1995–2006 versus 1999–2006 in Seamans and Zhu (2014)). Second, our empirical strategy is different in that we use the presence of classified managers in the pre-Cragslist period as proxy for newspapers' prior reliance on classifieds, while they use, for each year, the presence of classified managers in that year, which is likely to be endogenous to the entry of CL. Finally, we recognize and try to control for inconsistencies in the unit of measure of classified rates in the SRDS data (*i.e.* column inch, line or word), an issue that their analysis ignores.

25. The median classified-manager paper in our data in 2,000 has about 28 pages per issue, and the median price reported in 2,000 among classified-manager papers is $17.42 per column inch. Classified section layouts range from 6

TABLE 2

Effect of CL entry on classified ad quantities and prices

	Share classified pages		Log classified rates	
	(1)	(2)	(3)	(4)
Post-CL	−0.027**	−0.008	0.009	0.014
	(0.013)	(0.014)	(0.011)	(0.015)
Post-CL × Classified Mgr.		−0.042**		−0.010
		(0.021)		(0.019)
Baseline controls × Year FEs	Yes	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes	Yes
Newspaper FEs, Year FEs	Yes	Yes	Yes	Yes
Newspaper × Day-of-Week FEs	Yes	Yes	No	No
Newspaper × Unit FEs	No	No	Yes	Yes
Observations	6,368	6,352	14,980	14,869
Number of newspapers	255	254	1,407	1,395
R^2	0.76	0.76	0.98	0.98
Mean dependent variable	0.28	0.28	2.47	2.47

Notes: Regressions of share of classified pages and log classified rates on an indicator for the availability of a local CL website and its interaction with an indicator for the presence of a classified manager at baseline. Share classified pages is defined for the Newspapers.com subsample; classified rates are collected from SRDS but available for all E&P papers. The outcome is the average value per newspaper-weekday-year (for page share) or newspaper-unit of measure-year (for classified rate). Baseline controls include the share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout and share White/Black/Hispanic, all measured in the year 2000. Columns 1 and 2 additionally control for pages per issue recorded by Newspapers.com. Standard errors clustered by CL-area. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5.2. *Main newspaper outcomes*

Main results. We present the main results on the impact of CL entry on local newspapers' organization, readership and political coverage in Table 3. The specifications follow equations (1) and (2) in alternating columns. The first dependent variable is the number of jobs listed in E&P's staff section by newspaper and year. We find that CL-entry is associated with a significant decline of 1.1 fewer jobs on average, or 5% relative to the mean of 21.4 jobs (column 1). Turning to our main specification which allows the effect of CL entry to depend on newspapers' baseline reliance on classified ads, we find that this decline is driven entirely by newspapers with a classified manager at baseline (column 2). Comparing newspapers with a classified manager to ones without, the magnitude of the decline reaches three fewer jobs, or 14% relative to the mean. Conversely, the coefficient estimated for newspapers without a classified manager is small, insignificant and if anything, positive. The pattern is similar for CL's effect on newspaper readership. We find a significant decline in circulation per capita, with a magnitude of 2% on average, and 4% comparing newspapers with a classified manager at baseline to ones without (columns 3 and 4).

We also find evidence for significant declines in the volume of newspapers' coverage of politics using two alternative measures. One measure is the topic weight on politics obtained from a CorEx model estimated on a random sample of 2 million articles. We find a 4% average

to 12 columns per page, and broadsheet newspaper page lengths in the U.S. are either 21 or 22.75in. Taking the ends of these ranges yields an approximate loss of between roughly $2,650 and $5,725 per issue; with 52 Sunday issues (the traditionally heaviest day for classifieds) per year, our share-of-pages effect estimate implies a median revenue loss among classified-manager papers of at least $137K–$298K per year.

TABLE 3
Main newspaper outcomes

	Number of jobs		Circulation per capita		Politics coverage topic weight		Congress coverage names count (ihs)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	−1.114***	0.322	−0.004***	0.000	−0.012**	−0.000	−0.028	0.034
	(0.396)	(0.472)	(0.001)	(0.002)	(0.005)	(0.006)	(0.036)	(0.048)
Post-CL × Classified Mgr.		−3.027***		−0.008***		−0.024***		−0.119**
		(0.593)		(0.003)		(0.008)		(0.055)
Baseline controls × Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs, Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	22,377	22,177	22,543	22,316	7,139	7,067	7,430	7,375
Number of newspapers	1,451	1,438	1,454	1,439	863	855	884	878
R^2	0.91	0.91	0.98	0.98	0.52	0.52	0.93	0.93
Mean dependent variable	21.38	21.44	0.19	0.19	0.29	0.29	5.51	5.51

Notes: Regressions of selected newspaper-level outcomes on an indicator for the availability of a local CL website and its interaction with an indicator for the presence of a classified manager at baseline. The dependent variables are number of jobs (columns 1 and 2), circulation per capita (columns 3 and 4), weight on political coverage estimated from a CorEx topic model (columns 5 and 6), and the (ihs-transformed) number of articles containing the name of a congressional representative or candidate from the newspaper's state (columns 7 and 8). Baseline controls include share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout and share White/ Black/ Hispanic, all measured in the year 2000. Columns 7 and 8 additionally control for the (ihs-transformed) total number of total articles in relevant sections in the Newsbank database; see notes in Table 5 for details. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

decline in the topic weight on politics, and an 8% decline comparing newspapers with a classified manager at baseline to ones without (columns 5 and 6). A complementary measure is the number of articles that mention the name of a congressional representative or candidate running in the state in which the newspaper is headquartered. Looking at the ihs-transformed count of such articles, and controlling for the ihs-transformed count of total articles covered by NewsBank by newspaper and year, we find no significant effect of CL entry on average, but a significant 12% decline comparing newspapers with a classified manager at baseline to ones without (columns 7 and 8).

Figure 5 presents estimates for the dynamic effects of CL entry on these four main outcomes, following the method of de Chaisemartin and D'Haultfoeuille (2020). Placebo estimates for the pre-entry period (corresponding to first differences between consecutive pre-periods) are mostly insignificant, except for the 2nd and 3rd pre-period estimates for jobscount. This suggests some pre-trend in jobscount, though placebos for all further pre-periods are insignificant and close to zero.[26] The figure also shows that the estimated post-entry effects (corresponding to long differences relative to $t = -1$) tend to increase in magnitude with length of exposure—a pattern that is consistent with the dynamics of local CL take-up.

Robustness. These main results are robust to a number of alternative specifications. First, we consider alternative sets of controls. In Table A2, we present results without any controls (panel a), as well as results controlling for log population and number of Internet service providers but excluding the baseline controls × year FEs (panel b). Consistent with the notion that Internet penetration may have independent detrimental effects on newspaper outcomes, omitting the control for number of ISPs results in estimates of somewhat larger magnitude for some outcomes. In panel (c) on the other hand we include an additional control for Internet penetration: the average share of respondents with self-reported Internet access in the (pooled) GfK-MRI and NAES surveys by county and year. Even though this specification is valid only in county-years where a survey was conducted, which restricts the sample size significantly, our estimates become if anything more precise with the inclusion of this control. This supports the interpretation of the estimated effects as specific to CL-entry and separate from the generic effects of Internet penetration.

