Essays on the Political Economy of Mass Media

Milena Djourelova

TESI DOCTORAL UPF / Year 2021

THESIS SUPERVISOR
Ruben Durante
Department Department dâEconomia i Empresa

Acknowledgements

During my PhD I have been lucky to meet mentors and peers who have inspired me with their ideas and supported me in developing mine. My advisor and coauthor Ruben Durante has been a source of encouragement and I am grateful for his enthusiasm for my work. Maria Petrova and Ruben Enikolopov introduced me to political economy in their classes and reading groups and have had a great influence on my research interests. Brian Knight and David Strömberg welcomed me in Brown and in Stockholm respectively and provided invaluable comments on the chapters comprising this thesis. I also thank Greg Martin for coauthoring the final chapter, and the applied economics group at UPF for their always constructive feedback.

My colleagues and friends including Ines, Eva, Adil, Angelo, Sebastian, Josep, Ilja, Shohei, Rahil, Juan, Karolis, Vasi, Pilar, Lauri (& Fifi) have shared the many fun moments of PhD life in Barcelona and of writing this thesis, as well as the less fun ones. I thank them for their friendship.

Finally, I thank my sister Irena for her research assistance and media expertise, and my whole family in Bulgaria – with special mention of my grandparents Milka and Ivan and my parents Marinela and Nikolay – for their patience and support.

Abstract

This dissertation studies three questions on the political economy of mass media. The first chapter examines the effects of media slant on public policy views, exploiting an experiment created by a ban on the politically charged term "illegal immigrant" by the Associated Press newswire. The second chapter studies the strategic timing of controversial policies with respect to the media cycle, looking at the case US presidential executive orders. The third chapter investigates the impact of online competition on the operation of traditional media and the downstream political implications of these transformations.

Resum

Aquesta tesi estudia tres qüestions sobre l'economia política dels mitjans de comunicació. El primer capítol estudia els efectes del biaix dels mitjans de comunicació sobre les opinions de les polítiques públiques, aprofitant un experiment creat per la prohibició del terme políticament carregat d'immigrant illegal per la revista Associated Press. El segon capítol estudia el moment estratègic de polèmiques controvertides pel que fa al cicle mediàtic, examinant el cas de les ordres de l'executiu presidencial dels EUA. El tercer capítol investiga l'impacte de la competència dels mitjans de comunicació en línia sobre el funcionament dels mitjans tradicionals i les implicacions polítiques posteriors d'aquestes transformacions.

Preface

In this dissertation I study three questions related the influence of mass media on public policy views, the incentives for political actors created by this influence, and the effects of online competition on the operation and editorial choices of traditional media.

Chapter 1 studies the effects slanted language in mass media on public policy views. I examine this question in the context of the US debate on immigration, exploiting an abrupt ban on the politically charged term "illegal immigrant" in the dispatches distributed to media outlets by the Associated Press (AP) news wire. I use the timing of the ban combined with variation across media outlets in the extend to which they rely on AP-copy (AP-intensity), to study first – the effects of the ban on the media coverage of immigration, and second – its impact on public views on immigration policy. I find that one standard deviation higher AP-intensity leads to a 10 to 14% decline in use of the term "illegal immigrant", and a 0.7 percentage points decline is support for restrictive immigration and border security policies after the ban. The effect is driven by readers who are less politically engaged, and does not transfer to views on issues other than immigration. These findings imply that parties may use political rhetoric strategically to pull less engaged constituencies towards their platform. This chapter is an extended version of BSGE working paper Djourelova (2020).

Chapter 2 studies the strategic behavior of politicians when it comes to the release of potentially controversial policies. We examine this question by analyzing the timing of executive orders (EOs) signed by U.S. presidents over the past four decades. We find robust evidence that EOs are more likely to be signed on the eve of days when the news are dominated by other important stories that can crowd out coverage of EOs. This relationship only holds in periods of divided government when unilateral presidential actions are more likely to be criticized by Congress. The effect is driven by EOs that are more likely to make the news and to attract negative publicity, particularly those on topics on which president and Congress disagree. Finally, the timing of EOs appears to be related to predictable news but not to unpredictable ones, which suggests it results from a deliberate and forward-looking PR strategy. This chapter is a version of CEPR discususion paper Djourelova and Durante (2020), coauthored with Ruben Durante, and is conditionally accepted for publication in the *Amercian Journal of Political Science*.

Chapter 3 studies how competition from online platforms affects the organization, performance, and editorial choices of newspapers. We examine this question exploiting the staggered introduction of Craigslist - the world's largest online platform for classified advertising - across US 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 Craigslist, local papers experienced a significant decline in the number of newsroom and management staff. Cuts in editorial staff disproportionately affected reporters covering politics. These organizational changes led to a reduction in news coverage of politics, and resulted in a decline in newspaper readership which was not compensated by increased news consumption online or on other media. Finally, we find evidence that reduced news coverage of politics was associated with stronger reliance on party cues among voters, and benefited ideologically extreme candidates. This chapter is a version of CEPR discussion paper Djourelova et al. (2021), coauthored with Ruben Durante and Gregory Martin.

Contents

Chapter 1

MEDIA PERSUASION THROUGH SLANTED LANGUAGE

1.1 Introduction

Political actors choose carefully the words they put out in the media. In the US, Republicans and Democrats often use strikingly different language to describe the exact same issue, in an attempt to promote views favorable to their platform. Republican politicians and right-leaning media speak about the "China virus", the "death tax", and "illegal immigrants", while Democrats and left-leaning media refer to the same issues as "Covid-19", the "estate tax", and "undocumented immigrants".[1] If language has an impact on how the issue is perceived by the public, such partisan tactics could lead to polarized perceptions of the same factual reality (Alesina et al. 2020).

Yet, evidence on the persuasiveness of slanted language, i.e. on whether it can indeed sway readers in the intended direction, is lacking. From an empirical standpoint, assessing the causal impact of language is challenging for at least two reasons. First, votes-maximizing politicians and profit-maximizing media have an incentive to choose their slant to appeal to their audience's preference for like-minded content (Gentzkow and Shapiro 2006), making it difficult to disentangle cause and effect. Second, slanted language can be accompanied by other politically motivated choices, ranging from selective coverage of certain issues (Puglisi and Snyder 2011), to outright endorsement of policies or candidates (Chiang and Knight 2011), which are likely to independently affect the views of the audience.

To overcome these challenges, I propose a supply-side source of variation in media slant. I take advantage of the fact that many US media outlets source some of their content from the Associated Press (AP) – a newswire agency that gathers and distributes news to subscribing outlets. Since AP distributes a single news feed to all subscribers, their aim is to produce neutral coverage that appeals to outlets from all sides of the political spectrum (Fenby 1986). This philosophy has led to extremely strict and rigid guidelines for the use of politically sensitive language.

[1] The implicit policy positions behind these phrases are easy to recognize. "The China virus" is presumably an attempt to shift responsibility for the crisis to China (`https://edition.cnn.com/2020/03/20/politics/donald-trump-china-virus-coronavirus/index.html`). "Death tax" highlights the alleged unfairness of taxing the deceased, while "estate tax" draws attention to the wealth of the people it applies to (`https://www.businessinsider.com/death-tax-or-estate-tax-2017-10?r=US&IR=T`). "Illegal immigrants" underscores the transgression of crossing the border, while "undocumented immigrants" presents the issue of legal status as a formality (`https://www.al.com/news/2018/07/illegal_vs_undocumented_the_he.html`).

I exploit an abrupt reversal in AP's guidelines on the use of a specific politically charged term – "illegal immigrant". In April 2013, after years of resisting requests to revise its guidelines on the language on immigration, AP switched from officially *recommending* the term "illegal immigrant" to refer to people living in the US without legal authorization, to *banning* its use in AP wire dispatches.

The ban happened at a time when the issue of immigration, and the language used to talk about it, was extremely politicized. Figure 1 illustrates the partisan divide in use of the term "illegal immigrant" in political speech and in the media.[2] In Congress, Republican representatives use the label "illegal" about 50% of the time they mention the term "immigrant", while this frequency is less than 5% among Democrats. Similarly, the term appears twice as frequently in the right-leaning Fox News and Washington Times, compared to the left-leaning MSNBC and Washington Post.

Figure 1: Ideological charge of the term "illegal Immigrant"

(a) Congressional Speech

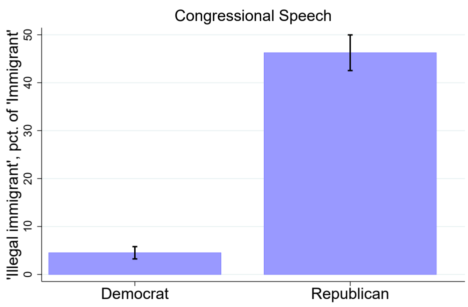

(b) TV

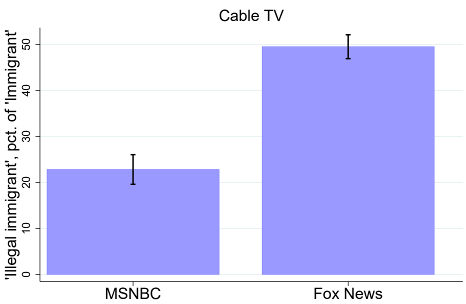

(c) Newspapers

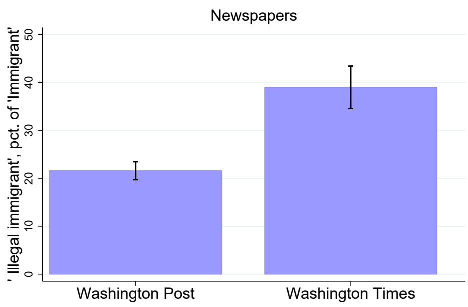

Notes: Frequency of mentions of "illegal immigrant" relative to "immigrant" in congressional speech, in cable TV (comparing MSNBC and Fox News) and in newspapers (comparing the Washington Post and the Washington Times) in the years 2009 to 2017. Data sources: Congressional Record, GDELT TV Archive and ProQuest respectively.

Beyond the clear political charge of the banned term, the setting of AP's ban has several features that make it attractive to study the causal effects of slanted language. First, the ban does not appear to be politically motivated and, according to AP's own statement, was part of a broader de-labeling policy in AP's Style Guide. For example, two months before the

[2]The reason for this divide can be traced back to deliberate party strategy. For example, the term "illegal immigrant" was advocated by Republican strategist Frank Luntz, who is famous for developing talking points for Republican candidates and for coining terms such as "death tax" (instead of "estate tax" or "inheritance tax") and "climate change" (instead of "global warming"). Luntz has urged Republicans to always use the term "illegal immigrant" and to put an emphasis on border security, calling the linguistic distinction between "illegal immigrant" and "undocumented immigrant" the "political battle of the decade" (Luntz 2007).

ban on “illegal immigrant”, AP banned the term “schizophrenic”, replacing it with “diagnosed with sciziphrenia”. This is further corroborated by the absence of pre-trend in AP’s slant on immigration, and the fact that the ban took effect in an extremely sharp fashion. Second, the majority of US legacy media, representing viewpoints across the ideological spectrum, are members of the AP. It is therefore unlikely that any *individual* outlet sways decisions on AP’s language rules. In other words, from the perspective of any individual outlet and the views of its readers, the ban can be viewed as producing exogenous variation in the editorial production function. Third and importantly for my empirical strategy, media outlets differ greatly in the extent to which they rely on AP’s input. This allows me to compare outlets with different degrees of use of this input, i.e. different *AP-intensity*, and the views of their respective readers before vs after the ban.

I start off my analysis by documenting how the ban affected AP’s content, using the text of all immigration-related AP dispatches released between 2009 and 2017. I find that, as intended, the ban caused the term “illegal immigrant” to instantaneously disappear from AP’s feed. At the same time, the volume of AP’s immigration coverage and other dimensions of AP’s slant on immigration, computed following the procedure of Gentzkow and Shapiro (2010a), remained largely unchanged. As a substitute for the label “illegal”, the new guidelines suggested the phrase “living in the county illegally” or “without legal permission”. However, text analysis reveals that these reformulations compensated for at most half of the decline in “illegal immigrant”. Hence, a significant part of the treatment in this natural experiment consists of substitution from “illegal immigrant” to “immigrant”, without any reference to legal status.

I then track how this change in AP’s language diffuses into the language of media outlets, using text data from more than 2200 print and online outlets. I employ a difference-in-difference strategy comparing the monthly number of “illegal immigrant” articles as a share of “immigrant” articles before and after the ban, in media outlets with different AP-intensity at baseline. Specifically, I measure AP-intensity as the share of “immigrant” articles published by each outlet in the 12 months prior to the ban that either credit AP explicitly, or are flagged by a plagiarism detection algorithm comparing their text to that of recent AP dispatches.[3]

My results suggest a large degree of diffusion – one standard deviation in AP-intensity causes a decline in the frequency of “illegal immigrant” articles by 14%. Put differently, outlets with positive AP-intensity decrease their use of the term by on average 28% compared to ones with zero AP-intensity, and for outlets in the top quartile of the AP-intensity distribution this decline reaches 60%. I find that on average, the diffusion effect is driven mostly by articles sourced from AP, as opposed to original ones. Hence, I find the same the (null) effects on volume of immigration coverage and on other aspects of immigration slant as document for AP dispatches, as well as the same gap in reference to legal status. A placebo test exploiting the intensity of another major newswire – *Reuters* – confirms that these results are not driven by general differences between outlets that rely on newswire content more or less.

Given the strong charge of the term “illegal immigrant”, one could expect differential reactions by left-leaning outlets, for which the ban is ideologically congruent, and by right-leaning ones, for which it is in-congruent. Indeed, I find a significant diffusion for both groups of outlets, but of a magnitude 2 times larger among left-leaning outlets compared to right-leaning ones. This is despite similar diffusion of AP-sourced articles and is instead driven by a spillover effect into the language used in originally produced content that is only present for left-leaning outlets.

I next exploit AP’s ban to identify the effect of exposure to “illegal immigrant” articles

[3]This procedure aims to capture the use of AP copy in cases when AP is credited as a source, and in ones in which AP is not credited (Cage et al. 2020).

on readers' views on immigration policy, using pre- and post-ban waves of the Cooperative Congressional Election Study (CCES). To identify the reduced form effect of the ban, I employ a difference-in-difference strategy comparing CCES respondents before and after the ban, in counties with different AP-intensity of locally circulated newspapers. Alternatively, to scale magnitudes in terms of the effect of "illegal immigrant" articles circulated in the respondent's county, I instrument their number (normalized by the number of "immigrant" articles) with the interaction of county-level AP-intensity and the timing of the ban. This strategy accounts for time-invariant effects of other county characteristics correlated with AP-intensity, but relies on the assumption that their effect on readers' views did not change in coincidence with the timing of the ban. I address this threat by controlling for a wide range of baseline county characteristics interacted with time.[4]

The results suggest that one standard deviation higher AP-intensity of locally circulated newspapers is associated with 0.7 percentage points lower support for increasing border security after the ban. For comparison, this corresponds to 1.8% of the gap in support for border security between Republican and Democrat respondents. It implies a persuasion rate, i.e. share of readers exposed to the treatment who changed their position (DellaVigna and Kaplan 2007; DellaVigna and Gentzkow 2010), in the range of 1.5 to 3.8%. These effects are robust to the inclusion of county controls, they occur in coincidence with the ban, and they do not reflect general reliance on newswire content as measured by *Reuters*-intensity.

While the above result applies to the sample of all CCES respondents, it is more pronounced for frequent print newspaper readers, who represent 33% of the sample. On the other hand, the effect is stronger among respondents with lower (self-reported) interest in politics, non-voters and independents. This is consistent with passive news consumers with weak priors on politically sensitive issues being more persuadable by slanted language. The effect is also more pronounced in counties with low shares of immigrants and Hispanics, where the issue of immigration is likely less salient and opinions on immigration policy more reactive to media framing.

I observe a similar shift in support for restricting immigration in 3 out of the 4 policy questions I am able to track across pre- and post-ban survey waves, as well in an index aggregating all immigration-related CCES questions including rotating ones. Specifically, I find significant effects on support for increasing border security, on allowing police to question suspected illegal immigrants, and on fining firms that employ illegal immigrants, and no significant effect on opposition to amnesty. On the other hand, a placebo exercise looking at other politically divisive policies suggests that these effects are specific to immigration. I find no significant change in responses on issues such as abortion, gay marriage or taxes and redistribution. On net, the relatively small and localized effect on immigration policy views appears (on average) insufficient to shift party support in national elections and I detect no effect on intentions to vote for Republican candidates or on electoral results. However I do find a significant effect on President Obama's approval, which suggests that the change in immigration policy views may have had some political repercussions in this period of heated debate over immigration reform.

[4]Importantly, my strategy includes time fixed effects, or alternatively, state-specific time fixed effects, which should absorb the effect of any major policy changes on policy views. County characteristics interacted with time should absorb any remaining heterogeneity in the effect of such confounds by socio-economic characteristics.

1.2 The Ban and It's Effect on AP's Language

1.2.1 Background

The term 'illegal immigrant" was dropped from AP's guidelines on April 3rd 2013. The decision was surprising since AP had previously resisted pressures from advocacy groups to change their language policy.[5] Up until the change was announced, AP's guidelines stated that "illegal immigrant" was the *preferred* term. The alternative endorsed by the left – "undocumented immigrant" – was not sanctioned because AP considers it legally inaccurate (as continues to be the case to this day). As a reason for the reversal on "illegal immigrant" AP mentioned the discussions surrounding this specific term but also a broader strategy to not label people that was being implemented in other areas.[6] For example, a month earlier AP had banned the term "schizophrenia", recommending instead "diagnosed with schizophrenia".

Appendix 1.5 presents the exact formulation of AP's guidelines before and after April 2013. As "illegal immigrant" was banned, the new guidelines proposed the following substitutes: "living / entering the county illegally / without legal permission".[7] Yet, AP executives recognized in their statement the possibility that these alternatives may "make it harder for writers" compared to the simple label "illegal". The ban took effect immediately in the online guidelines guidelines, which are also embedded in text editors (see Appendix Figure A1).

The ban was perceived as highly influential due to AP's dominant role in the US media landscape.[8] AP operates as a not-for-profit cooperative of about 1300 US newspapers and broadcasters. Members can have regular or associate status and in both cases have access to AP news and photos, share costs and contribute content. Regular membership, available to printed newspapers published in the US, also entails permission for the AP to distribute local news reports produced by the member.[9] While AP is the main newswire used in US media, Reuters offers competing services and is used as a substitute by some newspapers[10]. Therefore, in some parts of the analysis I exploit Reuters-intensity as the closest available placebo for AP-intensity.

1.2.2 Data

To analyze how the language used by AP changed in response to the ban and shed light on the nature of the treatment in this natural experiment, I obtain the text of all immigration-related AP dispatches released in the period July 1st 2009 to July 1st 2017. Specifically, I search the database of *Factiva* (`https://global.factiva.com`) for mentions of the word "immigrant" (singular or plural), restricting the source to "Associated Press Newswires". I record the date, headline, word-count and full text of each dispatch. This search results in 28,000 dispatches, 8,000 of which mention the phrase "illegal immigrant".

[5] `https://www.sfexaminer.com/national-news/society-for-professional-journalists-says-using-the-term-illegal-immigrant-is-unconsitutional/`

[6] `https://blog.ap.org/announcements/illegal-immigrant-no-more`

[7] According to the guidelines, the ban does not concern "illegal immigrant" used in direct quotes, nor the phrase "illegal *immigration*".

[8] Appendix Figure A2 presents the reaction of the Atlantic with the headline "The AP's Ban on 'Illegal Immigrant' Will Change How We Talk About Immigration" as one example. Full article available at `https://www.theatlantic.com/politics/archive/2013/04/ap-ban-illegal-immigrant/316701/`.

[9] `https://www.ap.org/about/annual-report/2016/AssociatedPress_2016FinancialStatements.pdf`

[10] `https://www.poynter.org/reporting-editing/2014/eighteen-months-after-dropping-ap-tribune-happy-with-reuters/`

1.2.3 Text Analysis Results

Descriptive evidence As a first check of how AP's language on the issue of immigration changed, figure 2 depicts the most frequent 3-grams encountered in headlines of "immigrant" dispatches before and after the ban. The label "illegal" clearly features prominently before the ban, and virtually disappears after.

To illustrate this difference in language more concretely, in Appendix 1.5 I present an example of two dispatches covering the same issue – a state law on immigrants' drivers licenses – released a month before and a month after the ban. Both dispatches have a neutral tone and highlight arguments from both the left and the right. However, the first dispatch talks about "non-citizens, including illegal immigrants", while the second talks about an "immigrant drivers' license bill".

Figure 2: Headlines of "immigrant" dispatches pre- and post-ban

(a) Pre-ban (b) Post-ban

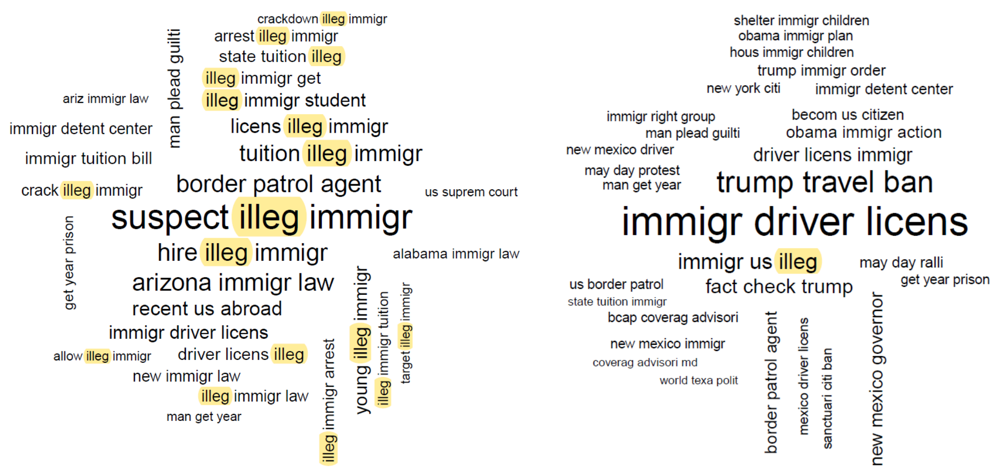

Notes: 50 most frequent tri-grams in the headlines of AP dispatches mentioning the word "immigrant", published before vs. after the ban.

The timing of this change coincides very precisely with the announcement of the ban. Panel (a) of Figure 3 shows the monthly number of AP dispatches mentioning the phrase "illegal immigrant" as percent of dispatches mentioning the word "immigrant". The share drops from an average of 50% in the period before April 2013, to less than 5% after, suggesting close to perfect compliance.[11] Panel (b) plots the frequency of the substitutes suggested by AP's guidelines. While their use increases sharply after the ban, the magnitude of this increase is relatively small (6 percentage points) and insufficient to compensate for the large decline in "illegal immigrant" (45 percentage points). At the same time, Panel (c) suggests that volume of "immigrant" dispatches did not change around the ban.

Substitutes for "illegal immigrant" I use two approaches to gather more exhaustive evidence on what language the banned term was substituted with in practice. First, I analyse the change in the words associations (i.e., the rate of co-occurrence) between "immigrant" and each other unigram contained in the AP-text corpus.[12] I plot the results for the top 50 correlates of

[11]Note that this figure includes mentions of "illegal immigrant" in direct quotes, which are not affected by the ban.

[12]I stem all words with the exception of "illegal" and "illegally" to account for the fact that while "illegal" was banned,"illegally" was, if anything, endorsed in the new guidelines.

Figure 3: Language of AP dispatches over time

(a) Share “illegal immigrant”

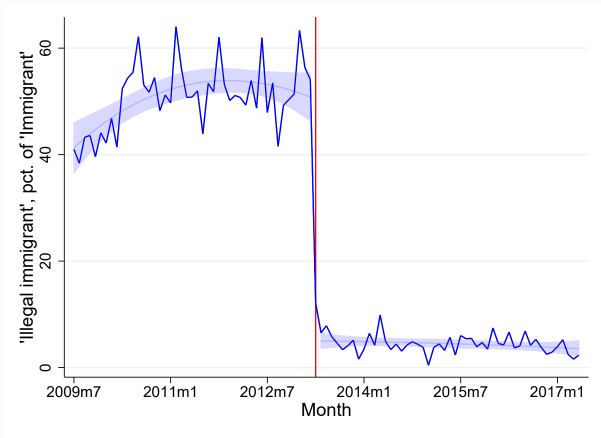

(b) Share of substitutes proposed by post-ban guidelines

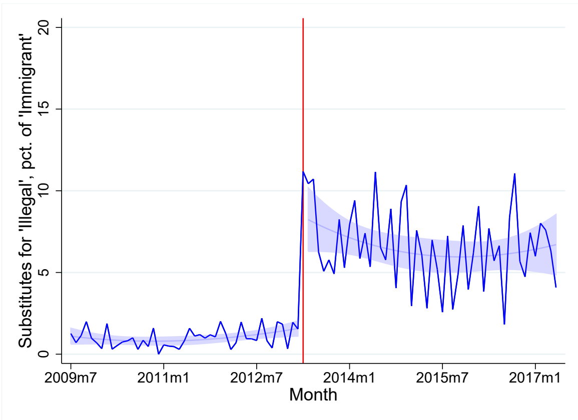

(c) Volume of “immigrant”

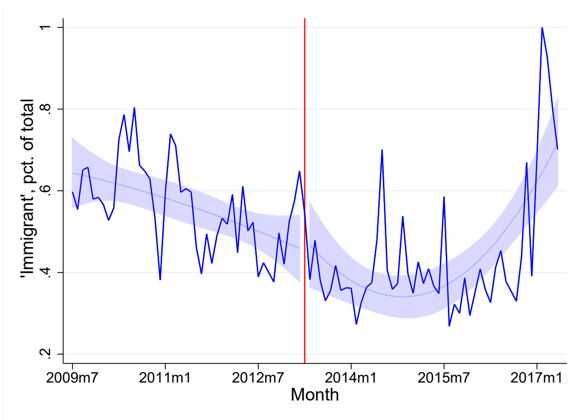

Notes: **Panel (a):** Monthly number of dispatches mentioning the phrase “illegal immigrant”, as percent of dispatches mentioning the word “immigrant”. **Panel (b):** Monthly number of dispatches mentioning the phrases “enter* / live* in the country illegally/ without legal permission”, as percent of dispatches mentioning the word “immigrant”. **Panel (c):** Monthly number of dispatches mentioning the word “immigrant”, as percent of total dispatches.

“immigrant” in Panel (a) of Figure 4, showing correlations in the full text on the left hand side, and correlations in headlines on the right hand side. The word “illegal” is clearly an outlier from the 45-degree line. Its correlation with “immigrant” drops dramatically from 0.66 (0.54) before the ban – the highest among all pre-ban correlates – to 0.21 (0.07) after the ban. Yet, the figure also suggests that no other unigram compensates for this decline. The closest candidate in article text is “illegally” – indeed, its correlation with“immigrant” increases significantly, but the magnitude corresponds to only about half of the decline of “illegal”. In headlines, the substitution is of an even smaller magnitude, likely due to the fact that the synonyms proposed by AP are inconvenient to use in a headline.

An alternative, and arguably more flexible approach to examine how language changes due to the ban, is to ask which words and phrases have the highest power in predicting whether a given AP dispatch was published before or after the ban. Let $f_{pl,before}$ and $f_{pl,after}$ denote the total number of times phrase p of length l (one to 3 words) is used before and after the ban, respectively. Let $f_{\sim pl,before}$ and $f_{\sim pl,after}$ denote the total occurrences of length-l phrases that are not phrase p – before and after the ban respectively. Let χ^2_{pl} denote Pearsonâs χ^2 statistic for each phrase:

$$\chi^2{}_{pl} = \frac{(f_{pl,before} f_{\sim pl,after} \hat{a} f_{\sim pl,after} f_{\sim pl,before})^2}{(f_{pl,before} + f_{pl,after})(f_{pl,before} + f_{\sim pl,before})(f_{pl,after} + f_{\sim pl,after})(f_{\sim pl,before} + f_{\sim pl,after})} \tag{1.1}$$

Panel (b) of Figure 4 presents the 20 words and phrases with highest χ^2_{pl}. "Illegal", "illegal immigrant" and "illegal immigrants" clearly emerge as the phrases most diagnostic of whether an article is published before or after the ban. Notably, "illegally" has only 1/4 of the predictive power of "illegal", confirming that synonyms indicating legal status were adopted only partially.

Figure 4: Words and phrases used in AP's "immigrant " dispatches before and after the ban

Correlates of the word "immigrant" before and after the ban

(a) Full text (inlc. headline) (b) Headlines only

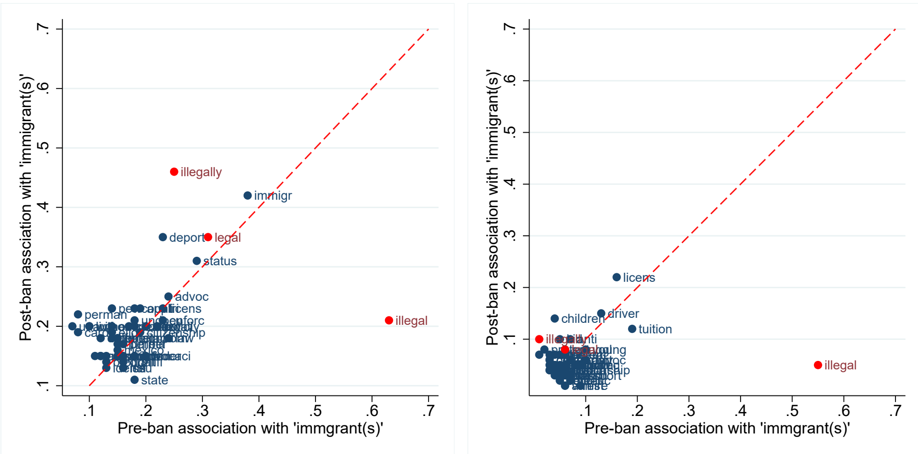

Phrases most predictive of post-ban publishing date

(c) Full text (inlc. headline) (d) Headlines only

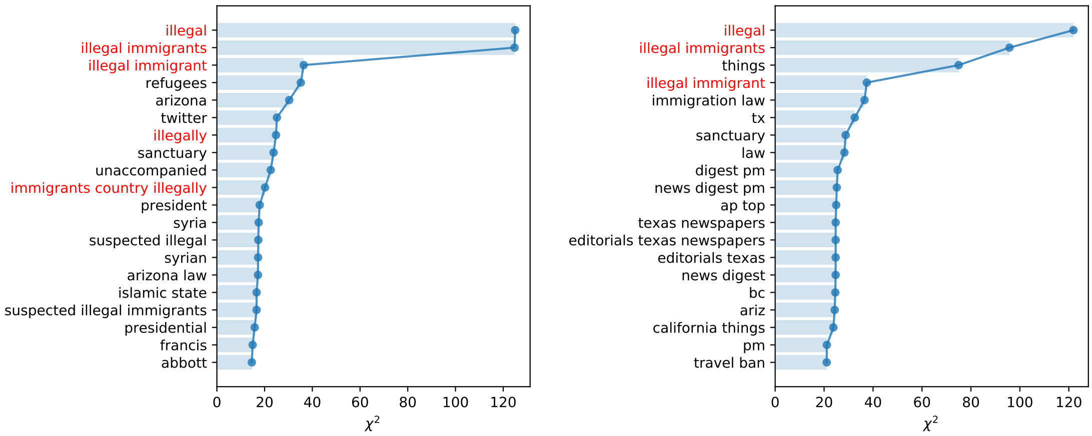

Notes: **Panel (a):** Top 50 unigrams with highest association with the word "immigrant" before and after the ban. Association defined as the rate of co-occurrence within the same dispatch. All unigrams are stemmed, with the exception of derivatives of "immigrant" and "illegal". **Panel (b):** Top 20 n-grams ($n \in 1, 2, 3$) in "immigrant" dispatches that are most predictive of a post-ban publishing date based on a χ^2 test-statistic.

To rule out the possibility that these results reflect a shift in topics occupying the news cycle over the sample period, in Appendix 1.5 I repeat the exercise separately for each of five topics estimated with a Latent Dirichlet Allocation (LDA) model.[13] The results discussed

[13]The estimated topics can be labeled as follows: law enforcement, immigration-related legislation, immigrants'

above are confirmed within each topic (with the exception of the χ^2 ranking within the topic of international affairs). Furthermore, I find that the distribution across topics is relatively similar in pre- and post-ban dispatches.

AP's overall immigration slant Finally, I examine whether the ban on "illegal immigrant" was part of a broader trend towards more liberal slant in AP's immigration coverage. To do so, I compute a measure of immigration-specific slant based on the similarity of AP's language to that used by Republicans vs Democrats in Congress when speaking about the issue. The procedure follows Gentzkow and Shapiro (2010a) in first selecting a set of phrases most predictive of partisanship in Congress, and then measuring their relative occurrence in AP's text over time (see Appendix 1.5 for more detail on the method). In order to isolate the influence of the ban from other dimensions of AP's language on immigration, I also compute a version of the index excluding the phrase "illegal immigrant" and its substitutes.

Panel (a) of Figure 5 presents the top partisan phrases that enter these measures: in line with intuition, phrases such as *"illegal immigrant"*, *"secure border"*, *"drug cartels"* are classified as Republican, while phrases such as *"undocumented immigrant"*, *"domestic violence"* and *"hate crime"* are classified as Democrat. Panel (b) shows and the evolution of the resulting 2 versions of the slant index over time. The version of the index that does not account for the ban on "illegal immigrant" follows closely the trend in AP's use of this phrase. This is intuitive since use of "illegal immigrant" is highly predictive of a Republican speaker in Congress (as already seen in Figure 1), and therefore receives a high weight in the measure of overall slant. However, once it is excluded and I focus on other dimensions of language, the trend in slant appears stable over time. Furthermore, neither measure of slant exhibits a clear pre-trend. This confirms the anecdotal evidence that the change in AP's language was sharp and unexpected. Appendix Figure C5 shows that similar results hold within each of the five LDA topics.

Taking these results together, the analysis of AP text suggests that: (1) As intended, the label "illegal" virtually disappears after the ban; (2) This decline is only partially compensated by the substitutes proposed by AP, while the remainder appears to omit any direct reference to legal status; (3) Other measurable dimensions of AP's language did not change significantly with the ban.

1.3 Diffusion

In next turn to analyzing the diffusion of AP's ban into the language of more than 2200 media outlets with different baseline reliance on AP copy ("AP-intensity").[14]

1.3.1 Data

Media content: Mentions of "immigrant" and "illegal immigrant" My main data source for media content is Newslibrary (`newslibrary.com`). I focus on print and online outlets that are covered continuously between July 1st 2009 and July 1st 2017 – there are 2566

integration and social issues, international issues such as the refugee crisis in Europe, and elections (Figure C1).

[14]This sample includes all US print and online outlets for which I am able to gather content data. However, the second stage analysis of readers policy views is restricted to the sample of print newspapers which allows me match geographic newspaper markets to the location of survey respondents. I obtain very similar results restricting the sample to print newspapers.

Figure 5: AP's slant on immigration

(a) Slant index computed for AP's immigration-related dispatches by quarter

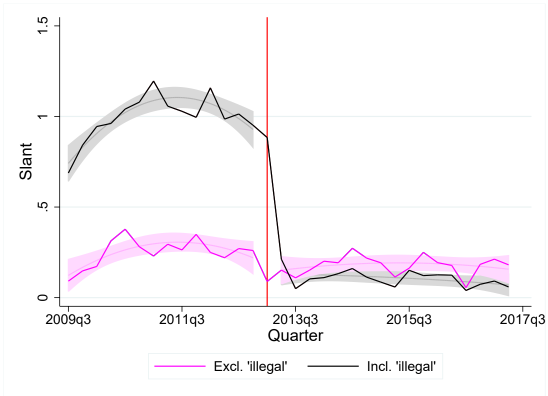

(b) Top partisan phrases in immigration-related Congressional speech

Phrases used more often by Republicans		
illegal immigrant	illegal immigration	enforce immigration
illegal alien	amnesty illegal	human smuggling
secure border	citizen legal	lottery program
federal government	taxpayer dollar	immigrant program
immigration law	visa lottery	american taxpayer
american people	insurance policy	social security
enforce law	yuma sector	country illegally
drug cartel	raise taxis	legal worker
free enterprise	immigration nationality	security number
illegal worker	national language	national medium
Phrases used more often by Democrats		
domestic violence	rhode island	charter school
violence woman	jewish american	federal employee
asian pacific	hate crime	immigrant student
pacific american	undocumented immigrant	visa program
victim domestic	american worker	heritage month
immigrant woman	sexual violence	house republican
young people	american community	comprehensive imm.
sexual assault	health care	violence sexual
rule pass	senate bill	protect victim
american woman	native american	homeland security

Notes:
Panel (a): Immigration-specific slant of AP dispatches over time. Higher values indicate more right-leaning slant. **Grey line:** Baseline measure of slant. **Magenta line:** Slant computed excluding phrases containing the term "illegal immigrant" or its substitutes. **Panel (b)**: Top partisan phrases derived from Congressional speech related to immigration.

such outlets. To cover some of the major US newspapers which are missing in Newslibrary, I supplement with data from ProQuest (`proquest.com`). This adds 125 newspapers.[15]

To construct measures of the language used in immigration coverage I search the database for articles that mention the phrase "illegal immigrant" (in singular or plural), and separately, for articles that mention the word "immigrant" (in singular or plural). This search results in about one million "immigrant" articles and 200,000 "illegal immigrant" articles. I record each article's date of publication, name of the publishing outlet, by-line, headline, word-count, and the text of the first paragraph.

Using this information, I compute for each outlet and each month the number of articles that mention "illegal immigrant" normalized by the monthly number of articles that mention "immigrant". I repeat the procedure with wordcount instead of number of articles, for the phrase "illegal immigration" as a percentage of "immigration", and for the potential synonyms "undocumented immigrant" and "unauthorized immigrant" normalized by "immigrant". Lastly, I collect mentions of the alternatives endorsed by AP – "living in / entering the country illegally" or "[...] without legal permission".

Identifying articles copied from AP I classify an article as sourced from AP if either one of two conditions is true: (1) AP is explicitly mentioned in the first paragraph (e.g. "according to AP"), or (2) a large portion of the text of the article is verbatim identical of to the text of a recent AP dispatch.

To capture the cases in which AP is credited explicitly, I search for mentions of "Associated

[15] Outlets that never mention the word "immigrant" in the sample period drop out of this sample.

Press" or "AP" in the lead paragraph or byline of the article. A similar procedure was employed by Gentzkow and Shapiro (2010a) to identify and, in their case, *exclude* news-wire content. Their audit of excluded articles suggests that "virtually all" articles identified in this way are indeed wire-copy. However, if media outlets use AP-content without explicit attribution, this procedure alone is likely to produce false negatives. Evidence on copying from the French news wire AFP suggests that this may indeed be a common occurrence (Cage et al. 2020).

Therefore, I additionally run the text of each article through a plagiarism-detection algorithm, comparing it to the set of dispatches released by AP on the previous day. The goal is to detect articles in which large portions of text are verbatim copy from an AP dispatch. In practice, my algorithm looks for overlap of sets of 5-grams that exceeds the threshold of 20% of text. I describe this procedure in detail and discuss summary statistics on crediting and plagiarism in Appendix 1.5.

AP-Intensity To proxy a media outlet's exposure to the change in AP's language, I measure the rate of copying from AP over the 12-months prior to the announcement of the ban. As a robustness check, I also select other pre-ban time windows further removed from the time of the ban (24 to 12 months or 24 to 36 months before the ban). I focus on the pre-ban period to avoid concerns about potential endogenous selection out of AP use in response to the change in AP's language policy. Indeed, Appendix Figure B3 suggests somewhat of a decline in the average extent of AP-use after the ban.

I then measure *AP-intensity* as the the number of articles copied from AP per 1000 articles in the selected pre-ban period – either credited to AP explicitly or flagged by plagiarism detection. Appendix Figure B2 presents the distribution of AP-intensity in my main sample of media outlets. This variable clearly features large variation, with the rate of AP-copying ranging from 0 to more than 750 AP-sourced articles per 1000. Since AP-intensity contains many zeros and has a skewed distribution, in the following analysis I use its inverse hyperbolic sine transformation.[16]

1.3.2 Empirical Strategy

To estimate the rate of diffusion from AP's language into that of AP-subscribing outlets, I implement a Difference-in-Difference strategy with continuous treatment. Specifically, I exploit the time-variation produced by the announcement of the ban and variation across media outlets in their exposure to the ban, proxied by AP-intensity. I estimate equations of the following form:

$$Illimm/Imm_{mt} = \alpha_m + \beta_t + \rho APintensity_m \times PostBan_t + \epsilon_{mt}, \tag{1.2}$$

where $Illimm/Imm_{mt}$ denotes the number of articles in media outlet m and month t that mention the phrase "illegal immigrant" as percent of articles mentioning "immigrant", AP_m is AP-intensity measured in the 12 months prior to the ban, $PostBan_t$ is a dummy for post-April 2013, and α_m and β_t are outlet- and calendar month FEs respectively. Standard errors are clustered at the outlet level. To account for the fact that $Illimm/Imm_{mt}$ is imprecisely estimated when the denominator, i.e. the number of "immigrant" articles is low, which is a frequent occurrence at monthly frequency, in my preferred specification this regression is

[16]About 20% of outlets in the baseline sample have positive AP-intensity (Table B1). This is the case for 40% of outlets in the sample of print newspapers (Table B2).

weighted by the number of “immigrant” articles.[17]

The identifying assumption in this strategy is that the frequency of ”illegal immigrant” articles in outlets with high AP-intensity vs outlets with low AP-intensity would have followed parallel trends in the absence of the ban. To examine the plausibility of this assumption as well as the timing of the effects, I estimate a dynamic version of the above equation, splitting the dummy for post-ban into a set of half-yearly leads and lags.

1.3.3 Results

Preliminary evidence Before proceeding to the estimation of the regression specified in 1.2, I examine visually the raw frequency of “illegal immigrant” articles in AP-intensive vs non AP-intensive outlets. Figure C7 shows these two series. While non AP-intensive media appear to gradually decrease their use of the term already prior to the ban, the use by AP-intensive media remains flat and quite high up until it exhibits a sharp decline coinciding with the ban. This pattern is in line with anecdotal evidence. For a long time, AP was resistant to demands to change their language policy, while in other media use of the term was gradually declining due to the controversy surrounding it. The figure also suggests that the ban was somewhat of an aggregate shock: even non AP-intensive media experience a decline at the time of the ban, albeit of a smaller magnitude. This is likely due to other outlets interpreting the ban – which was widely publicized – as a signal that the phrase “illegal immigrant” is no longer politically correct.Yet, the difference in magnitudes in the reactions of the two groups of outlets indicates that AP-intensity is a useful proxy for the degree of exposure to this aggregate shock.

Diffusion estimates Table 1 presents the main regression results corresponding to specification 1.2. I find a significant negative effect of the ban on use of the term – the magnitude suggests that one standard deviation increase in AP-intensity (=2.1) leads to 3 p.p lower frequency of “illegal immigrant” after the ban, or 14% relative to the mean. In addition to month fixed effects, in columns (3) and (7) I control for month times state fixed effects to absorb the effects of potentially confounding factors that vary over time and by state, such as the availability of state-specific newsworthy events related to illegal immigration. In columns (4) and (8) I control for outlet-specific linear time trends to account for possible differential trends depending on outlet characteristics correlated with AP-intensity. The estimates are stable to these controls, and if anything, *increase* in magnitude. Column (5) shows that despite not being officially banned by AP, use of the term “illegal immigration” also decreased (though by only half the magnitude of “illegal immigrant”).