Second, we address the concern that CL entry may be correlated with other location-specific shocks or trends by introducing location × year fixed effects. In Table A3, we introduce fixed effects at the level of state × year (panel a) and fixed-effects at the level of Designated Market Area (DMA) × year (panel b).[27] In both cases the main estimates remain very similar in magnitude and precision. The strongest test for a possible confounding role of location-specific factors is the inclusion of county × year FEs, which absorb any county-level trends or shocks.[28] This specification is identified only from variation in counties that contain at least one newspaper with a classified manager at baseline and at least one without, so that newspapers that are local monopolies in a county (the vast majority of U.S. newspapers) drop out of the sample. This leaves us with 168 counties and 412 newspapers, for which we can make *within-county* comparisons before and after the entry of CL, comparing newspapers with a classified ad manager to their neighbors without. Despite the demanding specification and the smaller sample size, the

26. This method allows us to estimate up to 10 pre-treatment placebos.

27. DMA is the definition of media market adopted by the FCC in regulations governing cross-ownership of newspapers and television stations. They typically contain multiple counties, including core metropolitan counties plus surrounding rural counties. There are 210 DMAs in the U.S.

28. County × year FEs also absorb the main effect of the post-CL indicator, as well as log population, number of ISPs and all baseline controls interacted with year FEs.

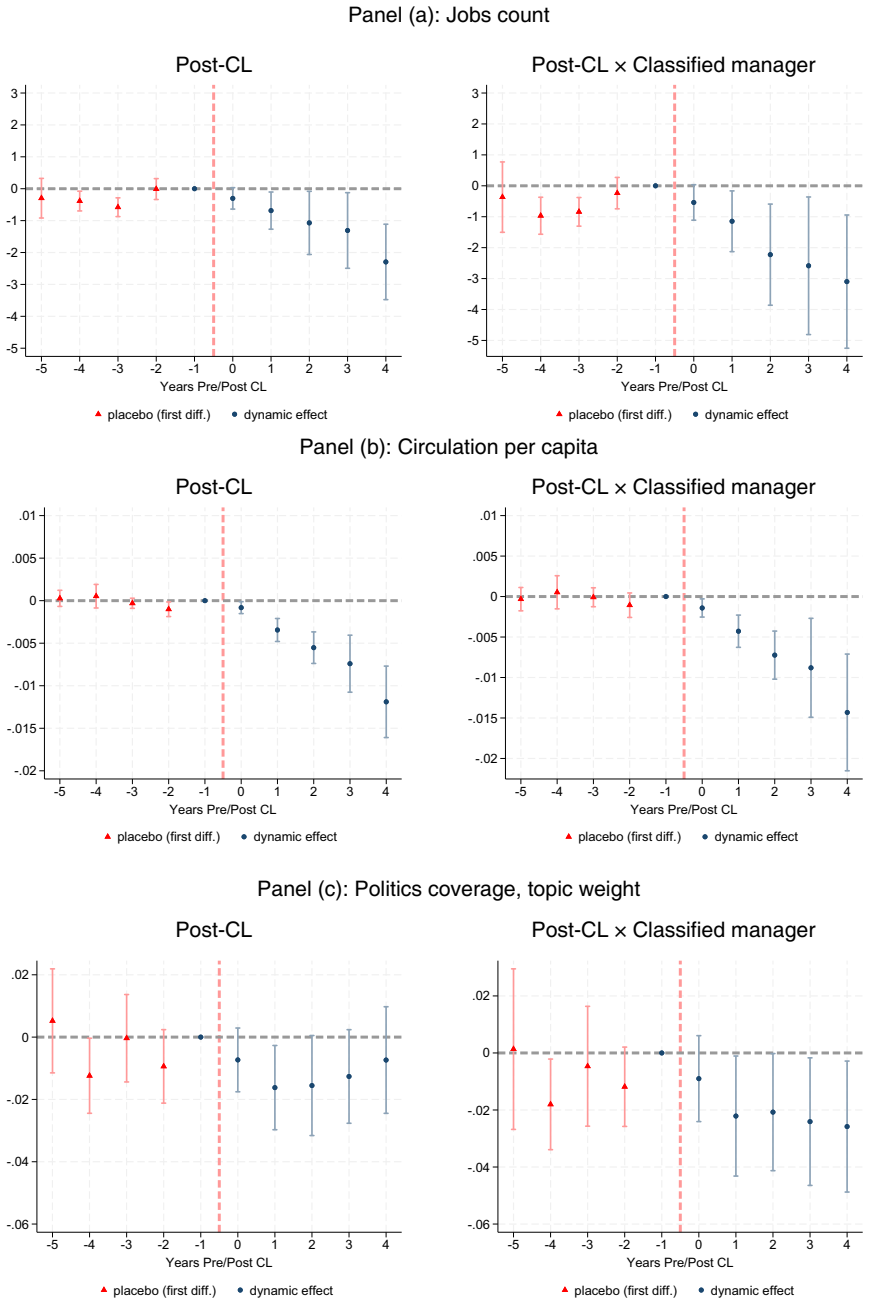

FIGURE 5
Main newspaper outcomes—Event studies

Notes: Pre-treatment placebos and dynamic effects of the entry of a local CL website on newspaper outcomes. The graphs present coefficients and 95% confidence intervals based on the DID_M estimator proposed in de Chaisemartin and D'Haultfoeuille (2020). Pre-treatment coefficients correspond to first-differences between consecutive periods. Post-treatment coefficients correspond to cumulative effects for the respective length of exposure to CL. The dependent variables are number of jobs in panel (a), circulation per capita in panel (b), estimated weight on political topics in panel (c), and the (ihs-transformed) number of articles containing the name of a congressional representative or candidate in panel (d). The left-hand panels define the treatment as any CL entry. The right-hand panels define the treatment as CL entry into the market of a newspaper with a classified manager at baseline. Controls include log population and number of Internet service providers. Panel (d) additionally controls for the (ihs-transformed) total number of articles in relevant sections recorded by Newsbank. Standard errors clustered by CL-area

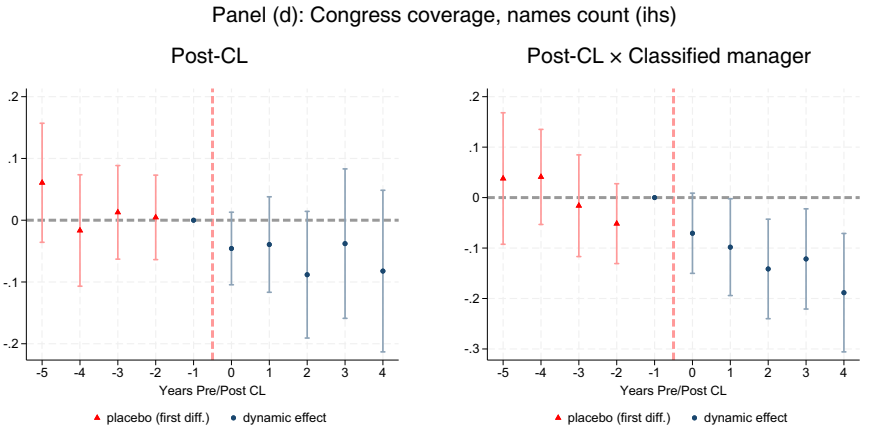

FIGURE 5
(Continued)

resulting estimates, presented in panel (c), suggest similar effect sizes for three out of the four main outcomes of interest.

Third, we consider specifications that model more precisely the dynamics of the effect of CL on newspaper outcomes. In panel (a) of Table A4, we replace the post-CL indicator with the number of years since CL entry, obtaining effects expressed in terms of an additional year of exposure to CL competition: 3.5% fewer jobs, 1% lower circulation per capita, 2.4% lower weight on political coverage, and 3.1% fewer articles covering congressional representatives or candidates. In panel (b) we instead split the post-CL indicator into an indicator for the 0–2 years post-CL entry and an indicator for more than 2 years post-CL entry. Consistent with the event-studies in Figure 5, the estimates comparing newspapers with a classified manager at baseline to ones without are up to twice as large in the long-run than in the short run.

Fourth, we address the concern that the presence of a classified manager may be correlated with newspaper characteristics such as size, which could confound the effects of CL entry. We address this issue in three ways. Our first approach is to consider a subsample of newspapers that is relatively homogeneous in size by excluding newspapers in the first and fourth quartile of the distribution of baseline readership. This narrows the range of baseline circulation per capita to between 0.1 and 0.2. Our main estimates are robust to this restriction (Table A5). Another approach is to explicitly allow for the impact of CL entry to vary by newspaper size by including as a control the interaction of baseline size and post-CL in addition to the interaction of classified manager and post-CL.[29] The estimates are robust to this inclusion (Table A6). Finally, we consider an alternative measure of classified reliance for which a mechanical correlation with newspaper size is unlikely: the average share of pages per issue devoted to classified ads. This measure is available from newspapers.com for a subsample of $\sim$ 250 newspapers.[30] In Table A7, we replace the indicator for the presence of a classified manager with an indicator equal to one if the newspapers' average share of classified pages exceeds the sample median of 0.3. Even

29. To avoid issues of mean reversion, we use circulation per capita as a measure of size when the dependent variable is number of jobs, and number of jobs when the dependent variable is circulation per capita.

30. We use the average number of pages of classifieds per issue from the newspapers.com archive as the numerator (see Supplementary Material, Appendix B.1.3 for details) and the average pages per issue reported by E&P as the denominator in constructing this share.

though we lose some precision in this significantly smaller sample, the estimates suggest the same direction and similar magnitudes as those in our baseline specification.

Fifth, in Table A8 we probe the robustness of our estimates to different sample and treatment definitions. While our main sample consists of all newspapers covered by E&P, including ones that exit or enter the market, in panel (a) we instead restrict the sample to a balanced panel of 1,300 newspapers that operate throughout the entire sample period of 1995–2010. In panel (b), we exclude the control group of newspapers that do not experience CL entry and only use variation in the timing of entry across newspapers that were eventually treated. In panel (c), we adopt a broader definition of CL and newspaper markets based on the geographic distribution of CL ads and newspapers' circulation (see Supplementary Material, Appendix Sections B.1.1 and B.1.2).

Finally, in Table A9 we show that the baseline estimates are robust to alternative levels of clustering of the standard errors—at the higher geographic level of state (panel a), or at the level of newspaper groups (panel b).

Spillovers. In Supplementary Material, Appendix Table A10(a) we examine the impact on papers not themselves exposed to CL, but located in a media market (DMA) in which other papers are exposed. To do so, we augment our baseline specification with the circulation-weighted leave-one-out share of newspapers in the same market that have experienced CL entry, that have a classified manager at baseline, and the interaction of the two. We find evidence of a small positive spillover in circulation, and no evidence of significant spillovers in terms of jobs or content.

Other outcomes. We also investigate whether CL affected newspapers in other ways. First, we test whether the entry of CL in an area affected the number of local newspapers. In the first two columns of Supplementary Material, Appendix Table A11, we estimate our two baseline specifications using as dependent variable the number of active newspapers headquartered in a county and year. The entry of CL does not appear to have any tangible effect on the number of newspapers, including in counties where newspapers had more to lose from increased competition for classified ads. Another possibility is that the revenue shock due to the entry of CL weakened newspapers' financial conditions, making them an easier target for acquisition. Using information from the E&P yearbooks on the group associated with each newspaper, in columns 3 and 4 of Supplementary Material, Appendix Table A11 we document that the entry of CL did not make local newspapers more likely to experience a change in ownership.[31] Finally, in Supplementary Material, Appendix Table A12, we explore whether, in response to the entry of CL, newspapers reduce the total number of pages, to cut costs, or increased subscription prices, to boost revenues. We find no decline in pages and a small negative effect on subscription prices with a magnitude of about 2%.[32]

31. A limitation of this analysis is that we use newspaper groups as listed by E&P and do not trace the ultimate cross-ownership across these groups.

32. The finding on subscription prices differs from the results in Seamans and Zhu (2014) who document a significant *increase* for newspapers more reliant on classified ads. This discrepancy could be due to the larger number of newspapers and years that our data cover (all daily newspapers for the entire period of 1995 to 2010), or to the fact that we use the pre-CL rather than contemporaneous presence of a classified manager as proxy for newspapers' reliance on classifieds. Another difference may be in the treatment of the subscription period the listed price refers to. We use E&P's yearly subscription price, which is available for most newspapers, and convert weekly or daily prices to the year level if the former is not available.

TABLE 4
Newspapers' jobs by type

Panel (a): Job categories

Num. jobs by category:	(1) Corporate/ General Mgmt.	(2) Advertising Mgmt.	(3) News Exec./ Editorial Mgmt.	(4) Other
Post-CL	0.025	0.085*	0.021	−0.026
	(0.079)	(0.049)	(0.348)	(0.123)
Post-CL × Classified Mgr.	−0.681***	−0.463***	−0.884**	−0.518***
	(0.133)	(0.080)	(0.444)	(0.191)
Baseline controls × Year FEs	Yes	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes	Yes
Newspaper FEs, Year FEs	Yes	Yes	Yes	Yes
Observations	20,327	20,327	20,327	20,207
Number of newspapers	1,436	1,436	1,436	1,436
R^2	0.80	0.80	0.89	0.82
Mean dependent variable	3.47	1.96	10.52	5.46

Panel (b): Editor job titles

Num. editors by job title:	(1) Politics	(2) Sports	(3) Entmnt.
Post-CL	0.036	0.037	−0.055
	(0.048)	(0.047)	(0.054)
Post-CL × Classified Mgr.	−0.158**	0.080	−0.081
	(0.063)	(0.085)	(0.093)
Baseline controls × Year FEs	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes
Newspaper FEs, Year FEs	Yes	Yes	Yes
Observations	19,071	19,071	19,071
Number of newspapers	1,413	1,413	1,413
R^2	0.74	0.59	0.76
Mean dependent variable	0.63	1.01	1.01

Notes: Dependent variables: number of jobs by E&P category (panel a) and number of editors by job-title derived topic (panel b). Regressions are conducted on an indicator for the availability of a local CL website and its interaction with an indicator for the presence of a classified manager at baseline. "Politics" is the number of editorial job titles that contain the keywords "politics/ political," "government," "Washington," "city" or "local." "Sports" is the number of editorial job titles that contain the keyword "sport/ sports." "Entertainment" is the number of editorial job titles that contain the keywords "entertainment," "lifestyle," "film," "music," "women," or "travel." Baseline controls include share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout, and share White/Black/Hispanic, all measured in the year 2000. Standard errors clustered by CL-area. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5.3. *Mechanism*

In this section, we present additional results on the impact of CL entry on newspaper's organization, readership and coverage and discuss our interpretation of the mechanisms that may be at play.

Jobs. An important question that the detailed data from E&P allow us to examine is what categories of workers were most affected by staff cuts. We can use the staff categories reported by E&P to determine whether a staff member holds a managerial or an editorial position, and, for editorial staff, we can in some cases identify the corresponding topical area (*e.g.* politics, sports, or entertainment).