Diffusion over time To verify that trends in the use of “illegal immigrant” in high- versus low-AP-intensity outlets did not start to diverge already prior the ban, I split the interaction of AP intensity and *Post Ban* into a set of interactions with half-yearly leads and lags. The results are plotted in Panel (a) of Figure 6. I find that if anything, the relative frequency of“illegal immigrant” increases up until the ban (in other words, trends were diverging rather than converging), at which point it falls abruptly. The decline is persistent, in line with the permanently low supply of “illegal immigrant” AP dispatches.

[17]The reason I normalize the dependent variable – number of “illegal immigrant” articles – by “immigrant” articles, is twofold. First, the coverage of the Newslibrary data is not universal and varies across outlets and years, which makes the absolute number of articles harder to interpret. Second, normalizing by the number of “immigrant” articles allows me to focus on changes in the *language* used to talk about the issue, conditional on the *volume* of coverage devoted to it.

Table 1: Diffusion of the ban depending on AP-intensity

	(1)	(2)	(3)	(4)	(5)
	'Illigal immigrant', pct. of 'Immigrant'				'Illegal immigration' pct. of 'Immigration'
PostBan × AP intensity	-1.490***	-1.462***	-1.426***	-1.737***	-0.976***
	(0.201)	(0.181)	(0.151)	(0.207)	(0.159)
AP intensity	1.716***				
	(0.215)				
PostBan	-12.497***				
	(0.757)				
Outlet FEs	No	Yes	Yes	Yes	Yes
Year-Month FEs	No	Yes	Yes	Yes	Yes
State × Year-Month FEs	No	No	Yes	Yes	No
Outlet-specific linear trend	No	No	No	Yes	No
Observations	133,349	133,347	133,329	133,329	106,412
Number of outlets	2271	2269	2269	2269	2150
R^2	0.15	0.42	0.49	0.53	0.34
Mean dep. var.	20.79	20.79	20.79	20.79	31.19

Notes: WLS weighted by number of number of "immigrant" articles in columns (1)-(4), and by number of "immigration" articles in column (5). Standard errors clustered by outlet.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Diffusion by bin of AP-intensity

In Panel (b) I relax another assumption implicit in equation 1.2 – that the effect is linear in AP-intensity. I estimate a more flexible specification discretizing the AP-intensity distribution. Specifically, I interact each quartile of the positive part of the AP-intensity distribution with $PostBan$, leaving outlets with zero AP-intensity as the reference category. The results suggest a roughly monotonic relationship in AP-intensity. The strongest effect comes from the top quartile, for which the effect amounts to a decline of 12 percentage points, or about than 60% relative to the mean.

Original vs AP-sourced articles If the documented decline in "illegal immigrant" articles is linked directly to the change in the language of AP dispatches, it should be driven by articles that are sourced from AP. To test this, in Panel (c) of Figure 6 I decompose the diffusion effect into articles copied from AP (with or without credit) vs. original content, and find that it is driven primarily by articles sourced from AP.

Robustness The result that outlets with higher AP-intensity decrease their use of the term"illegal immigrant" after the ban is stable to a number of alternative specifications and definitions of the variables of interest. In Table C1 I estimate specifications replacing the dependent variable with the number of "illegal immigrant" articles, dropping weights, replacing number of articles with word-count and with number of headlines, replacing continuous AP-intensity with a dummy for positive AP-intensity, and replacing *PostBan* with the time-series of "illegal immigrant" articles (normalized by "immigrant" articles) released monthly by AP.

In Table 2 I run the baseline regression with variations of the AP-intensity variable. Instead of accounting for both credited copying and plagiarism from AP, in column (2) I consider only the share of articles credited to AP, and in column (3) – only the share of articles flagged

Figure 6: Diffusion of the ban

(a) Dynamic diff-in-diff estimates

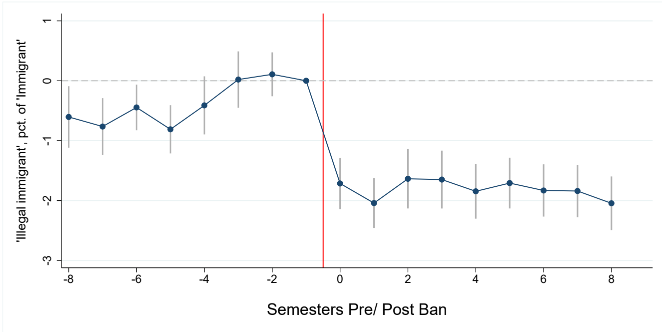

(b) Effects by quartile of AP-intensity

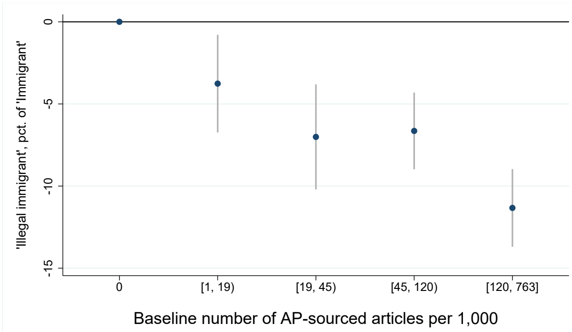

(c) Dynamic diff-in-diff estimates for AP-sourced vs original articles

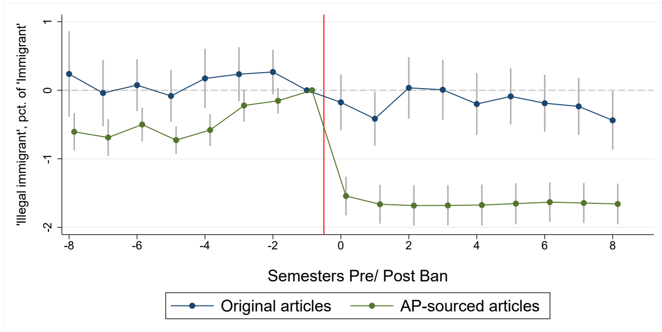

Notes: **Panel (a):** Coefficients and 95% confidence intervals from a regression of the frequency of "illegal immigrant" articles as percent of "immigrant" articles on full set of indicators for semester pre-/post-ban interacted with AP-intensity, controlling for outlet and year-month FEs. The omitted category is the semester before the ban. **Panel (b):** Coefficients and 95% confidence intervals from a regression of frequency of "illegal immigrant" articles as percent of "immigrant" articles on a full set of indicators for quartile of (positive) AP-intensity interacted with Post Ban, controlling for outlet and year-month FEs. The omitted category is AP-intensity $= 0$. **Panel (c):** Weighted by number of "immigrant" articles. Standard errors clustered by outlet.

by plagiarism detection. The two measures have a correlation of 0.56 and yield very similar results to the baseline. In column (4), rather than examining the sample of 'immigrant" articles, I consider articles on any topic and define AP-intensity as the share of total articles published the 12 months before the ban that credit AP. Finally, as a placebo exercise, in column (5) I consider use of Reuters rather than AP. Since the Reuters news-wire did not change their style rules regarding "illegal immigrant", prior reliance on Reuters should not be associated with the degree of reaction to the ban. Indeed, I find no change in use of the term depending on Reuters-intensity.

Heterogeneous diffusion by ideological leaning Media outlets decide on the extent to which they want to use AP dispatches and are free to edit their language as they wish. Therefore, a natural question is whether the diffusion effect differs by the ideological position of the outlet, i.e. by how congruent the ban is with their editorial policy. To answer this question I analyze a sub-sample of about 340 newspapers which I can match to the index of ideological leaning constructed by Gentzkow and Shapiro (2010a). I split this sample at the 33rd and 66th percentile with respect to this measure of ideological leaning, and label the 3 resulting groups of outlets as "left-leaning", "center" and "right-leaning".

In Panel (a) of Figure 7 I estimate the overall diffusion effect separately for each of the

Table 2: Alternative measures of AP-intensity

	(1)	(2)	(3)	(4)
	'Illigal immigrant', pct. of 'Immigrant'			
PostBan × AP-intensity: only AP credited	-1.437***			
	(0.191)			
PostBan × AP-intensity: only AP plagiarised		-1.434***		
		(0.209)		
PostBan × AP-intensity: only AP credited, *all articles*			-1.318***	
			(0.201)	
PostBan × Reuters-intensity: only Reuters credited, *all articles*				0.280
				(0.362)
Outlet FEs	Yes	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes	Yes
Observations	133,347	133,347	123,261	129,344
Number of outlets	2269	2269	2218	2421
R^2	0.42	0.42	0.39	0.40
Mean dep. var.	20.79	20.79	21.39	21.16

Notes: Replication of column (3) of table 1 with the following alternative measures of AP-intensity. Column (1): share of "immigrant" articles credited to AP. Column (2): share of "immigrant" articles flagged by a plagiarism algorithm. Column (3): share of all articles published in the 12 months before the ban that are credited to AP. Column (4): share of all articles published in the 12 months before the ban that are credited to Reuters. Standard errors clustered by outlet.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3 groups. In each case, I find a significant decline in "illegal immigrant" articles.[18] Yet, the magnitude of diffusion is significantly stronger among left-leaning outlets compared to center and right-leaning ones.

In Panel (b) I repeat the analysis looking only at articles sourced from AP. Interestingly, in this case I find very similar diffusion effects for the 3 groups of outlets, suggesting a limited role of screening and filtering out of AP-dispatches. Instead, in Panel (c) I focus on original articles and find significant differences – here the effect of "illegal immigrant" is negative and significant only for the group of left-leaning outlets, while for the 2 other groups I find, if anything, slightly positive point estimates. In other words, the ban has significant a spillover into original content, but only for outlets whose ideological position is congruent with this change in language.

The above results suggest that even right-leaning outlets comply with AP's ban to a significant extent, despite their (likely) ideological opposition to it. In other words, for the average right-leaning newspaper, inertia in the reliance on AP in content production appears to outweigh ideological considerations.

Another margin along which the ban could have faced resistance is the reaction of readers (Durante and Knight 2012). To examine possible effects on readership, I use data on circulation by newspaper and year available from the Alliance for Audited Media for about 600 newspapers. I estimate difference-in-difference regressions at this level, and with log circulation as a dependent variable. I furthermore split the sample into the 3 ideology categories defined above.

[18]Figure B4 presents the distribution of the difference in the share of "illegal immigrant" article post- versus pre-Ban by outlet. This difference is negative for 80% of outlets in the main sample, and for 90% of outlets in the restricted sample of print newspapers. This is indicative of mostly positive compliance, though its degree varies widely across outlets.

Figure 7: Diffusion by Ideology

(a) Effect on "illegal immigrant" articles

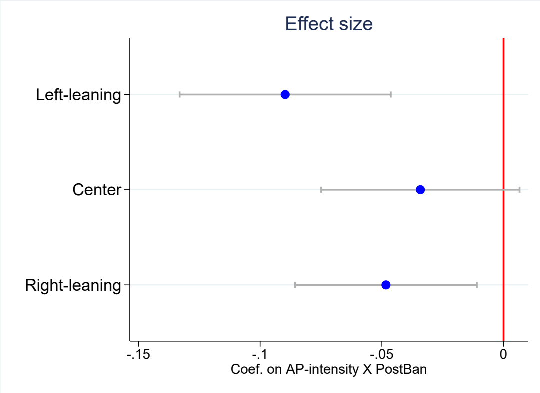

(b) "Illegal immigrant" articles sourced from AP

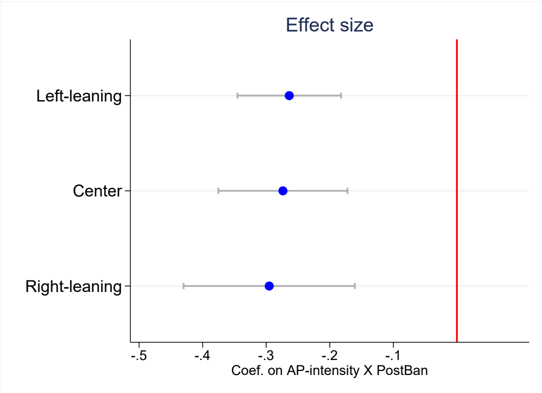

(c) Original "illegal immigrant" articles

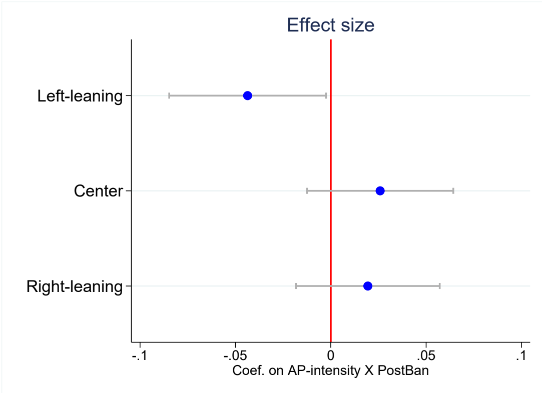

Notes: **Left-leaning**, **Center**, **Right-leaning** denote the sample of outlets in the 1st, 2nd and 3rd tercile of the distribution of the Gentzkow and Shapiro (2010a) slant index respectively.
The graphs present coefficients and 95% confidence intervals on the interaction of AP-intensity and PostBan from regressions restricted to one of the 3 samples at a time, with the following dependent variables (standardized to facilitate comparison of the coefficients). **Panel (a):** frequency of "illegal immigrant" articles as percent of "immigrant" articles; **Panel (b):** frequency of "illegal immigrant" articles sourced from AP as percent of "immigrant" articles; **Panel (c):** frequency of "illegal immigrant" articles *not* sourced from AP as percent of "immigrant" articles. Each regression controls for outlet and year-month FEs and is weighted by number of "immigrant" articles. Standard errors clustered by outlet.

Figure C2 shows that I detect no significant effect of the ban – neither on average, nor among right-leaning newspapers. This is consistent with the subtle nature and direction of the change – right-leaning readers would need to notice the *absence* of the term "illegal immigrant" in order to react by switching away from their preferred newspaper.

Other aspects of immigration coverage Having established that the ban on "illegal immigrant" in AP dispatches diffused into the *language* of media outlets, I turn to testing whether other measurable aspects of coverage were affected.

As with AP-dispatches, in Panel (a) of Figure 8 I find that the synonyms proposed by AP ("live(-ing)/enter(-ing) the country illegally / without legal permission") compensate some but not all of the decline in the phrase"illegal immigrant". The figure also suggests that this compensating effect tapered off over time.

Also consistent with AP-dispatches, in Panel (b) I find that the number of articles mentioning the word "immigrant" (normalized by total articles) was not affected by ban. The same null-effect applies to articles mentioning the word "immigration" over total articles.

Finally, I examine the effect of the ban on news outlets' immigration-specific slant. As with AP's text, I compute an index of slant based on the similarity of the language of a given news outlet to that of Republican vs Democrat representatives in Congress.[19] Panel (c) of Figure 8 presents the dynamic effect of the ban on slant, regressing the two versions of the index (computed at the level of news outlet $\times$ year) on the outlet's AP-intensity interacted with year fixed effects. The figure resembles the evolution of AP's slant (Figure 5) and suggests that the ban has a significant effect on the index that includes the phrase "illegal immigrant", but no clear effect on the version excluding the this phrase and its synonyms.

To sum up, several measurable features of the language of AP dispatches diffuse into the language used by media outlets, consistent with the result that copied articles drive the majority of the diffusion effect. Crucially, the volume of immigration coverage and its slant, apart from the component driven by the banned phrase, appear to be unaffected. Hence, AP's ban can be interpreted to cause a shock to the use of "illegal immigrant", while leaving other features of coverage largely unaffected.

1.4 Effects on Readers' Views on Immigration Policy

Having established that the ban produces significant variation in media outlets' use of the term "illegal immigrant", in this section I analyze how it affected public opinion on immigration policy. To this end, I compare pre- and post-ban responses in the CCES electoral survey for respondents living in counties with different AP-intensity of locally circulated newspapers.

1.4.1 Data

1.4.1.1 Aggregation to the county level.

Since the CCES survey does not ask *which* newspaper the respondent reads, I rely on county of residence to assess exposure to locally circulated newspapers. Therefore, the first step in this analysis is to aggregate my measures of newspapers' content to the county level. To do so, I obtain data on the geographic distribution of daily newspapers' circulation from Alliance of Audited Media (AAM). I use their Fall 2012 report, which includes circulation by newspaper and zip-code from the most recent audit prior to this date, and aggregate zip-code level data to the county level.[20] Finally, since AAM does not collect geographically disaggregated data for low-circulation newspapers, I impute these observations with data on total circulation from the Editor and Publishes yearbooks, assuming that small newspapers circulate mainly in the county of their headquarters.[21] [22]

I match this data to the sample of Newslibrary/ ProQuest media outlets based on the name, town and state of the newspaper. I then keep counties for which newspapers matched to the

[19]See appendix 1.5 for details on this procedure.

[20]For the largest nationally circulated newspapers AAM only reports circulation at the DMA level. For these cases I assign circulation to counties in proportion to voting-age population.

[21]The same procedure is used by Seamans and Zhu (2014).

[22]Similar results obtain using data on circulation by newspaper and county compiled by Snyder and Strömberg (2010). This data is based on circulation reported to the Audit Beaurau of Circulation (ABC, which later became "Alliance for Audited Media"), combined with the Standard Date and Rate Service for non-ABC newspapers. The drawback of this dataset is that it only covers the period up to 2006.

Figure 8: Other measures of immigration coverage over time

(a) Substitutes proposed by AP

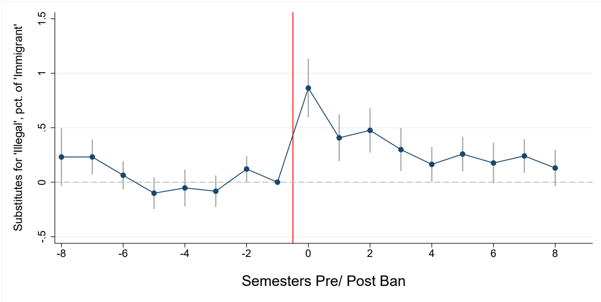

(b) Volume of immigration coverage

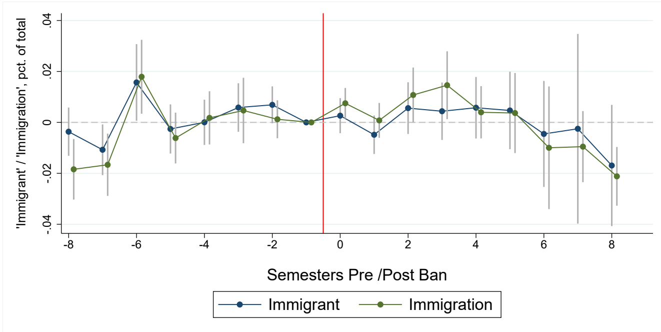

(c) Immigration slant

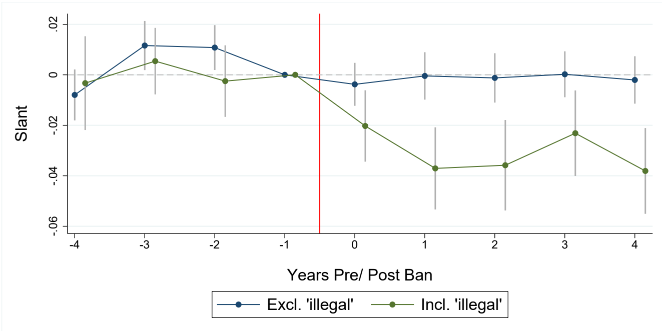

Notes: Coefficients and 95% confidence intervals on AP-intensity interacted with a full set of indicators for semester pre-/ post-ban. In **Panel (a)** the dependent variable is the number of articles mentioning the substitutes proposed by AP ("enter*/ live* in the country illegally/ without legal permission") as percent of "immigrant" articles. In **Panel (b)** the dependent variable is the number of articles mentioning the words "immigrant" or "immigration" as percent of total articles. In **Panel (c)** the dependent variable is the index of immigration slant computed including (green line) or excluding (blue line) the phrase "illegal immigrant" and its substitutes. All regressions control for outlet and year-month FEs. The omitted category is the semester before the ban. In Panel (a) the regression is weighted by number of "immigrant" articles and in Panel (b) by total articles. Standard errors clustered by outlet.

Newslibrary/ ProQuest sample account for at least 90% of total county circulation.[23] This ensures that the county-level data on newspapers' content is measured with reasonable precision. The resulting dataset contains about about 2300 counties (out of a total of 3000), and 800 daily newspapers (out of a total of 1200).

I aggregate AP-intensity to the county level by averaging the AP-intensity of newspapers circulated within the county (in number of AP-sourced articles per 1000), weighting each newspaper by its county-specific circulation. Formally:

$$AP_c = \frac{\sum_m (circ_{mc} \times AP_m)}{\sum_m circ_{mc}}, \tag{1.3}$$

where $circ_{mc}$ is circulation of newspaper m in county c. As in the outlet-level analysis, I take the inverse hyperbolic spline transformation of this variable. Panel (a) of Figure 9 presents the resulting geographic distribution of AP-intensity. Similarly, I aggregate the percentage of "illegal immigrant" relative to "immigrant" articles by county and year, again weighting by circulation:

[23]The results are robust to alternative thresholds, see table **??**.

$$Illimm/Imm_{cy} = \frac{\sum_m (circ_{mcy} \times Illimm/Imm_{my})}{\sum_m circ_{mcy}}. \tag{1.4}$$

Correlates of AP-intensity In order to understand the correlates of AP-intensity, I collect data on county-level demographic, economic and political characteristics. Data on annual county population is from ICHS. Data on the urban share of population is from the 2010 census. Racial composition, share college educated and share foreign-born are from the 2012 5-year American Community Survey, and the Republican vote share in the 2012 presidential election is from Dave Leip's Atlas. Finally, county-level newspaper circulation per capita is estimated with data from the Alliance of Audited Media combined with the Editor and Publisher yearbooks.