In panel (a) of Table 4, we estimate our main specification separately for the number of jobs reported by E&P in each staff category. We differentiate between corporate/general management, advertising management, news executives and editorial management, and a residual category of other staff which captures mostly production and technology. The results indicate that CL entry leads to cuts in all types of positions, with magnitudes ranging from 23% for advertising management (which includes classified ad managers) to 7% for other staff. Importantly, staff cuts are not limited to management positions but also affect news executives and editors, whose number declines by about 9%. Since we find that CL entry does not reduce the total number of pages published per issue (columns 1 and 2 of Supplementary Material, Appendix Table A12), this result is suggestive of a net increase in editorial workload.

In panel (b), we focus on the category of news executives and editors and use their individual job titles to gauge the news topics they are likely to cover—an approach similar to Fan (2013) and George and Waldfogel (2006). We code keywords related to three common topics: politics, sports, and entertainment.[33] The results indicate a significant decline in the number of dedicated political editors after CL's entry, while we find no significant impact for sports or entertainment. In line with the conceptual framework described in Section 2, this result suggests that, when facing financial difficulties, newspapers affected by CL opted to cut staff especially in areas like politics, for which producing quality content is more costly (Hamilton, 2016).[34]

Content. In Table 3, we have shown that CL entry is associated with a decline in the coverage of politics, measured as either the topic weight on politics or the number of articles covering local representatives and candidates for office. In panel (a) of Table 5, we present the complete set of results from the CorEx model used to estimate the first measure, which produces a probability distribution over five topics. The topics can be labelled as "politics," "sports," "entertainment," "obituaries," and "crime" (Table B2 presents the most representative words by topic). Consistent with the analysis of editors' job titles, we find no significant effect of CL entry on the weights associated with sports or entertainment coverage. We also find no effect on crime coverage. On the other hand, we find a sizable (though only marginally significant) positive effect on the weight of obituaries, arguably the only type of classified advertising that local newspapers continue to monopolize.[35] Supplementary Material, Appendix Table A15 probes the robustness and heterogeneity of these results to a version of the CorEx model that is anchored to separate political topics into 4 subgroups: national, congressional, local and foreign (see Supplementary Material, Appendix Table B3 for representative words). The results suggest significant declines in the topic weights associated with all of the political categories, and again, no significant change in other topics besides a marginally significant increase in obituaries.

Supplementary Material, Appendix Table A17 confirms the finding that CL entry appears to have reduced coverage of all levels of politics. In this table, we focus on the count of articles covering either national or local- and state-level politicians. We measure articles that include the names of the President, the names of leadership of both parties in each chamber of Congress, or

33. We classify job titles containing the keyword "sport[s]" as sports related, ones containing the keywords "entertainment," "lifestyle," "women," "travel," "film," and "music" as entertainment related, and ones containing the keywords "politics / political," "government," "Washington," "city," and "local" as politics related.

34. An alternative interpretation is that readers' elasticity with respect to the volume of political news is lower than the elasticity with respect to the volume of soft news, and that the risk of an even larger decline in readership motivated newspapers to cut politics disproportionately. Though later in this section we show that readers interested in politics did react to a decline in political coverage, we cannot test how the elasticity with respect to this type of content compares to the elasticity with respect to soft news.

35. Obituaries are generally paid for and written by families of the deceased, similarly to classified ads.

TABLE 5
Newspapers' coverage

Panel (a): Topic model weights

Topic probability weight	(1) Politics	(2) Sports	(3) Entertainment	(4) Obituaries	(5) Crime
Post-CL	−0.000	0.007	−0.007	−0.003	−0.001
	(0.006)	(0.007)	(0.007)	(0.009)	(0.005)
Post-CL × Classified Mgr.	−0.024***	0.012	−0.010	0.022*	0.001
	(0.008)	(0.008)	(0.008)	(0.012)	(0.008)
Baseline controls × Year FEs	Yes	Yes	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes	Yes	Yes
Newspaper FEs, Year FEs	Yes	Yes	Yes	Yes	Yes
Observations	7,067	7,067	7,067	7,067	7,067
Number of newspapers	855	855	855	855	855
R^2	0.52	0.44	0.47	0.56	0.45
Mean dependent variable	0.29	0.21	0.22	0.15	0.14

Panel (b): Coverage of congressional representatives and candidates: By time in the election cycle Split by timing

Congress coverage names count (ihs)	(1) After primary election	(2) Prior to primary election
Post-CL	0.005	0.040
	(0.047)	(0.064)
Post-CL × Classified Mgr.	−0.001	−0.144**
	(0.047)	(0.072)
Baseline controls × Year FEs	Yes	Yes
Log population, num. ISPs	Yes	Yes
Newspaper FEs, Year FEs	Yes	Yes
Observations	7,375	7,375
Number of newspapers	878	878
R^2	0.92	0.85
Mean dependent variable	2.73	4.95

Notes: Regressions of content measures on an indicator for the availability of a local CL website and its interaction with an indicator for the presence of a classified manager at baseline. Dependent variables: CorEx model topic weights by newspaper and year (panel a), and (ihs-transformed) number of articles containing the name of a congressional representative or candidate by newspaper and year (panel b). In column (1) of panel (b), the dependent variable is limited to mentions occurring after the primary election date for the given two-year election cycle. In column (2), the dependent variable is limited to articles published before the primary election date. Baseline controls include share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout, and share White/Black/Hispanic, all measured in the year 2000. Panel (b) additionally controls for the (ihs-transformed) total number of articles in relevant sections recorded by Newsbank. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

titles commonly held by state and local elected officials. We find substantial negative effects on coverage of all of these groups.[36]

In panel (b) of Table 5, we break out the effect of CL entry on congressional coverage, exploring its heterogeneity over the electoral cycle. We consider the number of articles mentioning

36. We also find declines in the number of articles using one of the keywords indicative of investigative journalism from Hamilton (2016): see Supplementary Material, Appendix Table A19. The effect on indicators of investigative journalism is consistent with the idea that politics coverage declined due to reductions in investment in costly investigative activities.

the names of local representatives and candidates for congressional office published during the general election period (between the primary election date and the end of the election year) in column (1), and the number of articles published in all other times in column (2). Interestingly, the decline in congressional coverage is entirely concentrated outside of general election periods.[37] Supplementary Material, Appendix Table A18 shows a very similar result using an alternative measure of congressional coverage: the decline in the number of articles mentioning generic Congress-related keywords ("Congress," "Senate," or "U.S. House") is roughly twice as large among such articles which were published in a primary election period and which also contain the keywords "primary" or "nomination."

This temporal pattern implies we should expect to find larger effects on candidate selection in primaries than on general election turnout. The results on political outcomes presented in Section 5.4 will align with this expectation.

Readership. We verify and extend the analysis of newspaper readership using data from the NAES and GfK-MRI surveys, which include questions related to media consumption.[38] In column (1) of Table 6, panels (a) and (b), we report the results for self-reported readership of non-national newspapers from individual-level regressions controlling for respondents' age, race and education. We find significant reductions of newspaper readership in both surveys, with a magnitude of 3–6% relative to the respective sample mean—comparable to the estimates for circulation per capita.

The decline in readership is consistent with at least two explanations. First, it is possible that newspapers respond to the shock to classified ad revenues by increasing their subscription prices, which in turn would lead to lower demand. Yet, as mentioned above, we find no evidence of an increase in subscription prices for newspapers affected by CL.

A second explanation is that readers respond to the changes in content brought about by CL's entry. One possibility is that the change toward less coverage of politics that we document in Table 5 alienated readers interested in this type of content. Alternatively, the fall in circulation may be driven by readers who were primarily interested in classified ads which, after the entry of CL, became relatively less appealing. Though in both cases some readers would ultimately be less exposed to news and political content, understanding which of these scenarios is more plausible can shed light on which segments of the population were most affected by the entry of CL.

To understand this question, it is useful to first get a sense of how many readers were interested in these different newspaper sections at baseline. Information on this is available for a sample of 100,519 respondents from the 1999 to 2001 waves of the GfK-MRI survey. The distribution of readers' preferences, depicted in Supplementary Material, Appendix Figure B5, indicates that most readers (63%) report reading the "General News" section (which includes politics), with the Sports and Business sections also being popular (38% and 37%, respectively). The Classified section is not far behind, however, with 34% of respondents reporting reading this section. It is, therefore, possible that a reduction in the value of print classifieds may drive the decline in circulation.

To understand what types of readers drove the drop in readership of local papers following the entry of CL, we examine heterogeneity in the readership effect by propensity to read the

37. Supplementary Material, Appendix Table A16 presents additional heterogeneity tests distinguishing between coverage of incumbents versus challengers, and House versus Senate representatives and candidates. We find no significant differences along these dimensions.

38. While the two surveys contain only limited information on readership of specific newspapers, we are able to differentiate between respondents who report most frequently reading a national newspaper, *i.e.* the *New York Times*, *U.S. Today* or the *Wall Street Journal*, and ones who most frequently read other newspapers.

TABLE 6
Self-reported newspaper readership

	(1) Read newspaper dummy: Full sample	(2) News propensity ≥ median, Classif. propensity < median	(3) News propensity < median, Classif. propensity ≥ median
Panel (a): GfK-MRI			
Post-CL	−0.006	0.005	−0.004
	(0.007)	(0.011)	(0.010)
Post-CL × Classified Mgr.	−0.014*	−0.029**	−0.003
	(0.008)	(0.012)	(0.013)
Respondent controls	Yes	Yes	Yes
Baseline controls × Year FEs	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes
County FEs, Year FEs	Yes	Yes	Yes
Observations	248,460	83,494	81,703
Number of counties	781	756	762
R^2	0.09	0.10	0.05
Mean dependent variable	0.43	0.54	0.32
Panel (b): NAES			
Post-CL	0.016	0.029	0.005
	(0.012)	(0.032)	(0.019)
Post-CL × Classified Mgr.	−0.057**	−0.122***	−0.019
	(0.026)	(0.044)	(0.025)
Respondent controls	Yes	Yes	Yes
Baseline controls × Year FEs	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes
County FEs, Year FEs	Yes	Yes	Yes
Observations	106,348	30,133	34,831
Number of counties	1,192	1,087	1,167
R^2	0.05	0.09	0.06
Mean dependent variable	0.75	0.82	0.71

Notes: Dependent variable: Indicator for self-reported newspaper readership (excluding national newspapers) in the GfK-MRI (panel a) and NAES surveys (panel b). Individual-level regressions on an indicator for the availability of a local CL website in the county of the respondent, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. Respondent controls include sex, age, an indicator for college degree and race indicators. Baseline controls include share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout and share White/Black/Hispanic, all measured in the year 2000. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

classified versus general news sections. Using the 1999–2001 waves of the GfK-MRI survey, we estimate an elastic-net penalized regression model to identify the individual characteristics that are most predictive of reading the general news and the classifieds sections, respectively.[39] Based on the model estimated in the 1999–2001 data, we then project two propensity scores for respondents in the following years of GfK-MRI, as well as respondents in the NAES survey. This

39. Regularized regression is useful here because GfK-MRI records more than 600 demographic features of respondents. The characteristics most strongly associated with general news reading are being white, having a post-graduate degree, having an income in the 75 K–150 K range, being retired, being married 25 years or more, and being aged 45–49. The characteristics most associated with reading classifieds are: being unemployed, living in a small to moderate sized county, having a high school diploma only or "some college," being engaged (to be married), and being 25–29 years old.

procedure allows us to assign to each respondent a probability for reading general news and one for reading classifieds. Projecting based on pre-CL data allows us to focus attention on differential changes among demographic types who would have been likely to read either classifieds or political news prior to CL entry, without the confound of the post-CL changes to newspapers' product. The projected propensity scores have fairly strong negative correlation, with correlation coefficient of -0.4 in the GfK-MRI dataset and -0.2 in the NAES dataset, indicating that the groups that tend to read each section are relatively distinct. In particular, younger, less educated, less wealthy, and unemployed respondents report high interest in classifieds, whereas older, wealthier, more educated, and retired respondents report high interest in general news.

With these propensity scores at hand, we re-estimate the individual-level readership regressions separately for two groups of respondents: (i) those with above-median probability of reading classifieds and below-median probability of reading general news, and (ii) those with below-median probability of reading classifieds and above-median probability of reading general news. The results, reported in columns (2) and (3) of Table 6, indicate that the decline in readership after the entry of CL is entirely driven by individuals with high news propensity and low classified propensity. Hence, though newspapers which at baseline offered the most classifieds were most affected by the CL shock, we find that readers of those papers interested in news rather than classifieds are the ones decreasing their readership.[40]

Taken together, these results support the view that the main driver of circulation declines was the indirect shift in news content induced by newspapers' revenue loss, rather than the direct effect of the obsolescence or disappearance of print classified ads.

Substitution to other news sources. In Supplementary Material, Appendix A2, we examine the possibility that readers may have substituted into other forms of news media. Overall, we find little evidence of such a substitution. Supplementary Material, Appendix Tables A13 and A14 show minimal effects of CL entry on visits to major news websites in the Comscore data or on self-reported consumption of other media. For example, looking at the number of visits to the top 100 news websites classified by Comscore, we can rule out an increase in response to CL entry larger than 2.5%.

5.4. *Political outcomes*

In the previous sections, we documented that newspapers affected by the entry of CL reported less about politics, in general, and candidates for congressional office, in particular. We also found that individuals in areas affected by CL reduced their readership of local newspapers. In this section, we examine the electoral implications of these changes. Given existing evidence on the relationship between exposure to political information and citizens' political decisions, it is plausible that changes in news content and newspaper readership may have ramifications for downstream political outcomes.

We focus on three groups of congressional election outcomes, motivated by the prior literature: (i) voter turnout (Gentzkow *et al.*, 2011), (ii) voters' propensity to rely on national partisan cues when voting for local candidates, measured by the incidence of split-ticket voting (Darr *et al.*, 2018; Trussler, 2021; Moskowitz, 2021), and (iii) the entry and electoral performance of ideologically extreme candidates (Hall, 2015; Hall and Lim, 2018; Autor *et al.*, 2020).

40. It is also important to note that the profile of readers who reduce their readership in response to the CL shock stands in contrast with the profile of readers who may be more likely to migrate online, *i.e.* the young and the less-educated (Gavazza *et al.*, 2019).

Turnout. Since electoral data for House and Senate elections are available at the county level, we estimate the county-level specification discussed in Section 4. We assume that newspapers affect voters' behaviour in the county in which they are based, which limits the sample to the 1,234 counties in which, according to the E&P data, at least one newspaper is headquartered. We stack House and Senate elections controlling for office × year fixed effects and, given the considerable differences in the number of voters across counties, weight observations by the county's total voting-age population.