Panel (b) of Figure 9 presents the univariate correlations of each of these variables with AP-intensity. AP-intensity is significantly negatively correlated with population size and density, with the share of college educated, with the share of foreign-born and with county-level newspaper circulation. This is consistent with the notion that smaller newspapers in less urban areas are more likely to resort to sourcing content from AP, rather than producing original reporting. A more urban, higher educated audience, as well as one with more immigrants may also have higher demand for original content, particularly on immigration.[24]

Figure 9: Geographic distribution of AP-intensity and county-level correlates

(a) Geographic distribution of AP-intensity (b) County-level correlates of AP-intensity

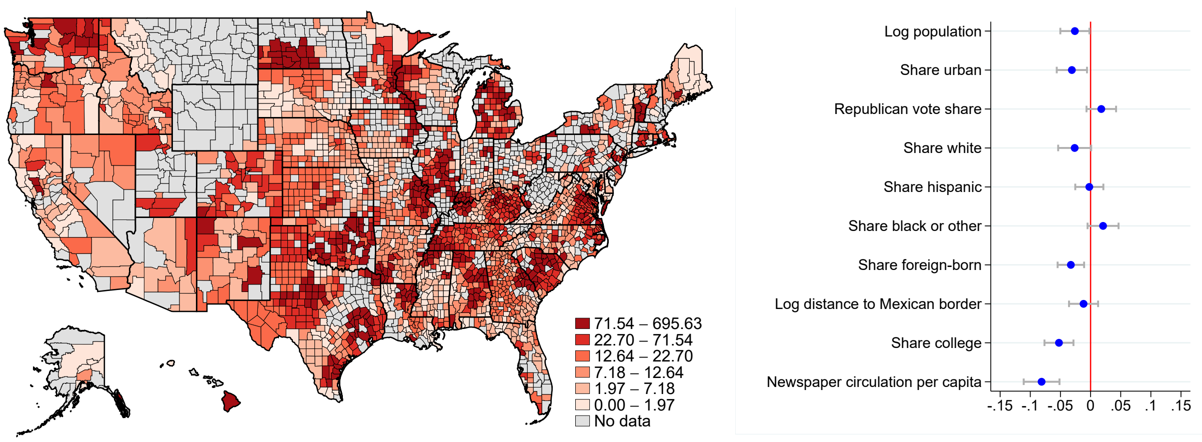

Notes: Panel (b): Coefficients and 95% confidence intervals from univariate regressions of each of the listed county characteristics on AP-intensity. All county characteristics are standardized to facilitate comparison of the magnitudes of the coefficients. Robust standard errors.

1.4.1.2 The CCES Survey

To assess how public opinion on immigration policy changed in response to the ban, I use a large nationally representative survey – the *Cooperative Congressoional Election Study* (CCES). CCES is a repeated cross-section with more than 50,000 respondents per wave (with smaller

[24]For illustration, the 5 newspapers with highest AP-intensity in my sample are the: *The Westerly Sun (RI)* (651 AP-sourced articles per 1000), *The Telegraph Herald (Dubuque, IA)* (653 articles per 1000), *The Logan Banner (WV)* (692 articles per 1000), *The Bismarck Tribune (ND)* (699 articles per 1000) and *The Breeze-Courier (Taylorville, IL)* (763 articles per 1000).

waves in some years), carried out roughly every 2 years.[25] The survey is administered online and a large portion of participants are YouGov panelists. Conveniently for my setting, a large share of survey respondents (33%) report that they regularly read a newspaper in print (i.e. that they have done so in the day before the survey).

Views on Immigration Policy Each CCES respondent is asked to select immigration policies she thinks the US government should undertake, out of a list of options. The set of policies differs in each wave with questions ranging from allowing the police to question suspect unauthorized immigrants to building a wall between the US and Mexico. Appendix 1.5 presents the full list of questions and their exact formulation in the survey. Two policies appear consistently in all survey years between 2009 and 2017: *"Increase the number of border patrols on the U.S.-Mexican border"* and *"Grant legal status to all illegal immigrants who have held jobs and paid taxes for at least 3 years, and not been convicted of any felony crimes"*.

For each policy, I code support for *restricting* immigration (e.g. increasing border control/ *not* granting amnesty) as 1, and opposition as 0. I also compute an index aggregating choices on *all* 9 immigration policies featured in the questionnaire in the respective year, including rotating ones. I recode each choice in the direction of restricting immigration, and take the average across all standardized choices (following Kling et al. (2007)).

Views on policies other than immigration and voting As placebo outcomes, I also use CCES questions that relate to policy issues other than immigration. Specifically, I create (1) A dummy variable for opposing a woman's right to choose to have an abortion under any circumstances; (2) A dummy variable for preferring to cut public spending rather than increase taxes; (3) A dummy variable for opposing gay marriage; (4) A dummy variable for believing that the state of the economy has gotten worse over the past year.

The survey also asks about the respondent's voting intentions for upcoming presidential, Senate and House elections, as well as about the respondent's approval of the president (i.e. president Obama who is incumbent throughout my baseline sample period). I code dummy variables equal to one if the respondent intends to vote for a Republican candidate for a given office, and a dummy equal to one if the respondent disapproves of Obama's performance in office.

I also obtain county-level Republican vote shares in the 2012 and 2016 presidential elections (i.e. share of votes for Romney and Trump respectively) from David Leip's Election Atlas.

1.4.2 Empirical Strategy

To identify the local average treatment effect of exposure to the phrase "illegal immigrant" on views on immigration policy, I estimate 2SLS equations of the following form:

$$X_{cy} = \alpha_c + \beta_y + \rho \widehat{Illimm/Imm}_{cy} + \epsilon_{cy}, \tag{1.5}$$

$$Illimm/Imm_{cy} = \alpha_c + \beta_y + \gamma AP_c \times PostBan_y + \epsilon_{cy} \tag{1.6}$$

where X_{cy} denotes immigration policy preferences of respondents in county c and year y, $Illimm_{cy}$ denotes the percent"illegal immigrant" relative to "immigrant" articles read in

[25] To the best of my knowledge, CCES is the only large-scale survey conducted between 2009 and 2017 that asks questions related to views on immigration policy.

that county and year, AP_c is the average AP-intensity of newspapers circulated in county c, $PostBan_y$ is an indicator equal to one for survey waves carried out after 2013, and α_c and β_y are county and survey-year fixed effects respectively. Standard errors are clustered by county.

The first-stage equation has the same form as the difference-in-difference specification from the previous section, but now estimated at the county and survey-year level (instead of media outlet and month). The excluded instrument for the potentially endogenous frequency of "illegal immigrant" articles $Illimm/Imm_{cy}$ is the interaction of AP-intensity with an indicator for the period after the ban $AP_c \times PostBan_y$. The approach is thus akin to a shift-share strategy where $PostBan$ is an aggregate shock and AP_c is local exposure to that shock.

Since the identifying variation is at the county $\times$ survey-year level, this equation can be estimated by aggregating individual survey responses up to that level. Alternatively, it can be estimated at the respondent-level. This has the advantage of allowing to control for respondent characteristics which are likely to correlate with immigration policy attitudes.

The identifying assumption is that the interaction of AP intensity with the timing of the ban affects policy views only through exposure to the term "illegal immigrant". Time-invariant county-characteristics correlated with AP intensity are absorbed by county fixed effects. Therefore, if observed or unobserved characteristics are to confound my results, their effect on attitudes would have to *change* at the same as the ban took effect. To account for this possibility, I examine the sensitivity of the estimates to controlling for a host of county characteristics measured at baseline and interacted with survey-year fixed effects (see Figure 9 for the list of controls and their correlation with AP-intensity).

Additionally, a strict causal interpretation of the LATE estimates requires the assumption that no other aspect of coverage changed in response to the ban, except for the share of "illegal immigrant" articles. This assumption is supported by my analysis of content although I can not rule out other, potentially non-measurable changes in content or tone. Yet, even under valuation of this assumption the reduced form intention-to-treat estimates for the effects of are valid with slightly different interpretation as the effect of the combined changes changes in content induced by the ban.

Finally, the effect identified by equation 1.5 is a local average treatment effect – it applies to readers of newspapers that change their language on immigration solely due to the change in the input supplied by AP. Such newspapers are likely to have a less pronounced stance on immigration policy and potentially more persuadable readers compared to the readers of always- or never-takers.

1.4.3 Results

First stage I start off by replicating the analysis of the diffusion of the ban for this new sample and unit of observation, i.e. aggregating newspapers' data to the county times year level (the 1st stage of equation 1.5). The results presented in Figure 10 suggest that in this sample 1 standard deviation increase in AP-intensity (= 1.5) is associated with 9.5% lower use of the term "illegal immigrant" after the ban.

Intention to treat effects I then turn to the reduced form intention-to-treat effect of the ban on support for restrictive immigration policies. In column (1) of table 3, I examine the effect on an index aggregating all immigration-related CCES questions, conditional on respondent characteristics and baseline county controls interacted with time. In columns (2) to (5) I examine each component of the index that I am able to look at separately, i.e. each question that is asked at least once before the ban and at least once after. With the exception of the question on

Figure 10: Diffusion over time: county × year level

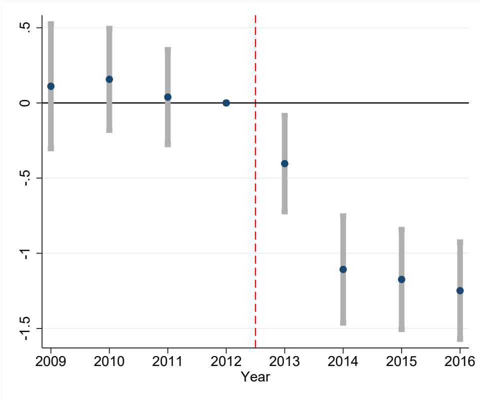

Notes: Point estimates and 95% confidence intervals on the interactions of AP-intensity with year, conditional on year and county FEs. Standard errors clustered by county.

amnesty, the results suggest a significant negative reduced form effect of the ban of support for restrictive policies. The magnitudes range from 1.2% to 2% reduction in support for a given policy for 1 standard deviation higher AP-intensity. Similar results obtain at the county-level, with the dependent variable collapsed by county times survey-year (Appendix Table C12).

Table 3: Views on immigration policy: Reduced form

	Reduced Form				
	(1) Index Restict Imm.	(2) Border	(3) No Amnesty	(4) Don't hire	(5) Question
PostBan × AP-intensity	-0.0112**	-0.0043***	-0.0006	-0.0075***	-0.0044**
	(0.005)	(0.002)	(0.002)	(0.002)	(0.002)
Respondent controls	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
Observations	162,456	162,456	162,456	74,705	119,552
Number of counties	2,113	2,113	2,113	1,924	2,066
R^2	0.27	0.14	0.16	0.13	0.22
Mean dep. var.	0.01	0.56	0.52	0.62	0.41

Notes: Reduced form OLS regressions. Respondent controls: age, age squared, gender, race, college, 1st or 2nd generation immigrant, and political ideology. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Local average treatment effects In table 4 I estimate the 2SLS version of equation 1.5 for the same set of outcomes, again conditioning on respondent characteristics and county controls interacted with time. Here, the coefficient of interest is the second stage effect of locally circulated "illegal immigrant" articles on support for restrictive immigration policies. The results mirror those of the reduced form – an increase in such articles has a significantly positive effect on support for restrictive policies (with the exception of amnesty). The magnitudes range from 0.9% to 1.4% increase in support for a given policy for 1 percentage point (or 4.8%) higher share of locally circulated "illegal immigrant" articles. These results are also confirmed at the county level (Appendix Table C13).[26]

Table 4: Views on immigration policy: 2SLS

	2SLS				
	(1) Index Restict Imm.	(2) Border	(3) No Amnesty	(4) Don't hire	(5) Question
'Illegal imm.', pct. of 'Imm.'	0.0116**	0.0045**	0.0006	0.0065**	0.0050**
	(0.005)	(0.002)	(0.002)	(0.003)	(0.002)
Respondent controls	Yes	Yes	Yes	Yes	Yes
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	19.67	19.67	19.67	16.62	11.31
Observations	162,456	162,456	162,456	74,705	119,552
Number of counties	2,113	2,113	2,113	1,924	2,066
R^2	0.22	0.10	0.12	0.09	0.16
Mean dep. var.	0.01	0.56	0.52	0.62	0.41

Notes: 2SLS regressions (upper panel), along with the corresponding 1st-stage coefficients (lower panel). Respondent controls: age, age squared, gender, race, college, 1st or 2nd generation immigrant, and political ideology. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Sensitivity to controls Since the border question is the only one (apart from the one on amnesty) that is asked in each CCES wave in the period of interest, I focus on this question for the remainder of this section. This has the advantage of holding the definition of the dependent variable constant over time, whereas the index aggregates policies of different severity in each wave, making comparisons over time harder to interpret.

Table 5 presents the reduced form and 2SLS effects on support for border security with alternative controls. In the first column, instead of including county and year fixed effects, I present the main effects of $PostBan$ and $AP{-}intensity$. The results mimic those from table 1. Consistent with the fact that AP-intensive outlets had a higher frequency of "illegal immigrant" articles before the ban, immigration policy views in such counties were more conserve before

[26] Appendix Table C6 presents the OLS equivalent of the relationship between share of locally circulated "illegal immigrant" articles and immigration policy views. The unconditional correlations (Panel a) are positive and highly significant throughout – as expected, in the cross section slant is strongly related to policy views but the direction of causality of this relation is not clear. Instead, and in contrast to the LATE estimates discussed above, in specifications saturated with fixed effects and controls, I instead find null OLS coefficients (Panel b). This difference is likely due to the nature of compliance to the ban. The group of group of compliers consists of newspapers that only change their language due to AP's input. In contrast to always-takers (who would have changed their language regardless) and never-takers (who do not change their language despite the ban), such newspapers are likely to have a less pronounced ideological stance on immigration in either direction, and hence, more persuadable readers compared to the general public.

the ban (main effect on AP-intensity is positive). As with its effect on use of "illegal immigrant", the ban appears to be somewhat of an aggregate shocks to views on border security (the main effect of $PostBan$ is negative), but it is amplified by AP-intensity. The coefficient on the interaction of PostBan with AP-intensity is stable to the inclusion of fixed effects and to county controls interacted with time, which absorb the possibly changing effect of these controls on readers' views. It is also robust to the inclusion of year $\times$ state fixed effects, which absorb the effect of any state-level policy changes – if anything, the 2SLS increase in magnitude.

Table 5: Support for increasing border control

	Reduced Form				2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	"Increase the number of border patrols on the US-Mexican border.": Selected						
PostBan $\times$ AP-intensity	-0.0044***	-0.0049***	-0.0046***	-0.0046**			
	(0.002)	(0.002)	(0.002)	(0.002)			
AP intensity	0.0047***						
	(0.002)						
PostBan	-0.0186***						
	(0.006)						
'Illegal imm.', pct. of 'Imm.'					0.0055**	0.0050**	0.0065**
					(0.002)	(0.002)	(0.003)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs $\times$ County controls	No	No	Yes	Yes	No	Yes	Yes
County FEs	No	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	No	Yes	Yes	Yes	Yes	Yes	Yes
Year $\times$ State FEs	No	No	No	Yes	No	No	Yes
First-Stage F stat.	.	.	.	.	9.90	23.63	12.19
First stage coef on PostBan $\times$ AP-intensity					-0.8859***	-0.9223***	-0.7104***
					(0.282)	(0.190)	(0.204)
Observations	162,057	161,943	161,490	161,490	161,943	161,490	161,490
Number of counties	2,236	2,122	2,113	2,113	2,122	2,113	2,113
R^2	0.12	0.14	0.14	0.14	0.10	0.10	0.10
Mean dep. var.	0.56	0.56	0.56	0.56	0.56	0.56	0.56

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent controls: age, age squared, gender, race, college, 1st or 2nd generation immigrant, and political ideology. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share. Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Reduced form over time To examine the dynamics of the reduced form effect, I estimate a regression including a full set of interaction of AP-intensity with indicators for survey waves, leaving the 2012 as the baseline category. In this analysis I can furthermore add the survey years 2007 and 2017, in order to examine longer-term trends. The results, presented in Figure 11 show no evidence of pre-trends â instead, the shift in policy views happens in the period after the ban, and remains roughly constant in following waves.

Reduced form by bin of AP-intensity In figure 12, I estimate a flexible version of the reduced form equation, splitting the distribution of AP-intensity into quartiles and interacting each one with an indicator for the period after the ban, leaving the first quartile as the baseline category. The results suggest that the effect is monotonic in AP-intensity.

Figure 11: Support for increasing border security: Reduced form effects over time

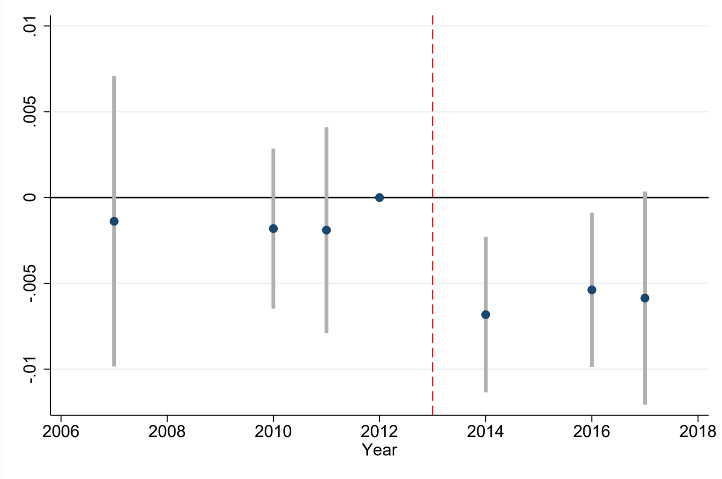

Notes: Point estimates and 95% confidence intervals on the interactions of AP-intensity with survey year, conditional on year and county FEs, respondent controls, and county controls interacted with year FEs. Respondent controls: age, age squared gender, indicators for race, college, and 1st or 2nd generation immigrant. County controls: log population, racial composition, share foreign born, share college degree, log income per capita, share urban, republican vote share (2012 pres. election) – 2012 levels interacted with year FEs. Standard errors clustered by county.

Figure 12: Support for increasing border security: Reduced form effect by quartile of AP-intensity

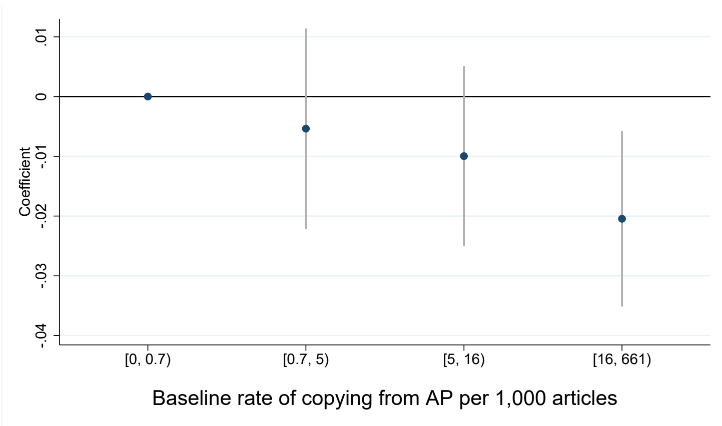

Notes: Point estimates and 95% confidence intervals on the interactions of AP-intensity with survey year, conditional on year and county FEs, respondent controls, and county controls interacted with year FEs. Respondent controls: age, age squared gender, indicators for race, college, and 1st or 2nd generation immigrant. County controls: log population, racial composition, share foreign born, share college degree, log income per capita, share urban, republican vote share (2012 pres. election) – 2012 levels interacted with year FEs. Standard errors clustered by county.

Robustness The above result is robust to alternative specifications, variable and sample construction,

In table 6 I test the robustness of the results to different versions of AP-intensity – using either attribution to AP or plagiarism detection to identify AP-sourced articles, and extending the definition to all articles, instead of ones on immigration. This yields very similar results to the baseline (columns 1 to 4 and 5 to 6). Instead, I find no differential effect of the ban depending on *Reuters*-intensity (column 4). This is reassuring since it suggests that the effect is specific to AP, rather than to the use of news wires in general.

Table 6: Support for increasing border control: Alternative measures of AP-intensity

	Reduced Form				2SLS		
	(1) Border	(2) Border	(3) Border	(4) Border	(5) Border	(6) Border	(7) Border
PostBan × AP-intensity: AP credited	-0.0047***						
	(0.001)						
PostBan × AP-intensity: Plagiarism detection		-0.0037**					
		(0.002)					
PostBan × AP-intensity: AP credited, *all articles*			-0.0028**				
			(0.001)				
PostBan × Reuters-intensity: Reuters credited, *all articles*				-0.0006			
				(0.003)			
'Illegal imm.', pct. of 'Imm.'					0.0057***	0.0047*	0.0050**
					(0.002)	(0.002)	(0.002)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	.	20.11	16.69	14.86
PostBan × AP-intensity: AP credited					-0.8255***		
					(0.184)		
PostBan × AP-intensity: Plagiarism detection						-0.8004***	
						(0.196)	
PostBan × AP-intensity: AP credited, *all articles*							-0.5729***
							(0.149)
Observations	161490	161490	148271	149681	161490	161490	148271
Number of counties	2113	2113	1767	1789	2113	2113	1767
R^2	0.14	0.14	0.14	0.14	0.10	0.10	0.10
Mean dep. var.	0.56	0.56	0.55	0.55	0.56	0.56	0.55

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent and county controls as in Table 5. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

As a further robustness check, in Table 7 I vary the threshold for inclusion of a county in the sample based on the share of the county's circulation covered by the Newslibrary & ProQuest content data (set at $\geq 90\%$ at baseline). Lowering this threshold increases measurement error as it leads to using data from counties for which the estimates of AP-intensity and share "illegal immigrant" articles omit newspapers with larger and larger market shares. Consistent with attenuation bias due to measurement error, the estimated effect of the ban decreases in magnitude as I lower the threshold, but is of the same sign and statistically significant throughout, including in the case of keeping all counties.

Finally, in Appendix Table C7 I show that similar results obtain using a dummy variable for AP-intensity above median instead of a continuous measure, clustering standard errors at the higher level of state instead of county, and restricting the sample to counties with one main newspaper, i.e. only one newspaper accounting for ¿=10% of total county circulation. The latter sample restriction holds for 40% of counties in the baseline sample and allows me

Table 7: Support for increasing border control: Alternative thresholds for share of county circulation covered by content data

	Reduced form				2SLS			
	(1) 1	(2) 0.75	(3) 0.50	(4) 0	(5) 1	(6) 0.75	(7) 0.50	(8) 0
PostBan × AP-intensity	-0.0066*** (0.002)	-0.0042*** (0.002)	-0.0039*** (0.001)	-0.0036** (0.001)				
'Illegal imm.', pct. of 'Imm.'					0.0062*** (0.002)	0.0054** (0.002)	0.0039** (0.002)	0.0035** (0.001)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	.	25.86	16.16	24.73	30.98
First stage coef. on PostBan × AP-intensity					-1.0660*** (0.210)	-0.7843*** (0.195)	-0.9845*** (0.198)	-1.0354*** (0.186)
Observations	98,993	178,584	202,574	240,638	98,993	178,584	202,574	240,638
Number of counties	1,685	2,209	2,320	2,834	1,685	2,209	2,320	2,834
R^2	0.15	0.14	0.14	0.14	0.10	0.10	0.10	0.10
Mean dep. var.	0.56	0.55	0.55	0.55	0.56	0.55	0.55	0.55

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent and county controls as in Table 5. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

to match each respondent to only one newspaper, avoiding aggregation of AP-intensity and content.

Heterogeneity by respondent characteristics In the above results I considered the sample of all CCES respondents. Yet, respondents who regularly read a newspaper are likely more exposed to the treatment. Therefore, in Figure 13 I split the sample into respondents who report that they have not read a newspaper in the past 24 hours, those who report that they have, and those who report that they have read a newspaper in print. This analysis has the caveat that the newspaper readership question was not asked in the 2012 wave, so that sample size and power are reduced. Yet, the results suggests a stronger magnitude of the effect among (self-reported) frequent newspaper readers.

I next examine the heterogeneity of the effects by respondents' political engagement and ideology. In Panel (a) of Figure 14 I split the sample into respondents who have voted in the most recent general election vs those who have not, and in Panel (b) I split the sample into respondents with low vs high (self-reported) interest in politics.[27] I find stronger treatment effects for voters who are not politically active and ones who are less interested in politics, in line with individuals with weak priors being more persuadable. This is also confirmed splitting the sample by a 3-point partisanship scale: in Panel (c) I find the strongest effects in the sample of independents.

Heterogeneity by county characteristics In Figure 15 I examine heterogeneity of the effect by the share of foreign born and share of Hispanic population in the county, splitting counties

[27] The exact wording of the question is as follows: *Some people seem to follow what's going on in government and public affairs most of the time, whether there's an election going on or not. Others aren't that interested. Would you say you follow what's going on in government and public affairs?*

Figure 13: Heterogeneity by newspaper readership

(a)

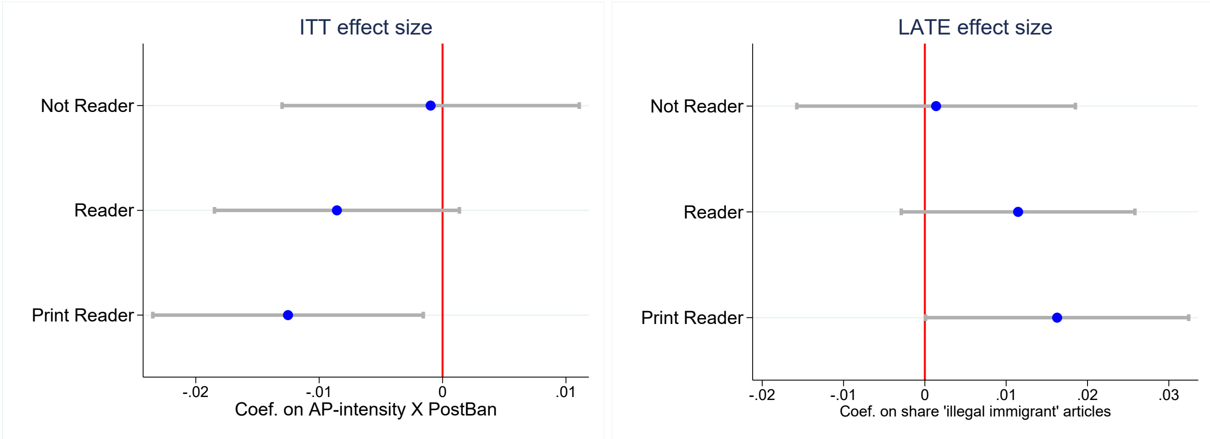

Notes: Each graph presents coefficients and 95% confidence intervals from a regression with support for increasing border security as dependent variable and sample restricted to counties with particular characteristics. In each case the dependent variable is standardized to facilitate comparison of the magnitudes.
Left hand side: Coefficients on the interaction of AP-intensity and PostBan (intention-to-treat). **Right hand side:** Coefficients on share 'illegal immigrant' articles (local average treatment effect).
All regressions control for respondent characteristics, county characteristics interacted with year FEs, county and survey year FEs. Standard errors clustered by county.

at the respective median value for the US. The results point to somewhat stronger effects in places with less immigrants and ones with less Hispanic population – a pattern that can again be interpreted in line with weaker priors for individuals with less direct content to the issue of immigration.

Views on other policies If these results reflect a general change in political leanings that by chance happens to be correlated with AP-intensity, we would expect that support for other policies endorsed by the Republican party is also affected in the same direction. In table 8 I present the results of a placebo exercise that tests for an effect on support for policies related to taxation, abortion, gay rights, and the respondent's assessment of the state of the economy. I find no significant effect of the on any of these outcomes.

Voting Was the change in immigration policy views enough to shift voting choices? The answer appears to be no – in table 9 I show that the ban had no effect on intentions to vote for the Republican candidate in elections for various offices. In table 10 I use electoral data rather than voting intentions reported in CCES, and confirm the null effect for presidential elections and for House midterm elections.

One interpretation of these results is that the effect on voters' views on immigration may not have been large enough to affect voting choices. I do however detect a statistically significant negative effect of the ban on disapproval of President Obama (columns 4 and 8 of table 9). This is in line with the previous results, given Obama's immigration reform agenda.

1.4.4 Magnitudes

In this section I discuss the magnitude of the estimated effect on immigration policy views. Expressed in terms of one standard deviation higher AP-intensity, the estimated treatment effect

Figure 14: Heterogeneity by political interest and participation

(a) Effects by political interest

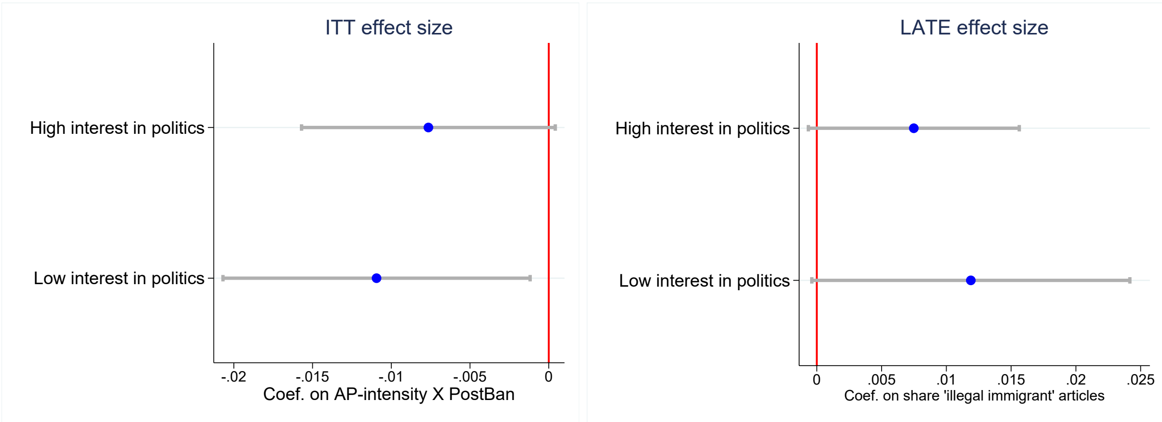

(b) Effects by voting participation

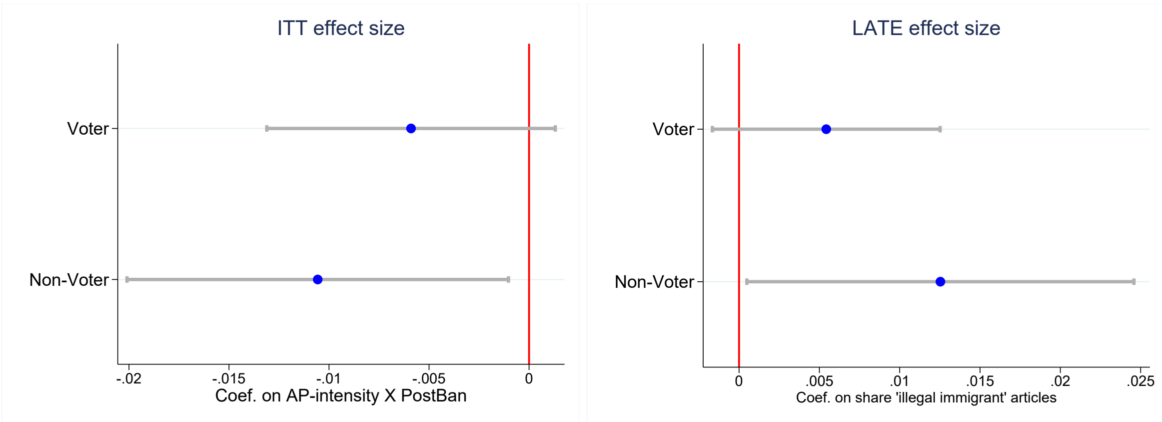

(c) Effects by political affiliation

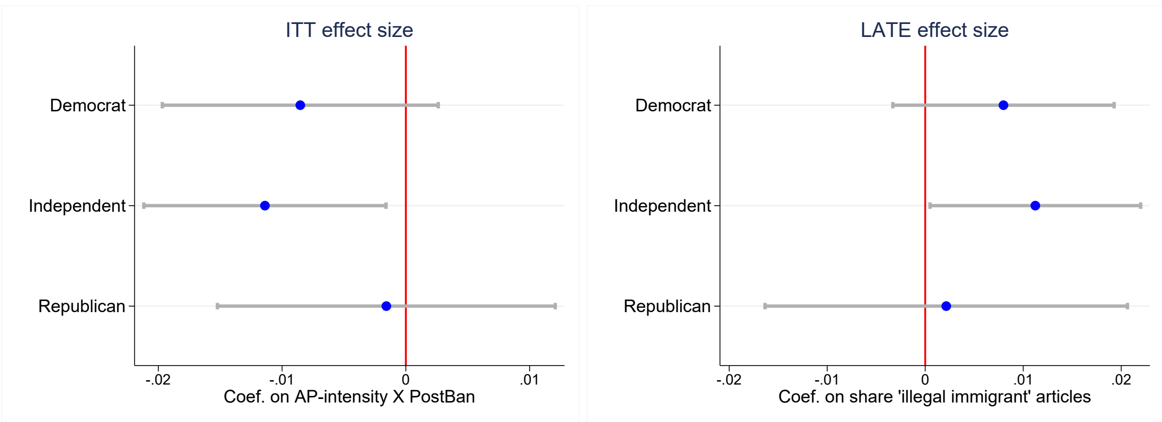

Notes: Each graph presents coefficients and 95% from a regression with support for increasing border security as dependent variable and sample restricted to respondents with particular characteristics. In each case the dependent variable is standardized for the purpose of comparison of the magnitudes. **Left hand side:** Coefficients on the interaction of AP-intensity and PostBan (intention-to-treat). **Right hand side:** Coefficients on share 'illegal immigrant' articles (local average treatment effect). All regressions control for respondent characteristics, county characteristics interacted with year FEs, county and survey year FEs. Standard errors clustered by county.

in the first stage amounts to 9.5% fewer "illegal immigrant" over "immigrant" articles per year. Relative to the mean in the ProQuest/ Newslibrary sample this implies about 9 fewer "illegal

Figure 15: Heterogeneity by share foreign-born / share Hispanic

(a) Effects by share foreign

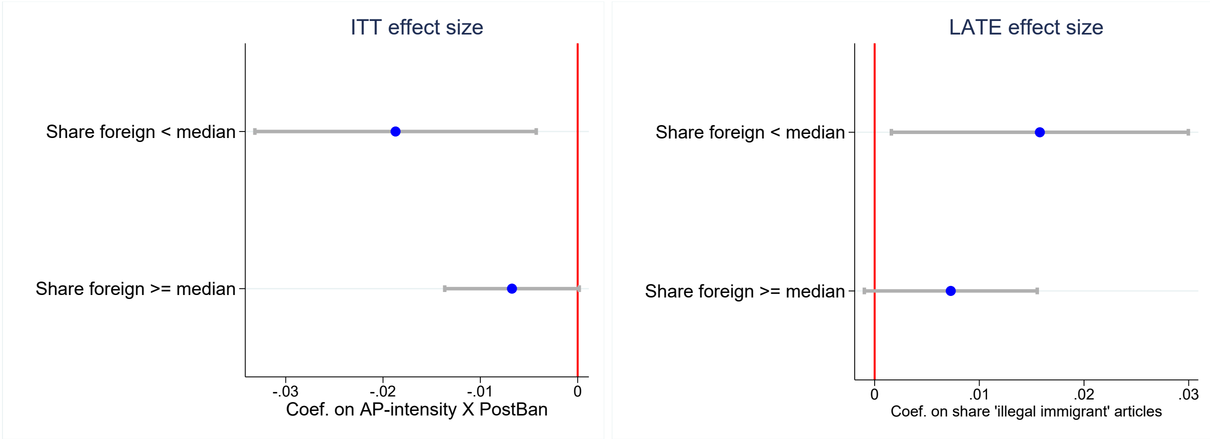

(b) Effects by share Hispanic

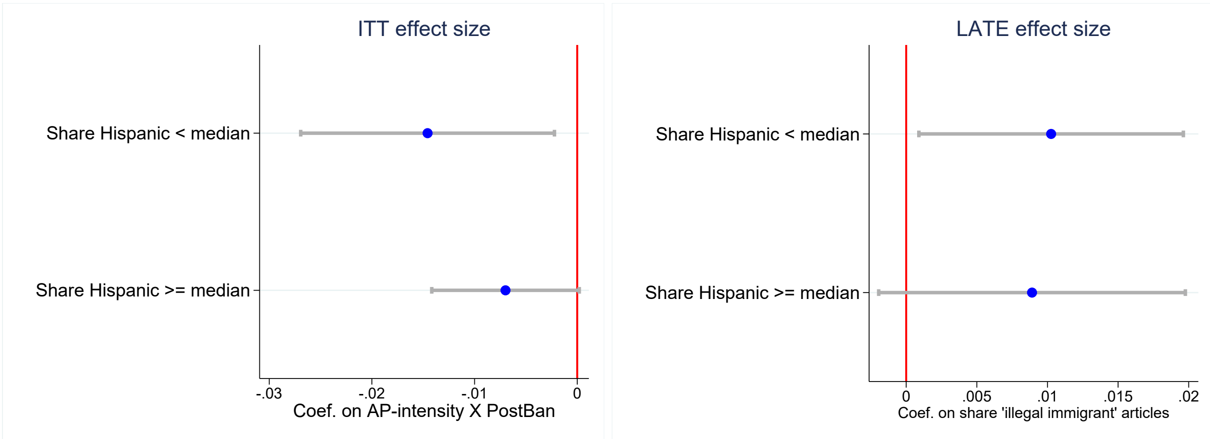

Notes: Each graph presents coefficients and 95% from a regression with support for increasing border security as dependent variable and sample restricted to counties with particular characteristics. In each case the dependent variable is standardized for the purpose of comparison of the magnitudes. **Left hand side:** Coefficients on the interaction of AP-intensity and PostBan (intention-to-treat). **Right hand side:** Coefficients on share 'illegal immigrant' articles (local average treatment effect). All regressions control for respondent characteristics, county characteristics interacted with year FEs, county and survey year FEs. Standard errors clustered by county.

immigrant" articles per year. [28] The corresponding intention to treat effect on support for border security is 0.7 percentage points in the sample of all survey respondents, or 0.9 percentage points in the sample of regular newspaper readers.

One way to benchmark this magnitude is to compare it to the gap between Republican and Democrat respondents in the CCES. Among Republicans, 76.8% support increasing border security, and this is the case for 37.6% of Democrats. Hence, the treatment corresponds to between 1.8% to 2.3% of the gap – a relatively mild effect.

To facilitate comparison to other studies in the media literature, it is also useful to express these magnitudes in terms of persuasion rates. The persuasion rate is defined as the share of people who change their behavior, or in this case – change their survey answer, in response to the treatment, out of the ones who could have potentially done so (DellaVigna and Kaplan 2007; DellaVigna and Gentzkow 2010). The persuasion rate for this treatment, that is, the share

[28]The coverage of these data is not universal, so that this should be taken as a lower bound.

Table 8: Views on other policies

	Reduced Form				2SLS			
	(1) Taxes	(2) Economy	(3) Abortion	(4) Gay marriage	(5) Taxes	(6) Economy	(7) Abortion	(8) Gay marriage
PostBan × AP-intensity	0.0004	0.0011	-0.0016	-0.0024				
	(0.001)	(0.001)	(0.001)	(0.002)				
'Illegal imm.', pct. of 'Imm.'					-0.0004	-0.0012	0.0017	0.0025
					(0.001)	(0.002)	(0.002)	(0.002)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	.	19.63	19.41	19.80	19.62
Observations	159,702	158,797	161,595	161,217	159,702	158,797	161,595	161,217
Number of counties	2,108	2,109	2,110	2,112	2,108	2,109	2,110	2,112
R^2	0.08	0.23	0.20	0.22	0.05	0.16	0.13	0.17
Mean dep. var.	0.45	0.42	0.47	0.42	0.45	0.42	0.47	0.42

Notes: "Taxes"= 1 if would rather cut public spending than increase taxes. "Economy"= 1 if believe the economy has gotten worse over the past year. "Abortion"= 1 if oppose always allowing women to have an abortion as matter of choice. "Gay marriage"= 1 if oppose gay marriage. Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent controls: age, age squared, gender, race, college, 1st or 2nd generation immigrant, and political ideology. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share. Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

of respondents who are dissuaded to support restrictive immigration policy, can be expressed as:

$$f = \frac{db}{de}\frac{1}{1-b_0}, \tag{1.7}$$

where b is support for restring immigration, e is exposure to "illegal immigrant" articles, and b_0 is the share of the population that would oppose restrictive immigration policy in absence of the treatment. With the coefficient estimated for the sample of all respondents, and taking into account that about 1/3 of them report that they read a newspaper and an average of 56% support restrictive immigration policy, this implies a persuasion rate of $f = (0.007)/(0.33 * 1) * (1/0.56) \approx 3.8\%$. With the coefficient estimated from the sample of newspaper readers, the implied persuasion rate is $f = (0.009)/(1 * 1) * (1/0.59) \approx 1.5\%$.[29]

This magnitude is in the lower end of the effects estimates in the media literature, consistent with the milder nature of the treatment compared to other studies. For comparison, Chiang and Knight (2011) estimate a persuasion rate of 6,5% for the effect of a (surprising) newspaper electoral endorsement on voting intentions for that candidate.

Finally, this analysis and the interpretation of the results has focused on print newspapers, as circulation data allows me to map survey respondents to their respective locally read newspapers. However, views on immigration policy are also affected by consumption of TV and Internet outlets, which may also have been affected by the ban. This matters for the interpreta-

[29] Here, as standard in the calculations of persuasion rates in the media literature, I am assuming that a newspaper reader reads every article. Relaxing this assumption, e.g. assuming that only a fraction of articles are actually read, would lead to a higher persuasion rate. On the other hand, as documented in section 1.2, the ban appears to be more salient when it comes to the language used in headlines. Assuming that readers are more likely to pay attention to headlines would therefore lead to a lower persuasion rate.

Table 9: Voting Intentions

	Reduced Form				
	(1) Turnout	(2) President Rep. vote	(3) House Rep. vote	(4) Senate Rep. vote	(5) Obama Disapprov.
PostBan × AP-intensity	-0.0016	-0.0024	-0.0014	-0.0048**	-0.0083***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
Respondent controls	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
Observations	148,182	99,149	119,737	83,560	157,365
Number of counties	2,103	1,996	2,056	1,939	2,109
R^2	0.46	0.46	0.47	0.48	0.52
Mean dep. var.	0.65	0.27	0.41	0.40	2.68

	Reduced Form				
	(1) Turnout	(2) President Rep. vote	(3) House Rep. vote	(4) Senate Rep. vote	(5) Obama Disapprov.
'Illegal imm.', pct. of 'Imm.'	0.0016	0.0018	0.0015	0.0063	0.0086**
	(0.002)	(0.002)	(0.002)	(0.004)	(0.004)
Respondent controls	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	20.40	26.93	19.50	9.27	19.82
Observations	148,182	99,149	119,737	83,560	157,365
Number of counties	2,103	1,996	2,056	1,939	2,109
R^2	0.08	0.32	0.41	0.42	0.47
Mean dep. var.	0.65	0.27	0.41	0.40	2.68

Notes: Intent to vote for Republican candidate in Presidential, House and Senate elections, and disapproval of President Obama. Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent controls: age, age squared, gender, race, college, 1st or 2nd generation immigrant, and political ideology. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

tion of the results to the extent that the AP-intensity of other media consumed in a given county is positively correlated with that of locally circulated newspapers. In that case, the results would be interpreted as a combined media exposure effect, rather than a per-article effect.