Columns (1) and (2) of Table 7(a) present the results. We find small and insignificant estimates for the impact of CL on turnout in congressional elections. We can rule out a negative effect of magnitude larger than 1.3 percentage points.[41] One explanation of this null result is that, as we show in Table 6, the reduction in local newspaper readership primarily affected individuals with characteristics associated with high electoral turnout (*i.e.* older, higher income, higher education). These readers were, for the most part, not marginal voters, and even if exposed to less local political news were unlikely to drop into abstention. This contrasts with existing work (Gentzkow, 2006; Ellingsen and Hernæs, 2018; Gavazza *et al*., 2019) which finds effects on turnout of substitution from newspapers into television or Internet; there the marginal readers tend to be younger, less educated, and lower income.

Split-ticket voting. Though changes in the information environment induced by CL may not impact *whether* people vote, they may affect *how* they vote. In particular, less exposure to coverage of candidates for office could reduce voters' ability to evaluate their specific platforms and valence attributes, and hence increase reliance on party labels.

Following the literature, we measure voters' tendency to deviate from party labels, or split-ticket voting, as the absolute difference in the Republican vote share in presidential versus concurrent congressional elections. We measure vote shares at the county level and we follow the exact same specification as for turnout, again stacking House (versus presidential) and Senate (versus presidential) elections.

The results, presented in columns (3) and (4) of Table 7(a) indicate that, following the entry of CL, voters become significantly less likely to split their vote between candidates of different parties. The effect is again driven by areas where newspapers were most vulnerable to CL's competition. The magnitude of the decline in split-ticket voting is about 13% of the sample mean.[42]

In Supplementary Material, Appendix Figure A1, we present the event-study corresponding to column (4) of Table 7. Since split-ticket voting is only defined for elections held in presidential years, we consider three election cycles before the entry of CL and two after. We define as "treated" counties that experience CL entry and in which more than 50% of newspapers (circulation-weighted) had a classified manager in the year 2000. We find no evidence of pre-existing trends, and a significant decline in split-ticket voting after the entry of CL in counties where newspapers were more reliant on classified ads.

Support for extreme candidates. Finally, we study whether the entry of CL favoured the emergence and success of ideologically extreme candidates. The hypothesis is that a coarser information environment makes the entry of more extreme candidates more likely and improves their electoral prospects by making it harder for voters to acquire information about candidates'

41. We also find null effects for turnout in presidential elections.

42. This result relates to similar findings by Darr *et al*. (2018) regarding the impact of newspaper closures on split-ticket voting. Our results indicate that closures are not a necessary condition, and that the impoverishment of local newspapers can produce similar consequences.

TABLE 7
Political effects

Panel (a): Turnout in Congressional elections and split-ticket voting

	Turnout House/Senate		Split-ticket vote	
	(1)	(2)	(3)	(4)
Post-CL	−0.003	0.000	−0.014***	−0.005
	(0.003)	(0.004)	(0.005)	(0.006)
Post-CL × Classified Mgr.		−0.005		−0.013*
		(0.004)		(0.007)
Baseline controls × Year FEs	Yes	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes	Yes
County FEs, Year-Office FEs	Yes	Yes	Yes	Yes
Observations	15,938	15,938	7,900	7,900
Number of counties	1,201	1,201	1,201	1,201
R^2	0.91	0.91	0.37	0.37
Mean dependent variable	0.45	0.45	0.11	0.11

Panel (b): Entry and performance of extreme candidates in House elections

	Extremist in primary		Extremist in general		Extremist wins general	
	(1)	(2)	(3)	(4)	(5)	(6)
Post-CL	0.024	−0.032	0.042*	0.003	0.003	−0.026
	(0.023)	(0.027)	(0.024)	(0.029)	(0.019)	(0.018)
Post-CL × Classified Mgr.		0.095***		0.066**		0.049**
		(0.030)		(0.030)		(0.023)
Baseline controls × Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Log population, num. ISPs	Yes	Yes	Yes	Yes	Yes	Yes
County FEs, Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
District FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,700	12,700	12,998	12,998	12,868	12,868
Number of counties	1,201	1,201	1,201	1,201	1,191	1,191
Number of districts	439	439	439	439	437	437
R^2	0.56	0.56	0.57	0.57	0.79	0.79
Mean dependent variable	0.79	0.79	0.75	0.75	0.37	0.37

Notes: Regressions of electoral outcomes on an indicator for the availability of a local CL website and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. Panel (a): The level of observation is county × election. Observations are weighted by voting-age population. Standard errors clustered by CL-area. Panel (b): The level of observation is electoral district - county cell × election. Observations are weighted by the share of the voting-population in the district-county cell relative to the district. Standard errors clustered by district. Baseline controls include share of urban population, share college educated, rental share of housing, log income per capita, median age, turnout and share White/Black/Hispanic, all measured in the year 2000. Standard errors clustered by CL-area. Significance levels: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$.

ideological positioning.[43] As a concrete example, Skocpol and Williamson's (2016) study of the Tea Party movement argues that coverage of the movement by national media outlets focused on national polling data, promoting the "wrong-headed notion that the Tea Party appealed to centrist

43. This is in line with correlational findings by Hall and Lim (2018) who document that the advantage of extreme candidates in Congressional primaries is concentrated in areas with low news coverage, and consistent with the theoretical model of Matějka and Tabellini (2021), which predicts that increases in the cost of acquiring political information will induce candidates to cater more to relatively extreme segments of the population, because these extreme segments are willing to pay higher costs of acquiring information than are moderates.

independents (pp. 147)." The thinness of on-the-ground media coverage of Congressional primaries—in which Tea Party-affiliated candidates consistently challenged Republican incumbents from the right—allowed the movement to maintain an electorally beneficial ambiguity about its ideological aims.

To test this hypothesis, we use a measure of candidate ideology (CFScore) available for all candidates in House elections from Bonica (2016) (see Section 3.2 for details of the measure).[44] We examine three outcomes: (i) the probability that an ideologically extreme candidate runs in a primary election; (ii) the probability that an ideologically extreme candidate *wins* a primary election and thus runs in the general election, and (iii) the probability that ideologically extreme candidates win the general election.

Since these outcomes are defined at the electoral district level, we use the version of our baseline specifications defined at the level of county × district cells (see Section 4). All regressions include district × redistricting regime fixed effects to absorb the effect of changes in district boundaries, so that we exploit variation over time within a constant district boundary. It is important to stress that the geographic level of electoral districts, at which the outcomes in this analysis are observed, is much coarser than the level of the treatment (*i.e.* county). To the extent that the assignment of the local treatment to an aggregated outcome introduces measurement error, it should bias our coefficient toward zero.

The results are presented in Table 7(b). Columns (1) and (2) indicate that following the entry of CL and in districts where newspapers were more reliant on classified ads, the probability that a candidate with an extreme CFScore runs in a primary election increases significantly by about 9.5 percentage points, or 12% of the mean. We also find a significant 8.8% increase in the likelihood that such candidates win the primary and run in the general election (columns (3) and (4)), as well a significant 13% increase in the likelihood that an ideologically extreme candidate wins the general election (columns (5) and (6)).

In Supplementary Material, Appendix Figure A2, we present event-studies corresponding to each of the three outcomes. Since in this case we consider all House elections (held in both presidential and mid-term years), we can estimate dynamic effects with a larger number of leads and lags. The results show no significant pre-trends in the four election cycles prior to CL entry and a clear increase in the probability that extremists participate in and win in primaries in the three election cycles after CL entry, in districts with classified-reliant newspapers.[45] The event-study for the probability of extremists winning general elections is on the other hand noisier and in this case does not match the two-way fixed-effects estimate in Table 7; we therefore interpret this last result with caution.