1.4.5 Mechanisms

1.4.6 Persuasion vs social signaling

The results discussed above are consistent with a persuasion mechanism in which exposure to the phrase "illegal immigrant" affects readers' intrinsic views on immigration policy. This interpretation is supported by the heterogeneity of the effect by respondents' political engagement, interest and partisanship – the treatment affects the views of individuals who do not have

Table 10: Electoral results

	Reduced Form				2SLS			
	(1) Rep. Share President	(2) Turnout President	(3) Rep. Share House	(4) Turnout House	(5) Rep. Share President	(6) Turnout President	(7) Rep. Share House	(8) Turnout House
PostBan × AP-intensity	0.0008	0.0003	-0.0019	0.0008				
	(0.001)	(0.000)	(0.002)	(0.001)				
'Illegal imm.', pct. of 'Imm.'					-0.0006	-0.0002	0.0021	-0.0009
					(0.001)	(0.000)	(0.002)	(0.001)
Respondent controls	No	No	No	No	No	No	No	No
Year FEs × County controls	No	No	No	No	No	No	No	No
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	.	66.18	66.18	25.26	26.05
Observations	4698	4698	4662	4702	4698	4698	4662	4702
Number of counties	2349	2349	2331	2351	2349	2349	2331	2351
R^2	0.98	0.98	0.86	0.90	-0.04	-0.01	-0.05	-0.04
Mean dep. var.	0.64	0.57	0.65	0.39	0.64	0.57	0.65	0.39

Notes: Republican vote shares and turnout in presidential elections (2012 and 2016) and in House midterm elections (2010 and 2014). Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share. Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

a strong partisan attachment or direct exposure to the issue of immigration and hence likely have a more malleable stance on immigration policy.

An alternative interpretation to that of persuasion may be that the ban served as a signal that expressing anti-immigration views is no longer socially acceptable (), making people less likely to state this response in a survey even if they do intrinsically hold anti-immigration views. This would be a plausible interpretation if the announcement the ban received news coverage correlated with AP-intensity. To test if that is the case, I search "illegal immigrant" articles published in 2013 for mentions of the keywords ("ban" or "Style Guide") and ("AP" or "Associated Press"). I find a total of 37 such articles in my sample of 2200 outlets. The correlation between coverage of the ban and AP-intensity is positive but low (0.08). Furthermore, a large fraction of these articles take a critical stance on AP's decision and are hence unlikely to shift norms in the direction suggested by my results.[30]

1.4.7 Issue vs equivalency framing

The counterfactual to exposure to "illegal immigrant" is a mix consisting of an equivalence frame (substitution from "illegal immigrant" to "[immigrant who] entered the country illgeally") and an issue frame (substitution from "illegal immigrant" to "immigrant"). While the setting of AP's ban does not lend itself to differentiating between these two treatments, two pieces of evidence suggest that the issue frame component is likely to explain a large portion of the observed effect.

First, as discussed in section 1.2, text analysis of AP dispatches suggests that this type of substitution occurred more frequently due to its compactness compared to the equivalency

[30]To name some examples, the Miami Herald, the Arizona Daily Sun and the Chattanooga Times Free Press ran articles with the following respective titles: *"AP should not stop with âillegal immigrantsâ"*, *"AP banning clear thinking"*, and *"Decision to ban 'offensive words' means banning thoughts as well"*.

frame proposed by AP, and that was especially the case in headlines.[31]

Second, survey experimental studies suggest that views on immigrants and immigration policy are generally reactive to issue frames (e.g. highlighting the aspect of crime or not), while the evidence on equivalency frames (e.g. ”undocumented” vs ”illegal immigrant”) is more mixed (Merolla and Ramakrishnan 2016). It therefore seems plausible that the issue-frame component plays an important role in explaining the observed effects.

1.5 Conclusion

This paper has documented a large degree of diffusion of language from news wires to media outlets. Changes in news wire language rules, which are determined centrally rather than in consideration of the political leanings of the owners or readers of a particular media outlet, are therefore a useful source of variation to estimate the effects of media slanted language on readers.

Applying this strategy, I find evidence consistent with exposure to the term “illegal immigrant” in local media shifting preferences towards more restrictive immigration policy. This effect is driven by passive readers with less pronounced political views and lower direct exposure to immigration, consistent with a persuasion mechanism.

This evidence provides proof of concept for the hypothesis that ideologically slanted language can have a persuasive impact. Yet, this evidence is limited to the setting of unauthorized immigration and to exposure to one particular term. More work is needed to understand the external validity of this mechanism of media influence.

[31] This result highlights a potentially important aspect of politically slanted language – that frames pushed through short and catchy phrases are more readily adopted by the media compared to nuanced narratives.

Appendix B: Background

AP Style Guide entry on "illegal immigrant"

Pre-Ban

illegal immigrant Used to describe those who have entered the country illegally, it is the preferred term, rather than *illegal alien* or *undocumented worker*.
Do not use the shortened term *illegals*.

Pre-Ban

illegal immigration Entering in a country in violation of civil or criminal law. Except in direct quotes essential to the story, use *illegal* only to refer to an action, not a person: *illegal immigration*, but not *illegal immigrant*. Acceptable variations include *living in* or *entering a county illegally* or *without legal permission*.
Expect in direct quotations, do not use the terms *illegal alien, an illegal, illegals* or *undocumented*.
Do not describe people as violating immigration laws without attribution.
Specify wherever possible how someone entered the country illegally and from where. Crossed the border? Overstayed a visa? What nationality?
People who were brought into the county as children should not be described as having immigrated illegally. For people guaranteed a temporary right to remain in the U.S. under the Deferred Action for Childhood Arrivals program, use *temporary resident status*, with details on the program lower in the story.

Figure A1: AP's Style Guide embedded in a text editor

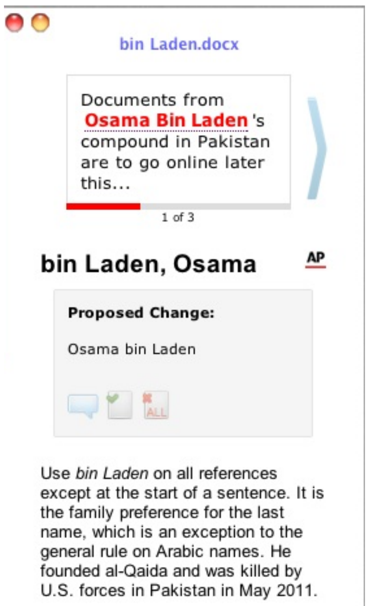

Figure A2: The ban reported in the Atlantic

Atlantic   Popular   Latest

The AP's Ban on 'Illegal Immigrant' Will Change How We Talk About Immigration

That faint sound you hear is Senate reporters from the AP, *The New York Times*, and beyond smacking their delete keys, rethinking their agenda setting aloud, and figuring out how we talk now, amidst a serious legislative discussion

ALEXANDER ABAD-SANTOS | APR 2, 2013 | POLITICS

Examples of AP dispatches before and after the ban

<u>Pre-Ban</u> (18-Mar-2013): *Senate panel OKs letting non-citizens, including illegal immigrants, get driver's licenses*

ST. PAUL, Minn. (AP) â- Bills that would let illegal immigrants get a Minnesota driver's license are moving forward at the Capitol. The Senate Transportation and Public Safety Committee on Monday passed a bill to ease restrictions on driver's licenses for non-U.S. citizens. A House committee endorsed a similar bill last week. Sen. Bobby Joe Champion, a Minneapolis Democrat, says his bill would make Minnesota roads safer by funneling more drivers through the state's driving test and making it easier for them to buy automobile insurance. Republicans say the change could lead to unintended consequences, like illegal immigrants using state IDs to vote. The bill passed 10-7, with all Democrats in favor and all Republicans voting against it.

<u>Pre-Ban</u> (16-Apr-2013): *Immigrant driver's license bill takes step forward in Oregon Senate committee work*

SALEM, Ore. (AP) – An Oregon Senate committee has advanced a bill granting four-year driver's licenses to people who can't prove they're legally in the United States. The Senate Business and Transportation Committee approved the measure Monday on a 4-2 vote. The bill would allow immigrants who have lived in Oregon for at least a year and meet other requirements to apply for driver's cards without proving legal presence. The card would be valid for only four years– half as long as a standard Oregon license– and would state "driving privilege only." Supporters say it will make Oregon roads safer because there would be fewer untrained and uninsured drivers, but opponents say it could create a culture of crime in the state. The bill goes to a legislative budget committee.

Appendix B: Data

Computation of immigration slant

In this section I describe the procedure for computing an index for the immigration-specific slant of AP dispatches released in each quarter, and that of the articles published by each news outlet in a given year. This follows closely the method developed by Gentzkow and Shapiro (2010a).

AP's slant over time I start off with the set of all Congressional speeches for the period 2009-2012 (i.e., before AP's ban) that mention the word "immigrant".[32] First, after pre-processing the text (removing stop-words and lemmatising), I rank the 500 bi-grams that are most predictive of the speakers' party based on the Pearson's χ^2 statistic. I keep the ones encountered in the similarly pre-processed corpus of AP dispatches at least 10 times – this results into 363 phrases. Table **??** lists the phrases with highest χ^2 that are encountered more often in Republican vs Democrat speech respectively, i.e. the phrases with highest partisanship. In one version of the slant measure, in this step I further exclude the phrase "illegal immmigrant" and its substitutes "country illegally" / "border illegally".

For each phrase p and for each congressperson c I compute the relative frequency of the phrase in the congressperson's speech as $\tilde{f_{pc}} = f_{pc}/\sum_p f_{pc}$. I then regress the relative frequency of the phrase by congressperson on a continuous measure of the congressperson's ideology. Specifically, I measure ideology as the first dimension of the DW-nominate score provided by *Voteview* – a widely used index of ideology derived from roll-call voting. I obtain phrase-specific intercept and slope coefficients a_p and b_p.

Finally, I compute the relative frequency of each phrase in AP dispatches released in a given quarter – $\tilde{f_{pq}}$ – and regress $(\tilde{f_{pq}} - a_p)$ on b_p. The resulting slope coefficient is the quarter-specific measure of slant.

As a within-sample validation of this measure, I also compute the analogous index by congressperson and correlate it with true ideology. The correlation is 0.59 for the baseline version of slant, and 0.62 for the version excluding "illegal immigrant" and its substitutes. Taking the square of these coefficients, this implies that respectively 34% and 38% of the variation in these measures is due to variation in ideology, with the rest due to noise.

Slant of news outlets over time To compute immigration-specific slant by newspaper and year I follow exactly the procedure outlined above with three modifications. First, I replace the AP dispatch corpus with a corpus containing the text (headline + first paragraph) of each news article mentioning the word "immigrant". Second, given the high volume of data I set the threshold for a phrase's occurrence in the corpus to 50, which results into 338 phrases. Third, in the final set of regressions I operate at the level of news outlet $\times$ year rather than by quarter.[33]

Plagiarism detection algorithm

In this section I describe the algorithm I use to identify "immigrant" articles that are copied from AP but do not necessarily credit AP.

[32] I focus on the pre-ban period in light of the result that the ban may have affected the language used in Congressional speech. On the other hand, this has the disadvantage of possibly omitting phrases that emerged after 2013. I obtain similar results taking Congressional speech for the entire sample period.

[33] For some newspaper-years none of the selected phrases is contained in the news article's text – in this case the slant index is set to missing.

The first step of the algorithm is to assign to each article a set of AP dispatches that could potentially have been used in the writing of the article. I focus on AP dispatches released in the day before publication and mentioning the word “immigrant”.[34] This is a simplified version of the procedure used in Cage et al. (2020), which first clusters articles by the event they cover, and then forms the set of potentially plagiarized articles as those that cover the same event and are published prior to the article of interest. The second step in the algorithm is to compute a measure of verbatim copying. I pre-process all texts by removing punctuation and stop-words, stemming, and tokenizing into 5-grams. I then measure the share of the article’s text that is identical to each paired dispatch and take the maximum over all paired dispatches. I label an article as copied from AP if the maximum text overlap exceeds 20% (equivalent to 70 characters, relative to the mean text length of 350).

Figure B1: AP-copy rate: Attribution and plagiarism

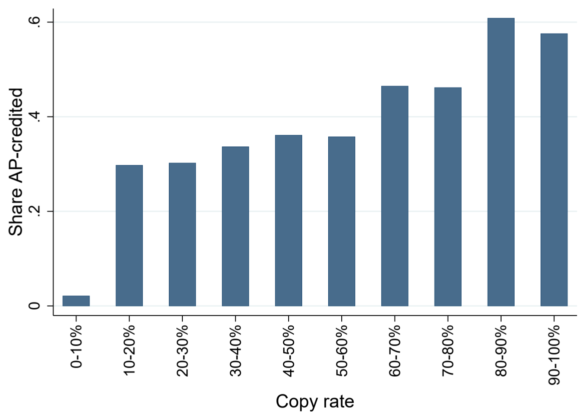

Figure B1 presents the relationship between copying and crediting AP, plotting the average share of credited articles by bin of the copy-rate distribution (i.e. by share of text overlapping with an AP dispatch). It is notable that even among articles whose lead paragraph is virtually identical to an AP dispatch (with 90-100% identical text), the rate of crediting AP never exceeds 60%. In other words, relying on attribution to AP alone would have missed a substantial volume of copied articles. When collapsed at the media outlet level however, the correlation between the two measures is 0.83.

Immigration questions in the CCES

What do you think the U.S. government should do about immigration? Select all that apply.

- Fine US businesses that hire illegal immigrants.
  (-07, -12, -14, -17)

- Grant legal status to all illegal immigrants who have held jobs and paid taxes for at least 3 years, and not been convicted of any felony crimes.
  (-07, -10, -11, -12, -14, -16, -17)

[34] I do not use contemporaneous (same-day) AP-dispatches because the origin of the content is more ambiguous in this case – text similarity could be due the media outlet copying AP, or to AP redistributing content produced by a member outlet.

- Increase the number of border patrol on the US-Mexican border.
  (-07, -10, -11, -12, -14, -16, -17)
- Build a wall between the US and Mexico.
  (-07, -17)
- Allow police to question anyone they think may be in the country illegally.
  (-10, -11, -12, -14, -17)
- Prohibit illegal immigrants from using emergency hospital care and public schools.
  (-12)
- Deny automatic citizenship to American-born children of illegal immigrants.
  (-12)
- Identify and deport illegal immigrants.
  (-14, -16, -17)
- Grant legal status to people who were brought to the US illegally as children, but who have graduated from a U.S. high school.
  (-16)

———————————————————————-

Descriptive statistics

Table B1: Summary statistics: Main sample (all media outlets)

Variable	Mean	Std. Dev.	Min.	Max.	N
Immigrant/ Total articles (pct)	0.809	2.179	0	100	232,729
Illigal immigrant / Immigrant articles (pct)	18.584	28.628	0	100	143,696
Illigal immigrant / Immigrant articles (pct) – Pre-Ban	24.761	32.017	0	100	64,561
Illigal immigrant / Immigrant articles (pct) – Post-Ban	13.546	24.4	0	100	79,135
Immigrant & AP-sourced / Immigrant (pct)	2.115	9.918	0	100	143,821
AP-intensity (asinh)	0.886	1.881	0	7.331	216,709
1[AP-intensity $>$ 0]	0.195	0.397	0	1	216,709
AP intensity, plagiarised (asinh)	0.558	1.541	0	7.331	216,709
AP intensity	0.668	1.62	0	6.908	216,709
AP-int, credited – all articles (asinh)	0.875	1.856	0	7.399	213,984
Reuters-int, credited – all articles (asinh)	0.359	0.991	0	7.154	239,328

Appendix C: Additional Results

Text analysis of AP dispatches

Figure B2: Histogram of AP-intensity, in # articles per 1000

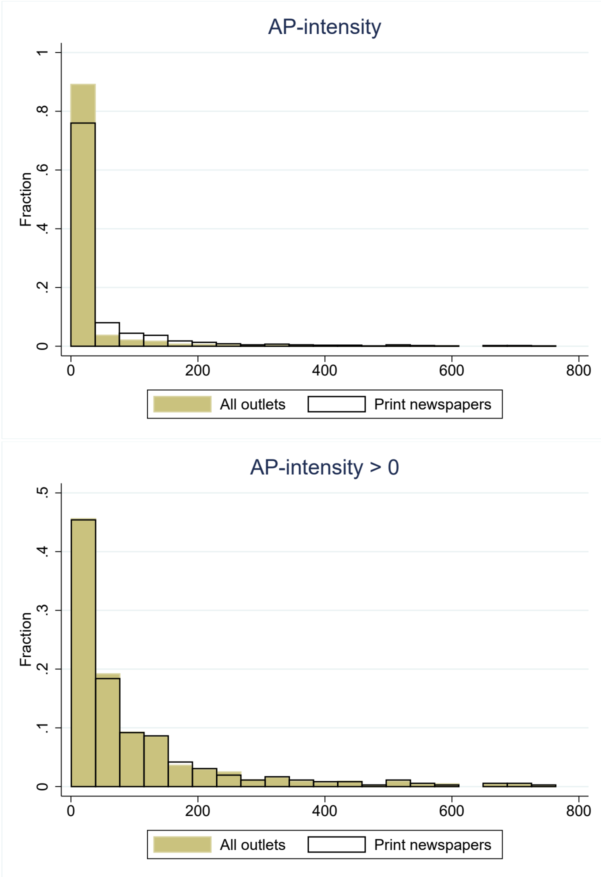

Figure B3: Percent of immigrant articles sourced from AP over time

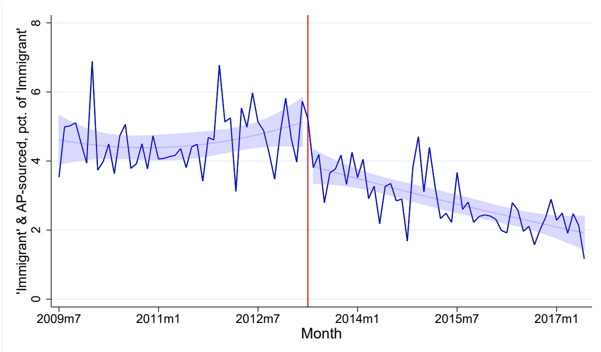

Notes: Monthly number of “immigrant” articles sourced from AP, as percent of all “immigrant” articles.

Figure B4: Histogram of the difference between post-ban and pre-ban share of "illegal immigrant" articles by outlet

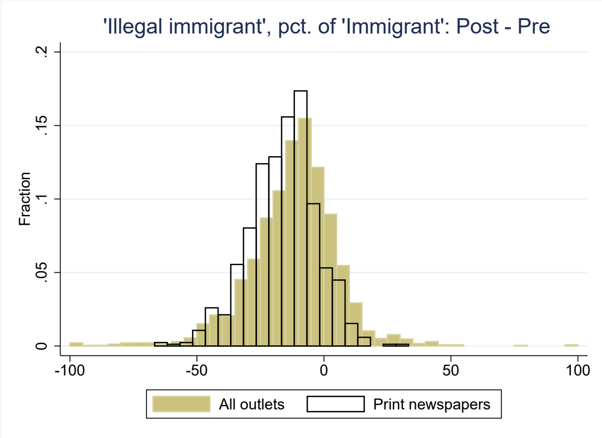

Notes:

Table B2: Summary statistics: Restricted sample of print newspapers

Variable	Mean	Std. Dev.	Min.	Max.	N
Immigrant/ Total articles (pct)	0.61	0.984	0	100	73,041
Illigal immigrant / Immigrant articles (pct)	22.499	27.418	0	100	65,593
Illigal immigrant / Immigrant articles (pct) – Pre-Ban	30.261	30.423	0	100	29,474
Illigal immigrant / Immigrant articles (pct) – Post-Ban	16.165	22.818	0	100	36,119
Immigrant & AP-sourced / Immigrant (pct)	3.725	12.595	0	100	65,666
AP-intensity (asinh)	1.968	2.407	0	7.243	77,204
1[AP-intensity $>$ 0]	0.431	0.495	0	1	77,204
AP intensity, plagiarised (asinh)	1.192	2.083	0	7.140	77,204
AP intensity, credited (asinh)	1.584	2.193	0	6.908	77,204
AP-int, credited – all articles (asinh)	2	2.37	0	7.399	71,040
Reuters-int, credited – all articles (asinh)	0.455	0.942	0	6.568	73,248

Table B3: Summary statistics: CCES sample (respondent-level)

Variable	Mean	Std. Dev.	Min.	Max.	N
Restrict Imm. Index	0.016	0.729	-1.032	1.112	186,252
Border	0.548	0.498	0	1	186,252
Amnesty	0.524	0.499	0	1	186,252
Question	0.399	0.49	0	1	137,569
Don't hire	0.611	0.488	0	1	92,492
Illigal immigrant / Immigrant articles (pct)	21.303	12.265	0	96.727	180,749
Illigal immigrant / Immigrant articles (pct) – Pre-Ban	29.388	11.245	0	96.439	87,562
Illigal immigrant / Immigrant articles (pct) – Post-Ban	13.707	7.337	0	96.727	93,187
AP-intensity (asinh)	3.062	1.538	0	7.232	187,380
AP-int, credited (asinh)	2.303	1.703	0	7.128	187,380
AP-int, plagiarised (asinh)	2.51	1.498	0	6.814	187,380
AP-int, credited – all articles (asinh)	2.431	2.153	0	7.339	172,007
Reuters-int, credited – all articles (asinh)	0.917	1.193	0	5.916	173,726

Table B4: Summary statistics: CCES county-level sample

Variable	Mean	Std. Dev.	Min.	Max.	N
\ shortstack{Restrict Imm.\ Index}	0.133	0.463	-1.032	1.112	12,057
Border	0.588	0.313	0	1	12,057
Amnesty	0.59	0.314	0	1	12,057
Question	0.465	0.32	0	1	8,936
Don't hire	0.65	0.311	0	1	6,537
Illigal immigrant / Immigrant articles (pct)	23.388	13.532	0	96.727	10,919
Illigal immigrant / Immigrant articles (pct) – Pre-Ban	31.664	12.352	0	96.439	5,511
Illigal immigrant / Immigrant articles (pct) – Post-Ban	14.956	8.565	0	96.727	5,408
AP-intensity (asinh)	3.185	1.612	0	7.232	12,112
AP-int, credited (asinh)	2.287	1.766	0	7.128	12,112
AP-int, plagiarised (asinh)	2.717	1.565	0	6.814	12,112
AP-int, credited – all articles (asinh)	2.149	2.261	0	7.339	10,347
Reuters-int, credited – all articles (asinh)	0.589	0.902	0	5.916	10,495

Figure C1: Topics of "immigrant" AP dispatches

(a) Topic 1 "Enforcement"

(b) Topic 2 "Legislation"

(c) Topic 3 "Integration"

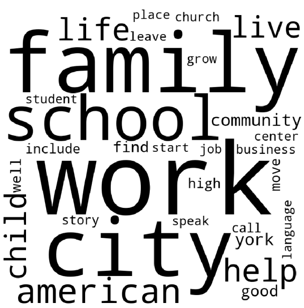

(d) Topic 4 "Europe"

(e) Topic 5 "Elections"

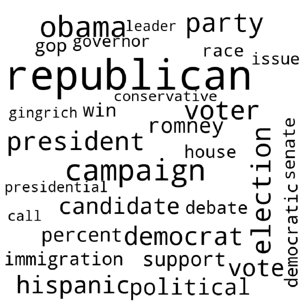

Notes: Word-clouds describing the 5 topics obtained with a Latent Dirichlet allocation (LDA) model applied to the corpus of AP dispatches mentioning the word "immigrant". The corpus excludes the phrase "illegal immigrant" and its synonyms. The number of topics is chosen for interpretability.

Figure C2: Topics distribution pre- and post-ban

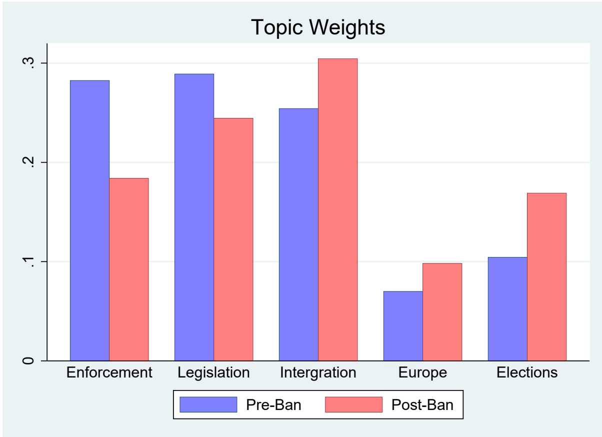

Notes: Distribution of topic weights before vs after the ban.

Figure C3: Correlates of the word "immigrant" before and after the ban; Estimated separately by topic

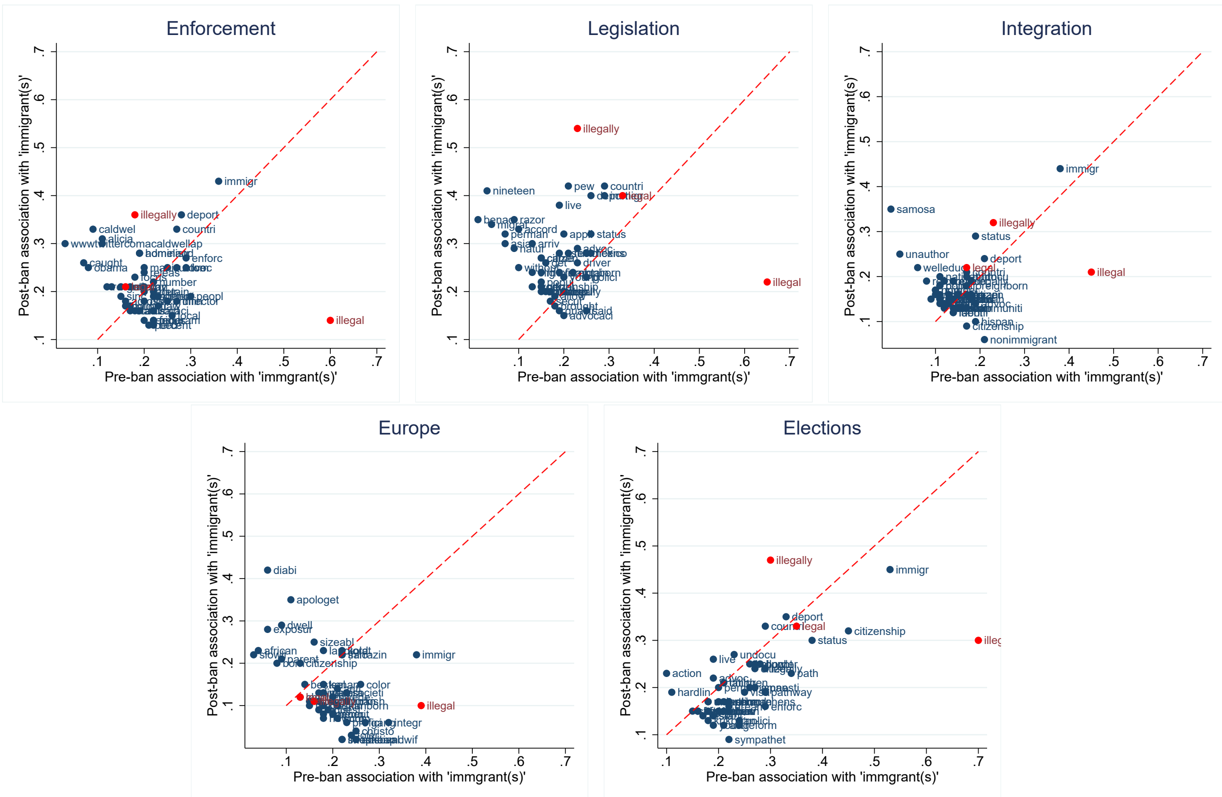

Notes: Top 50 unigrams with highest correlation with the word "immigrant", before and after the ban. Correlations defined based on rate of occurrence within the same article. Derivatives of 'immigr' and 'illeg' are not stemmed for illustration purposes. Estimated separately for each of 5 topics derived from an LDA topic model.

Figure C4: Phrases most predictive of post-ban publishing date; Estimated separately by topic

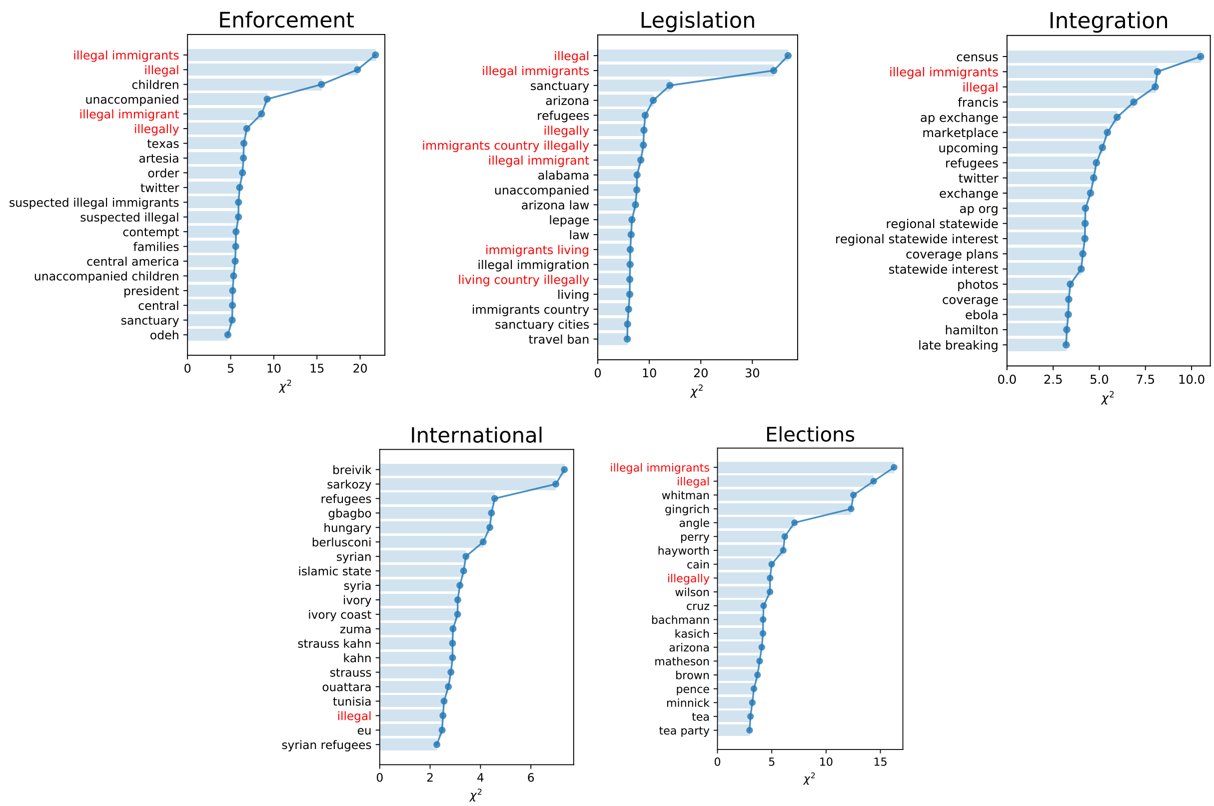

Notes: Top 20 n-grams (n $\in$ 1,2,3) in "immigrant" dispatches that are most predictive of a post-ban publishing date based on a χ^2 test-statistic. Estimated separately for each of 5 topics derived from an LDA topic model.

Figure C5: Slant Index for AP's immigration-related dispatches by quarter; Estimated separately by

(a) Enforcement

(b) Legislation

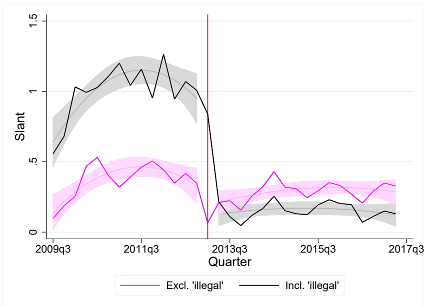

(c) Integration

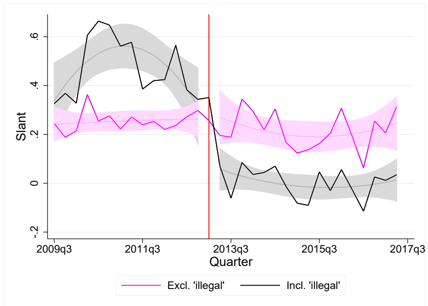

(d) International

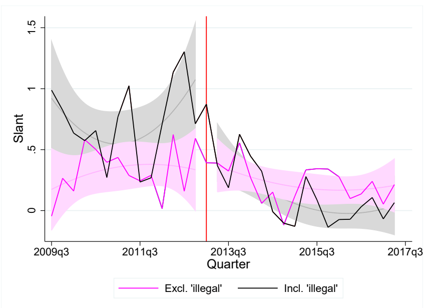

(e) Elections

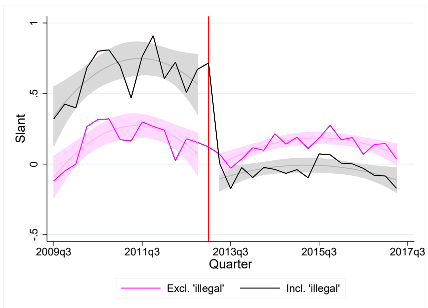

Notes: Immigration-specific slant of AP-dispatches over time. Higher values indicate more right-leaning slant. **Grey line:** Baseline measure of slant. **Magenta line:** slant computed excluding phrases containing the term "illegal immigrant" or substitutes. Estimated separately for each of 5 topics derived from an LDA topic model.

Diffusion

Figure C6: Share of “illegal immigrant” articles in AP-intensive vs non AP-intensive outlets

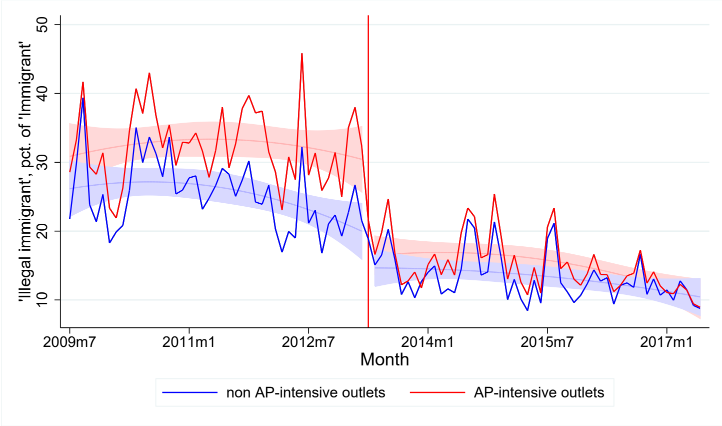

Figure C7: Effects of the ban on newspaper readership

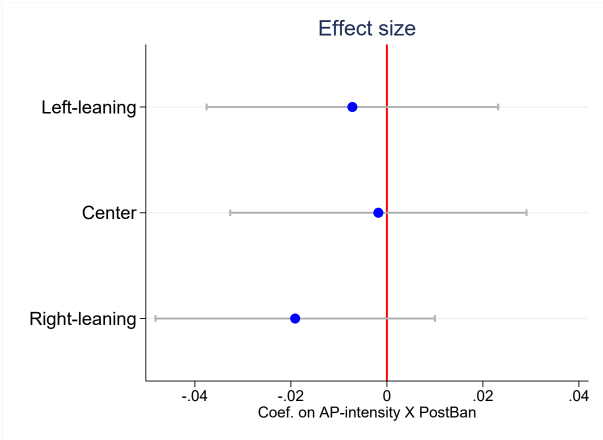

Left-leaning, **Center**, **Right-leaning** denote the sample of outlets in the 1st, 2nd and 3rd tercile of the distribution of the Gentzkow and Shapiro (2010a) slant index respectively.
The graphs present coefficients and 95% confidence intervals on the interaction of AP-intensity and PostBan from regressions restricted to one of the 3 samples at a time. The dependent variable is log circulation. Each regression controls for outlet and year FEs and is weighted by number of “immigrant” articles. Standard errors clustered by outlet.

Table C1: Alternative specifications

	Not normalized	Unweighted	Word-count	Headlines	AP dummy	Elasticity
	(1)	(2)	(3)	(4)	(5)	(6)
	Log(1 + 'Illegal Immigrant')	'Illigal immigrant', pct. of 'Immigrant'				
PostBan $\times$ AP intensity	-0.058***	-1.607***	-1.755***	-1.071***		
	(0.005)	(0.124)	(0.209)	(0.285)		
PostBan $\times$ I[AP-int $>$ 0]					-5.541***	
					(1.059)	
$(Illimm/Imm)_{AP}$ $\times$ AP intensity						0.121***
						(0.022)
Outlet FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	216,709	133,347	124,232	18,976	133,347	131,920
Number of outlets	2271	2269	2160	1414	2269	2269
R^2	0.56	0.21	0.36	0.24	0.42	0.42
Mean dep. var.	0.34	19.52	19.17	14.46	19.52	20.93

Notes: Replication of column (3) of table 1 with the following modifications: (1) Replacing the dependent variable with the log of 1 + number of "illegal immigrant" articles and dropping weights; (2) Regression without weights; (3) Replacing number of articles with word-count; (4) Replacing articles with number of headlines; (5) Replacing continuous AP-intensity with a dummy for positive AP-intensity; (6) Replacing *PostBan* with the time-series of "illegal immigrant" articles (normalized by "immigrant" articles) released monthly by AP. Standard errors clustered by outlet.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C2: AP-sourced vs. original articles

	(1) **AP-credited**	(2) **AP-plagiarised**	(3) **not AP-sourced**
PostBan $\times$ AP intensity	-1.033***	-0.200***	-0.276
	(0.124)	(0.024)	(0.171)
Outlet FEs	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes
Observations	133,469	133,468	133,376
Number of outlets	2269	2269	2269
R^2	0.41	0.11	0.39
Mean dep. var.	0.78	0.35	18.81

Notes: WLS weighted by number of number of "immigrant" articles. Standard errors clustered by outlet.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C3: Heterogeneity by ideology

	Slant < p33 (Dem)	p33 ≤ Slant < p66 (Indep)	Slant ≥ p66 (Rep)
	(1)	(2)	(3)
		'Illegal immigrant', pct. of 'Immigrant'	
PostBan × AP intensity	-1.817***	-0.800*	-1.258**
	(0.443)	(0.482)	(0.491)
Outlet FEs	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes
Observations	9,705	9,886	9,608
Number of outlets	106	110	113
R^2	0.55	0.48	0.50
Mean dep. var.	17.55	22.41	25.22

Notes: WLS weighted by number of "immigrant" articles. Standard errors clustered by outlet.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C4: Heterogeneity by ideology: AP-sourced vs original articles

	Slant < p33 (Dem)	p33 ≤ Slant < p66 (Indep)	Slant ≥ p66 (Rep)
	(1)	(2)	(3)
		'Illegal immigrant' & AP-sourced, pct. of 'Immigrant'	
PostBan × AP intensity	-1.140***	-1.773***	-1.939***
	(0.178)	(0.332)	(0.447)
Outlet FEs	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes
Observations	9,718	9,894	9,619
Number of outlets	106	110	113
R^2	0.47	0.52	0.48
Mean dep. var.	0.98	2.35	1.67
	Slant < p33 (Dem)	p33 ≤ Slant < p66 (Indep)	Slant ≥ p66 (Rep)
	(1)	(2)	(3)
		'Illegal immigrant' & not AP-sourced, pct. of 'Immigrant'	
PostBan × AP intensity	-0.825**	0.579	0.484
	(0.393)	(0.432)	(0.473)
Outlet FEs	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes
Observations	9,711	9,890	9,612
Number of outlets	106	110	113
R^2	0.51	0.41	0.49
Mean dep. var.	15.35	18.95	22.62

Notes: WLS weighted by number of "immigrant" articles. Standard errors clustered by outlet.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C5: Synonyms of “illegal immigrant” and volume of immigration coverage

	(1) AP-approved synonyms pct. of ’Immigrant’	(2) ’Undocumented immigrant’ pct. of ’Immigrant’	(3) ’Immigrant’ pct. of total articles	(4) ’Immigration’ pct. of total articles
PostBan $\times$ AP intensity	0.317***	0.001	-0.002	0.002
	(0.061)	(0.126)	(0.006)	(0.005)
Outlet FEs	Yes	Yes	Yes	Yes
Year-Month FEs	Yes	Yes	Yes	Yes
Observations	133,188	133,330	204,175	204,180
Number of outlets	2269	2269	2162	2162
R^2	0.20	0.34	0.55	0.48
Mean dep. var.	5.06	8.74	0.62	0.51

Notes: WLS weighted by number of ”immigrant” articles in column (1), and by total articles in columns (2) and (3). Standard errors clustered by outlet. AP-approved synonyms are ”living in the country illegally/ without legal permission”, ”enter(-ing/-ed) the country illegally/ without legal permission”. Standard errors clustered by outlet.