Taken together, these findings are consistent with the notion that the transformations in the media landscape triggered by the entry of CL had downstream effects on electoral politics. The coarser information on candidates for office may have reduced voters' ability to screen them and thereby increased their equilibrium ideological divergence.

6. CONCLUSION

Hamilton (2004) lays out the basic economics of the news-gathering business: high fixed costs—in the form of reporting staff who must develop expertise in their subjects and form long-term relationships with their sources—combined with a non-excludable product lead

44. We focus on the House because we can define outcomes at the finer level of electoral district, rather than at the coarse level of state as would be the case for Senate elections.

45. The relatively rapid effect of CL on candidate entry is in line with the short timelines and minimal formal barriers to entry of candidate-centric U.S. House campaigns; according to Herrnson (2008), strategic candidates focus on local circumstances and "respond to opportunities that arise in specific election years (pp. 42)" in making entry decisions.

generally to under-provision of news relative to the social optimum. Counteracting this unhappy equilibrium to some degree are reporters' professional norms, which value the production of "hard news" and investigative journalism over cheaper-to-produce and sometimes more popular "soft news."

For a time in the 20th century, local monopoly papers were able to extract sizable profits from the advertising business. Reporters employed by those papers captured some of these rents in the form of resources dedicated to reporting of local political news and other "hard" topics valued by journalists themselves (rather than readers or advertisers). The growth of advertising profits, in fact, can be directly tied to the emergence of the ideal of an independent press staffed by professional journalists, in contrast to the 19th-century norm of newspapers operated as propaganda organs of local party organizations (Petrova, 2011).

The emergence of competition in the advertising business from new internet-based entrants in the early 2000s upset this tenuous balance, eliminating the economic profits which had supported investments in money-losing but high-prestige reporting. We show that the entry of one particularly important competitor, the classified advertising platform CL, had severe impacts on newspapers' staffing levels and production of political news coverage.

The CL effect is not simply a consequence of changes to the demand for news induced by internet availability; rather, it appears to operate by reducing newspapers' ability to invest in local reporting resources. Papers that were especially reliant on classified advertising in the pre-CL period saw much larger changes on these dimensions than comparably internet-exposed but less classified-dependent papers. The loss of advertising revenues at these papers seems to have reduced political coverage and in particular coverage of local representatives, an area with large positive externalities but also large private costs for newspaper operators.

Consistent with existing work on media effects on political outcomes, we find that there were measurable social consequences of this change in the production of news content. Voters in areas served by papers affected most by the CL shock saw their Congressional elections become more nationalized, which we interpret as a consequence of thinner information about the local incumbent's behaviour. Changes in the media landscape may thus be an important driver of the overall trend towards nationalization of elections in the U.S. (Hopkins, 2018).

The change in voters' information about candidates also had consequences for ideological polarization. We show that the reduction of representative-specific information led to greater entry and better electoral performance by relatively extreme candidates at the expense of their more moderate peers. This change provides evidence of newspapers' role in providing information about candidates' ideological and issue positioning. The contraction in coverage generated by CL's entry diminished voters' ability to distinguish between moderate and extreme candidates, reducing the electoral penalty to taking positions "out-of-step" with district preferences (Canes-Wrone *et al*., 2002). The coarseness of voters' information environment thus provides a plausible explanation for the documented failure of Downsian convergence in candidate positioning in the U.S. House (Fowler and Hall, 2016) and for the relatively weak relationship between candidate positions and vote intentions (Tausanovitch and Warshaw, 2018).

Our results have implications for our understanding of the link between advertising market structure and the market for news. They highlight the fragility of compensating the production of a public good—politically relevant information—with proceeds from bundled advertising. Technological innovation that unbundles the two products, as CL did for classified advertising, can have spillover effects on the news market, with significant and lasting consequences for the quality of representation and political polarization.

In this paper, we focus on the implications of the transformations triggered by CL's entry for congressional politics, relying on available measures of election outcomes and candidates' ideology. Yet, it is likely that similar mechanisms apply, potentially to an even greater extent, to

political issues at the more local level for which reporting outside of local newspapers is even more scarce. This question presents one possible avenue for future research.

Acknowledgments. We thank the editor, four anonymous referees and numerous seminar and conference participants for helpful comments. David Ampudia Vicente, Marco Lo Faso, Lorena Mita, and Aaron Schade provided excellent research assistance. We also thank Irena Djourelova, Anna Griñó, Anastasiya Nebolsina, Adriana Oliveres and Silvia Terzulli for their help with digitizing the Editor & Publisher data. Ruben Durante acknowledges financial support from the European Union's Horizon 2020 research and innovation program [Grant 759885].

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Data Availability

The data and code underlying this research is available on Zenodo at https://doi.org/10.5281/zenodo.10120113.

REFERENCES

ANGELUCCI, C. and CAGÉ, J. (2019), "Newspapers in Times of Low Advertising Revenues", *American Economic Journal: Microeconomics*, **11**, 319–3643.
ANGELUCCI, C., CAGE, J. and SINKINSON, M. (2024), "Media Competition and News Diets", *American Economic Journal: Microeconomics*, forthcoming.64666
AUTOR, D., DORN, D., HANSON, G., *et al.* (2020), "Importing Political Polarization? The Electoral Consequences of Rising Trade Exposure", *American Economic Review*, **110**, 3139–3183.
BESLEY, T. and BURGESS, R. (2002), The Political Economy of Government Responsiveness: Evidence from India", *Quarterly Journal of Economics*, **117**, 1415–1451.
BHULLER, M., HAVNES, T., MCCAULEY, J., *et al.* (2024), "How the Internet Changed the Market for Print Media", *American Economic Journal: Applied Economics*, **16**, 318–358.
BONICA, A. (2016), *Database on Ideology, Money in Politics, and Elections* (Stanford, CA: Stanford University Libraries). https://data.stanford.edu/dime.
CALVERT, R. L. (1985), "Robustness of the Multidimensional Voting Model: Candidate Motivations, Uncertainty, and Convergence", *American Journal of Political Science*, **29**, 69–95.
CAMPANTE, F., DURANTE, R. and SOBBRIO, F. (2018), "Politics 2.0: The Multifaceted Effect of Broadband Internet on Political Participation", *Journal of the European Economic Association*, **16**, 1094–1136.
CANES-WRONE, B., BRADY, D. W. and COGAN, J. F. (2002), "Out of Step, Out of Office: Electoral Accountability and House Members' Voting", *American Political Science Review*, **96**, 127–140.
CHAN, J. and GHOSE, A. (2014), "Internet's Dirty Secret: Assessing the Impact of Online Intermediaries on HIV Transmission", *MIS Quarterly*, **38**, 955–976.
Comscore (2010), "WRDS ComScore, Wharton Research Data Services". https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/comscore/ (accessed April, 2023).
CUNNINGHAM, S., DEANGELO, G. and TRIPP, J. (2019), "Craigslist's Effect on Violence Against Women" (Working Paper).
DARR, J. P., HITT, M. P. and DUNAWAY, J. L. (2018), "Newspaper Closures Polarize Voting Behavior", *Journal of Communication*, **68**, 1007–1028.
DE CHAISEMARTIN, C. and D'HAULTFOEUILLE, X. (2020), "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects", *American Economic Review*, **110**, 2964–2996.
ELLINGSEN, S. and HERNÆS, Ø. (2018), "The Impact of Commercial Television on Turnout and Public Policy: Evidence from Norwegian Local Politics", *Journal of Public Economics*, **159**, 1–15.
E&P (2010), *Editor & Publisher International Year Book, 1995–2010*.
FALCK, O., GOLD, R. and HEBLICH, S. (2014), "Elections: Voting Behavior and the Internet", *American Economic Review*, **104**, 2238–2265.
FAN, Y. (2013), "Ownership Consolidation and Product Characteristics: A Study of the US Daily Newspaper Market", *American Economic Review*, **103**, 1598–1628.
FCC (2008), "Federal Communications Commission Form 477 Additional Data, 1999-2008". https://www.fcc.gov/general/fcc-form-477-additional-data (accessed April, 2023).
FCC (2016), "Federal Communications Commission Report FCC 16-107" (Tech. rep).
FOWLER, E. F. (2020), "Strategy over Substance and National in Focus?: Local Television Coverage of Politics and Policy in the United States", in *The Routledge Companion to Local Media and Journalism* (Routledge) 8.
FOWLER, A. and HALL, A. (2016), ""The Elusive Quest for Convergence" Available at SSRN 2262600.
GALLAGHER, R. J., REING, K., KALE, D., *et al.* (2017), "Anchored Correlation Explanation: Topic Modeling with Minimal Domain Knowledge", *Transactions of the Association for Computational Linguistics*, **5**, 529–542.
GAVAZZA, A., NARDOTTO, M. and VALLETTI, T. (2019), "Internet and Politics: Evidence from U.K. Local Elections and Local Government Policies", *Review of Economic Studies*, **86**, 2092–2135.
GENTZKOW, M. (2006), "Television and Voter Turnout", *Quarterly Journal of Economics*, **121**, 931–972.

GENTZKOW, M. and SHAPIRO, J. M. (2010), "What Drives Media Slant? Evidence from U.S. Daily Newsapers", *Econometrica*, **78**, 35–71.
GENTZKOW, M., SHAPIRO, J. M. and SINKINSON, M. (2011), "The Effect of Newspaper Entry and Exit on Electoral Politics", *American Economic Review*, **101**, 2980–3018.
GEORGE, L. M. and WALDFOGEL, J. (2006), "The New York Times and the Market for Local Newspapers", *American Economic Review*, **96**, 435–447.
GfK-MRI (2010), "MRI-Simmons: Survey of the American Consumer 1999-2010", GfK - Mediamark Research & Intelligence. https://www.mrisimmons.com/our-data/national-studies/survey-american-consumer/.
HALL, A. B. (2015), "What Happens When Extremists Win Primaries?", *American Political Science Review*, **109**, 18–42.
HALL, A. B. and LIM, C. (2018), "Ideology and News Content in Contested US House Primaries" (Tech. rep., Working Paper.) http://andrewbenjaminhall.com/news.pdf.
HALL, A. B. and THOMPSON, D. M. (2018), "Who Punishes Extremist Nominees? Candidate Ideology and Turning Out the Base in US Elections", *American Political Science Review*, **112**, 509–524.
HAMILTON, J. (2004), *All the News That's Fit to Sell: How the Market Transforms Information into News* (Princeton, NJ: Princeton University Press).
HAMILTON, J. T. (2016), *Democracy's Detectives: The Economics of Investigative Journalism* (Cambridge: Harvard University Press).
HERRNSON, P. S. (2008), *Congressional Elections: Campaigning at Home and in Washington* (5th ed) (Washington, DC: CQ Press).
HOPKINS, D. J. (2018), *The Increasingly United States: How and Why American Political Behavior Nationalized* (Chicago, IL: University of Chicago Press).
KROFT, K. and POPE, D. G. (2014), "Does Online Search Crowd Out Traditional Search? Evidence from Craigslist", *Journal Labor Economics*, **32**, 259–303.
LARCINESE, V. and MINER, L. (2018), "Was Obama Elected by the Internet? Broadband Diffusion and Voters Behavior in US Presidential Elections" (Working Paper No. 6882, CESifo).
LEIP, D. (2020), "Atlas of American Elections, 1995–2010". https://uselectionatlas.org/ (accessed April, 2023).
LELKES, Y., SOOD, G. and IYENGAR, S. (2017), "The Hostile Audience: The Effect of Access to Broadband Internet on Partisan Affect", *American Journal of Political Science*, **61**, 5–20.
LOS SANTOS, B. D., HORTAÇSU, A. and WILDENBEEST, M. R. (2012), "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior", *American Economic Review*, **102**, 2955–2980.
MAHONE, J., WANG, Q., NAPOLI, P., *et al.* (2019), "Who's Producing Local Journalism? Assessing Journalistic Output Across Different Outlet Types" (Tech. rep).
MATĚJKA, F. and TABELLINI, G. (2021), "Electoral Competition with Rationally Inattentive Voters", *Journal of the European Economic Association*, **19**, 1899–1935.
MCCHESNEY, R. W. and NICHOLS, J. (2011), *The Death and Life of American Journalism: The Media Revolution That Will Begin the World Again* (New York, NY: Bold Type Books).
MOSKOWITZ, D. J. (2021), "Local News, Information, and the Nationalization of US Elections", *American Political Science Review*, **115**, 114–129.
NAES (2008), "Annenberg Public Policy Center: National Annenberg Election Survey 2000, 2004, 2008", University of Pennsylvania. https://www.annenbergpublicpolicycenter.org/2008-naes-telephone-and-online-data-sets/ (accessed April, 2023).
NewsBank (2010), "Access World News Database". https://www.newsbank.com/libraries/military/solutions/access-world-news.
PETERSON, E. (2021), "Paper Cuts: How Reporting Resources Affect Political News Coverage", *American Journal of Political Science*, **65**, 443–459.
PETROVA, M. (2011), "Newspapers and Parties: How Advertising Revenues Created an Independent Press", *American Political Science Review*, **105**, 790–808.
ROTH, J., SANT'ANNA, P. H. C., BILINSKI, A., *et al.* (2023), "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature", *Journal of Econometrics*, **235**, 2218–2244.
SEAMANS, R. and ZHU, F. (2014), "Responses to Entry in Multi-Sided Markets: The Impact of Craigslist on Local Newspapers", *Management Science*, **60**, 476–493.
SKOCPOL, T. and WILLIAMSON, V. (2016), *The Tea Party and the Remaking of Republican Conservatism* (Oxford, UK: Oxford University Press).
SNYDER, J. M. and STRÖMBERG, D. (2010), "Press Coverage and Political Accountability", *Journal of Political Economy*, **118**, 355–408.
SRDS (2006), *Standard Rate and Data Service: Newspaper Advertising Source, 1995–2006*.
STARKMAN, D. (2014), *The Watchdog That Didn't Bark: The Financial Crisis and the Disappearance of Investigative Journalism* (New York, NY: Columbia University Press).
STRÖMBERG, D. (2001), "Mass Media and Public Policy", *European Economic Review*, **45**, 652–663.
TAUSANOVITCH, C. and WARSHAW, C. (2018), "Does the Ideological Proximity Between Candidates and Voters Affect Voting in U.S. House Elections?", *Political Behavior*, **40**, 223–245.
TRUSSLER, MARC. (2021), "Get Information or Get in Formation: The Effects of High-Information Environments on Legislative Elections", *British Journal of Political Science*, **51**, 1529–1549.
WITTMAN, D. (1983), "Candidate Motivation: A Synthesis of Alternative Theories", *The American Political Science Review*, **77**, 142–157.