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Views on immigration policy

Table C6: OLS correlations between slant and policy views

	OLS correlations				
	(1) Index Restict Imm.	(2) Border	(3) Amnesty	(4) Don't hire	(5) Question
'Illegal imm.', pct. of 'Imm.'	0.0050***	0.0028***	0.0037***	0.0018***	0.0034***
	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Observations	169,545	169,545	169,545	75,413	126,422
Number of counties	2,251	2,251	2,251	2,128	2,217
R^2	0.00	0.00	0.01	0.00	0.01
Mean dep. var.	0.01	0.56	0.53	0.62	0.41

	OLS correlations				
	(1) Index Restict Imm.	(2) Border	(3) Amnesty	(4) Don't hire	(5) Question
'Illegal imm.', pct. of 'Imm.'	0.00053	-0.00004	0.00005	0.00019	-0.00007
	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
Respondent controls	Yes	Yes	Yes	Yes	Yes
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
Observations	162,456	162,456	162,456	74,705	119,552
Number of counties	2,113	2,113	2,113	1,924	2,066
R^2	0.27	0.14	0.16	0.13	0.22
Mean dep. var.	0.01	0.56	0.52	0.62	0.41

Notes: OLS regressions of policy attitudes on the share of locally circulated "illegal immigrant" articles as percent of "immigrant" articles.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C7: Alternative Specifications

	Reduced form			2SLS		
	(1) APint $\geq$ median	(2) Clustered by state	(3) 1-newspaper counties	(4) APint $\geq$ median	(5) Clustered by state	(6) 1-newspaper counties
PostBan $\times$ I[AP-int > median]	-0.0133**					
	(0.005)					
PostBan $\times$ AP-intensity		-0.0046**	-0.0053***			
		(0.002)	(0.002)			
'Illegal imm.', pct. of 'Imm.'				0.0049**	0.0048*	0.0050**
				(0.002)	(0.003)	(0.002)
Respondent controls	Yes	Yes	Yes	Yes	Yes	No
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	11.78	9.27	7.77
Observations	168,409	168,409	87,880	168,409	168,409	87,880
Number of counties	2,125	2,125	841	2,125	2,125	841
R^2	0.08	0.08	0.08	0.04	0.04	-0.00
Mean dep. var.	0.56	0.56	0.56	0.56	0.56	0.56

Notes: OLS reduced form regressions in the left hand side panel, 2SLS regressions in the right hand side panel. Respondent controls and county controls as in Table 5.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C8: Heterogeneity by newspaper readership

	Reduced Form			2SLS		
	(1) Not Reader	(2) Reader	(3) Print Reader	(4) Not Reader	(5) Reader	(6) Print Reader
PostBan $\times$ AP-intensity	-0.0005	-0.0042*	-0.0062**			
	(0.003)	(0.003)	(0.003)			
'Illegal imm.', pct. of 'Imm.'				0.0007	0.0057	0.0080**
				(0.004)	(0.004)	(0.004)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes	Yes	Yes
County and Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	9.52	10.12	10.04
Observations	58174	66777	40192	58174	66777	40192
Number of counties	1844	1805	1596	1844	1805	1596
R^2	0.15	0.16	0.17	0.10	0.11	0.10
Mean dep. var.	0.54	0.56	0.59	0.54	0.56	0.59

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Reader = 1 if read newspaper in the past 24 hours. Print reader = 1 if read print newspaper in the past 24 hours.
Respondent controls and county controls as in Table 5.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C9: Heterogeneity: Interest in Politics and Voting

	Reduced Form		2SLS	
	(1) Voter	(2) Non-voter	(3) Voter	(4) Non-voter
PostBan × AP-intensity	-0.0029	-0.0053**		
	(0.002)	(0.002)		
'Illegal imm.', pct. of 'Imm.'			0.0027	0.0063**
			(0.002)	(0.003)
Respondent controls	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes
County and Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	23.46	14.31
Observations	85618	63056	85618	63056
Number of counties	1950	1830	1950	1830
R^2	0.19	0.12	0.14	0.07
Mean dep. var.	0.58	0.52	0.58	0.52

	Reduced Form		2SLS	
	(1) High interest	(2) Low interest	(3) High interest	(4) Low interest
PostBan × AP-intensity	-0.0037*	-0.0055**		
	(0.002)	(0.002)		
'Illegal imm.', pct. of 'Imm.'			0.0036*	0.0059*
			(0.002)	(0.003)
Respondent controls	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes
County and Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	21.61	16.98
Observations	86443	71617	86443	71617
Number of counties	1939	1875	1939	1875
R^2	0.20	0.10	0.15	0.05
Mean dep. var.	0.61	0.50	0.61	0.50

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent controls and county levels as in Table 5.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C10: Heterogeneity by ideology

	Reduced Form			2SLS		
	(1) Dem.	(2) Indep.	(3) Rep.	(4) Dem.	(5) Indep.	(6) Rep.
PostBan × AP-intensity	-0.0042	-0.0056**	-0.0007			
	(0.003)	(0.002)	(0.003)			
'Illegal imm.', pct. of 'Imm.'				0.0039	0.0056**	0.0009
				(0.003)	(0.003)	(0.004)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs × County controls	Yes	Yes	Yes	Yes	Yes	Yes
County and Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	.	17.93	22.39	11.92
Observations	60047	59697	42101	60047	59697	42101
Number of counties	1679	1815	1744	1679	1815	1744
R^2	0.07	0.11	0.08	0.02	0.05	0.03
Mean dep. var.	0.39	0.57	0.77	0.39	0.57	0.77

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent and county controls as in Table 5.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C11: Heterogeneity by share foreign born and share Hispanic

	Reduced Form			2SLS
	(1) Share Imm. $<$ median	(2) Share Imm. $\geq$ median	(3) Share Imm. $<$ median	(4) Share Imm. $\geq$ median
PostBan $\times$ AP-intensity	-0.0091**	-0.0034*		
	(0.004)	(0.002)		
'Illegal imm.', pct. of 'Imm.'			0.0077**	0.0036*
			(0.004)	(0.002)
Respondent controls	Yes	Yes	Yes	Yes
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes
County and Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	19.76	13.55
Observations	19498	142958	19498	142958
Number of counties	1052	1061	1052	1061
R^2	0.15	0.14	0.07	0.11
Mean dep. var.	0.61	0.55	0.61	0.55

	Reduced Form			2SLS
	(1) Share Hisp. $<$ median	(2) Share Hisp. $\geq$ median	(3) Share Hisp. $<$ median	(4) Share Hisp. $\geq$ median
PostBan $\times$ AP-intensity	-0.0072**	-0.0035*		
	(0.003)	(0.002)		
'Illegal imm.', pct. of 'Imm.'			0.0050**	0.0044
			(0.002)	(0.003)
Respondent controls	Yes	Yes	Yes	Yes
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes
County and Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	.	.	29.79	8.20
Observations	28906	133550	28906	133550
Number of counties	1055	1058	1055	1058
R^2	0.15	0.14	0.09	0.11
Mean dep. var.	0.59	0.55	0.59	0.55

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. Respondent and county controls as in Table 5.
Standard errors clustered by county. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Views on immigration policy: County-level

Table C12: Views on immigration policy: Reduced form; county-level

	Reduced Form				
	(1) Index Restrict Imm.	(2) Border	(3) No Amnesty	(4) Don't hire	(5) Question
PostBan $\times$ AP-intensity	-0.0206*	-0.0086**	-0.0035	-0.0117**	-0.0082*
	(0.011)	(0.003)	(0.004)	(0.005)	(0.004)
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	9,239	9,239	9,239	3,536	7,232
Observations	2,104	2,104	2,104	1,768	2,040
Number of counties	0.39	0.35	0.38	0.58	0.41
R^2	0.26	0.61	0.59	0.67	0.49

Notes: Reduced form OLS regressions in the left hand-side panel. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C13: Views on immigration policy: 2SLS; county-level

	2SLS				
	(1) Index Restrict Imm.	(2) Border	(3) No Amnesty	(4) Don't hire	(5) Question
'Illegal imm.', pct. of 'Imm.'	0.0244*	0.0102**	0.0042	0.0112**	0.0108*
	(0.013)	(0.005)	(0.004)	(0.005)	(0.006)
Year FEs $\times$ County controls	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
First-Stage F stat.	36.83	36.83	36.83	29.08	20.27
First-Stage coef. on					
PostBan $\times$ AP-intensity	-0.8408***	-0.8408***	-0.8408***	-1.0508***	-0.7588***
	(0.138)	(0.138)	(0.138)	(0.195)	(0.168)
Observations	9,239	9,239	9,239	3,536	7,232
Number of counties	2,104	2,104	2,104	1,768	2,040
R^2	-0.05	-0.09	-0.01	-0.13	-0.11
Mean dep. var.	0.26	0.61	0.59	0.67	0.49

Notes: 2SLS regressions (upper panel), along with the corresponding 1st-stage coefficients (lower panel). County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C14: Support for increasing border security: county-level

	Reduced Form			2SLS			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	"Increase the number of border patrols on the US-Mexican border.": Selected						
PostBan × AP-intensity	-0.0083**	-0.0079**	-0.0086**	-0.0098**			
	(0.003)	(0.003)	(0.003)	(0.004)			
AP intensity	0.0093***						
	(0.003)						
PostBan	-0.0186						
	(0.012)						
'Illegal imm.', pct. of 'Imm.'					0.0082**	0.0102**	0.0121**
					(0.004)	(0.005)	(0.005)
Year FEs × County controls	No	No	Yes	Yes	No	Yes	Yes
County FEs	No	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	No	Yes	Yes	Yes	Yes	Yes	Yes
Year × State FEs	No	No	No	Yes	No	No	Yes
First-Stage F stat.	.	.	.	.	44.40	36.83	31.70
First-Stage coef. on PostBan × AP-intensity					-0.9661***	-0.8408***	-0.8140***
					(0.145)	(0.138)	(0.144)
Observations	9,407	9,274	9,239	9,224	9,274	9,239	9,224
Number of counties	2,245	2,112	2,104	2,101	2,112	2,104	2,101
R^2	0.01	0.34	0.35	0.36	-0.07	-0.09	-0.11
Mean dep. var.	0.61	0.61	0.61	0.61	0.61	0.61	0.61

Notes: Reduced form OLS regressions in the left hand-side panel, 2SLS regressions in the right hand-side panel. County controls: log population, share urban, racial composition, share foreign born, share college degree, log income per capita, newspaper circulation per capita and Republican vote share. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Chapter 2

MEDIA ATTENTION AND STRATEGIC TIMING IN POLITICS

Joint with Ruben Durante (ICREA, UPF, Barcelona GSE, IPEG and CEPR)[1]

2.1 Introduction

Mass media play a crucial role in informing citizens about government policies, allowing them to hold politicians accountable for their actions (Besley and Burgess 2002b; Snyder and Strömberg 2010). Yet, due to limited news space and audience attention, other newsworthy events can crowd out information that is relevant to evaluate government's behavior (Eisensee and Strömberg 2007). Taking this aspect into account, a sophisticated politician may have an incentive to time unpopular measures to moments when the media and the public are distracted by other news, so as to minimize public scrutiny of her actions.

There are many examples of political actions carried out or announced in coincidence with other newsworthy events, both in the U.S. and abroad. For example, on August 25th 2017 - the day North Korea launched several ballistic missiles and the day before hurricane Harvey struck Texas - president Trump enacted several controversial measures including pardoning Joe Arpaio, a former sheriff accused of racial profiling, and issuing a ban against transgender soldiers in the military.[2] In Russia, Putin's government announced a rise in the retirement age and an increase in the value added tax on the day of the inauguration of the 2018 FIFA World Cup which the country was hosting.[3] In Italy, Berlusconi's passed an emergency decree that freed hundreds of politicians with pending corruption charges on the day Italy qualified for the final of the 1994 FIFA World Cup.[4] Trying to anticipate and exploit the structure of the news cycle is also a well-known practice among political spin doctors.[5]

In this paper we examine study the strategic behavior of United States presidents focusing on one particular type of policy action: the signing of presidential executive orders (henceforth

[1]Orestis Exarchos, Nikola Kiprijanovski and Giulia Tosetti provided excellent assistance with coding the content of executive orders and TV news segments.

[2]https://www.theatlantic.com/politics/archive/2017/08/trump-news-dump-transgender-arpaio-gorka-harvey/538116/ (accessed on March 3rd, 2021)

[3]https://www.bloomberg.com/news/articles/2018-06-14/russia-plans-to-raise-retirement-age-increase-value-added-tax (accessed on March 3rd, 2021)

[4]http://www.archiviolastampa.it/component/option,com_lastampa/task,search/mod,avanzata/action,viewer/Itemid,3/page,1/articleid,0746_01_1994_0190_0001_15725553/anews,true/ (accessed on March 3rd, 2021)

[5]Ronald Reagan's communications assistant, David Gergen, once stated that "...if you've got some news that you don't want to get noticed, put it out Friday afternoon at 4pm" (cited in Gibson (1999)).

EOs). The ability of U.S. presidents to direct government through EOs derives from Article II of the U.S. Constitution which states that the president has the power to "take care that the laws be faithfully executed" - that is, to guide the execution of existing legislation. However, since EOs have the same value as federal laws and do not require Congressional ratification, in practice they have been often used to "guide" policy in a direction other than that intended by Congress, especially when the latter is not politically aligned with the president.

The signing of presidential EOs represents an ideal setting to analyze the question of strategic timing for at least two reasons. First, unlike other types of legislation, U.S. presidents have full discretion over *when* EOs are issued, so that there is ample scope to actively manipulate their timing. Second, though legislating through EOs offers the president a way to push his agenda and circumvent Congress, it can also generate criticism from Congress. The potential negative publicity associated with such criticism can create an incentive for the president to avoid media attention. Since the newsworthiness of an EO is usually short-lived, timing its signing to coincide with other news worthy events may be one viable strategy to minimize such negative publicity.

To test this hypothesis empirically we collect information on the timing and content of every EO signed by U.S. presidents between 1979 and 2016, and combine it with data on the content of daily evening news on the major U.S. broadcast TV networks. Following previous work on U.S. media (Eisensee and Strömberg 2007; Durante and Zhuravskaya 2018), we capture the presence of other important stories that may crowd out news about EOs with a daily measure of "news pressure". This is defined as the total airtime devoted to the top three stories featured on each news channel, excluding any stories related to EOs, and adjusting the length to keep the total duration of a newscast constant. Hence, higher levels of news pressure indicate days on which other important stories dominate the news cycle and on which EOs are more likely to go unnoticed.

We start by analyzing the relationship between news pressure, news coverage of EOs, and presidential approval ratings. First, we document that EOs tend to get covered by the media, although the majority of them - namely those that are relatively less significant or contentious â do not make the news. The news coverage of EOs is concentrated on the day an EO is signed and the following day and, importantly, is crowded out by other important stories proxied by news pressure.

Looking at how the public reacts to EOs, we find that EO-related news are associated with a decline in presidential approval ratings only in periods of divided government, i.e. when the Congress majority and the president belong to different parties, but not in periods of unified government. This is consistent with the idea that the public may react negatively to the president's use of EOs *if* Congress expresses criticism against it (Christenson and Kriner 2017b). Indeed, we document that EOs signed under divided government are substantially more likely to concern topics of prior disagreement between president and Congress, so that the scope for such criticism is arguably greater.[6]

We then turn to the analysis of the determinants of the timing of EOs. Our empirical strategy is based on daily time series regressions of an indicator for the signing of at least one EO in a given day, on lags and leads of news pressure, controlling for seasonality and the president's time in office. [7] We find that EOs are significantly more likely to be signed on the eve of

[6] One plausible explanation for these differences is that under divided government, if the president wants to legislate on issues on which Congress disagrees, he can only do so through unilateral action while under unified government he can push his agenda through a friendly Congress.

[7] Other work has focused on more general drivers of the use of unilateral power that are at play at higher levels of aggregation (Moe and Howell 1999; Howell 2003; Chiou and Rothenberg 2014). In our high-frequency

days characterized by high levels of news pressure, and this effect is only present in periods of divided government. The magnitude is sizeable: a standard deviation increase in next-day news pressure ($\approx$ 2.5min) is associated with a 1.1-percentage-point increase in the probability that at least one EO is signed on a given day, which corresponds to a 11% increase from a baseline probability of 10%. This result is robust to the use of different specifications, different measures of news pressure, and to the inclusion of a range of controls.

We then examine what type of EOs and what type of news are driving the relationship between news pressure and the timing of EOs under divided government. We find no effect for EOs that are routine or ceremonious in nature, i.e. those on government operations, and those with low significance (as estimated by Chiou and Rothenberg (2014)). Similarly, we find no effect for EOs that are unlikely to make the news, i.e. those that are not reported by the the Associated Press news wire which generally covers all newsworthy stories. Instead, the effect is driven by EOs that are ex-ante more likely to attract criticism for over-stepping presidential authority - i.e., on topics on which the president and Congress have disagreed more frequently in the prior months.

In terms of the type of news, the hypothesis of *forward-looking* strategic timing implies that only predictable news events can be targeted strategically to sway public opinion, while the same should not occur with unpredictable news. To test this prediction, we use dictionary-based text analysis methods to classify each news segment as being associated with anticipation (e.g., political campaign events, economic news, sports) or with surprise (e.g., accidents, natural disasters, violent crime), and construct two separate measures of news pressure. We find that the timing of EOs coincides with high levels of next-day news pressure related to anticipation but not to surprise. This finding is corroborated by a placebo exercise which exploits the occurrence of unpredictable events - such as major earthquakes, terrorist attacks and mass shootings. While these events lead to high news pressure, they are not associated with a higher probability of EO signing.

Finally, to shed light on why presidents may time EOs to next-day rather than same-day news pressure, we examine the systematic differences in the type of news coverage EOs receive on the day they are signed vs. the following day. We document that next-day coverage is significantly more likely to feature reactions from Congress (which, under divided government, tend to be negative), less likely to feature statements by the president, and is overall more negative in tone. Hence, targeting next-day news may be a sensible strategy if the goal is to minimize the publicity of such negative reactions.

Our work relates to several streams of literature. First, it contributes to previous work on limited attention (Gabaix et al. 2006), and to recent studies on the use of strategic timing by corporations (DellaVigna and Pollet 2009), NGOs (Couttenier and Hatte 2016), the military (Durante and Zhuravskaya 2018), and regulators (Potter 2017, 2019; Garz and Maass 2020). Our paper provides the first systematic evidence that similar tactics are employed by elected officials to limit public scrutiny of their actions.

Second, our research contributes to a large literature in political economy on the role of mass media in democratic societies, which documents that well-functioning media are key to discipline politicians and bolster political accountability (Snyder and Strömberg 2010; Besley and Burgess 2002b; Ferraz and Finan 2008). Our findings suggest that, even in the presence of free and independent media, politicians' strategic behavior can hinder citizens' ability to effectively monitor elected officials.[8]

specification these more aggregate factors are largely absorbed by calendar fixed effects.

[8]In this regard, our results also relate to recent findings by Balles et al. (2018); Kaplan et al. (2018) which document that, when media attention is captured by non-political events, U.S. representatives are more likely to

Finally, our paper relates to a large literature in political science on the use of presidential executive powers, and on the institutional factors that drive or constrain it. One view in this literature is that, since the threat of Congressional or judicial overturn is not credible (except for extreme cases of overreach), public opinion is the main factor that limits president's unilateral action (Posner and Vermeule 2010; Baum 2004; Reeves and Rogowski 2018; Christenson and Kriner 2019).[9] Several studies based on survey experiments have explored how the public reacts to the use of executive power, finding strong support for the view that EOs carry a risk of public backlash. For example, Reeves and Rogowski (2018) show that the same policy proposal draws significantly less support if enacted through executive order than through a federal law.[10] Crucially for the interpretation of our results, public opinion is not always opposed to the use of EOs (Christenson and Kriner 2017a), but is rather activated by Congressional criticism (Christenson and Kriner 2017b).[11] While there is evidence that public opinion – and the ability of Congress to influence it – constrains unilateral power, our paper enriches this framework by documenting that presidents may attempt to circumvent this constraint through strategic behavior.

The rest of the paper is organized as follows. In section 2.2 we describe the data and the construction of our main variables. Section 2.3 presents preliminary evidence on the news coverage of EOs. In section 2.4 we discuss our empirical strategy and present the main evidence of strategic timing. Section 2.5 presents heterogeneity analysis of the main effect. In section 2.6 we discuss possible mechanisms. Section 2.7 concludes.

2.2 Data

Our analysis combines a wide range of data. First, we gather comprehensive information on the signing date and content of all EOs issued by U.S. presidents over the past four decades. Second, to investigate the relationship between the timing of EOs and the news cycle, we collect data on the news stories featured in the daily evening newscasts of the major U.S. broadcast TV networks. In various parts of the analysis, we also use data on: i) presidential positions and roll call voting on bills considered in Congress ii) coverage of EOs on the Associated Press news wire, iii) the occurrence of major earthquakes, terror attacks and mass shootings, iv) the volume of Google searches related to EOs, v) president's approval ratings. Table A1 presents summary statistics for all main variables.

2.2.1 Sample

The sample period of our analysis spans 1979 to 2016. Throughout the analysis we distinguish between periods of divided vs. unified government, i.e. aligned vs misaligned party control of White House and Congress (detailed in Table A2). In Appendix A2 we show that the main difference between EOs issued under unified versus divided government is that in the latter case EOs are substantially more likely to concern topics of prior disagreement between president and Congress.

vote in line with the preferences of special interests as opposed to those of their constituents.

[9]Related work studies unilateral action at the state-level, in this case highlighting the constraining role of legislatures and the strategic incentives arising from these constraints (Bolton and Thrower 2021; Sellers 2017).

[10]Other studies by the same authors further corroborate these findings (Reeves and Rogowski 2016, 2015; Reeves et al. 2017).

[11]Similar evidence is available from Christenson and Kriner (2017c) on the impact of criticism of EOs by the judiciary.

2.2.2 Executive Orders

Date, subject and topic. Comprehensive data on the universe of EOs are available from the American Presidency Project[12]. The data include information on the date of issuance, a short summary and the full text of each EO. From the text we identify a set of keywords indicative of the subject of each EO, which we then use to find related news stories. To do so we use two distinct procedures. First, we instructed a research assistant to read the summary of each EO and identify two to three words or phrases particularly descriptive of the subject matter. Second, we consider the entire corpus of EO-s texts in our sample, and perform an automated keyword selection based on term-frequency/inverse document frequency (tf-idf) – a standard statistic used to identify terms descriptive of a document within a corpus.[13] For each EO, we consider as "keywords", the five uni- or bi-grams with highest tf-idf score. Table A3 presents examples or the (stemmed) keywords obtained using these two alternative procedures. We use manually coded keywords in our baseline analysis, and keywords from the automated procedure in robustness checks.

We also use information on the broad topic of each EO, as well as the topics of Congressional bills, coded into 20 categories by the Comparative Agendas Project[14]. Figure 1 reports the topic distribution of the 1647 EOs in our sample.

Figure 1: Distribution of EOs by Topic

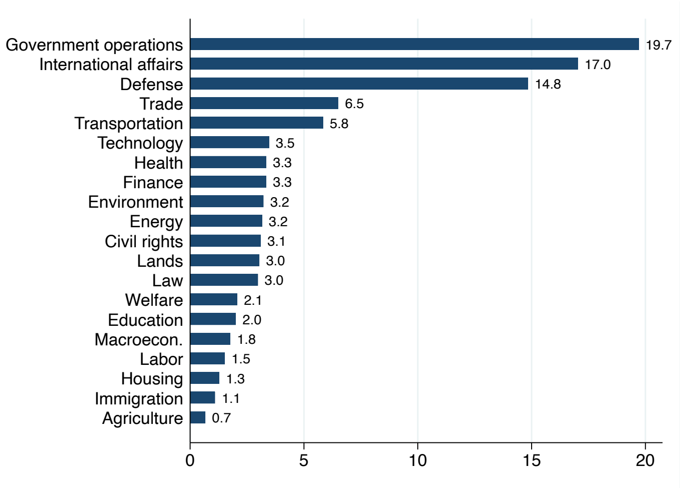

Distribution of EOs by major topic, as classified by the Comparative Agendas Project.

EO Significance. For a measure of the political significance of each EO we use the index proposed by Chiou and Rothenberg (2014), available for EOs signed before 2003. The index is based on a hierarchical item response model, applied to data from 19 sources including historical overviews of EOs, national newspapers, general news magazines, politics and policyâfocused magazines, and top law reviews.

[12] http://www.presidency.ucsb.edu/ (accessed on March 3rd, 2021)

[13] Intuitively, tf-idf increases with the frequency of the term within a document but decreases with the number of documents in the corpus in which the term appears, thereby discounting terms that are less useful to distinguish one document from the rest. In the case of EOs, procedural terms commonly used in EOs (e.g. "executive", "amendment", "continuation") are heavily discounted. For a lengthier discussion of the tf-idf method see Gentzkow et al. (2018) and Grimmer and Stewart (2013).

[14] https://www.comparativeagendas.net (accessed on March 3rd, 2021)

Congressional voting and presidential positions. To proxy the degree of disagreement between Congress and president on the topic of a given EO, we measure differences in the support for bills on that topic recently considered by Congress. Roll-call votes and presidential positions (clear public statements by the president on specific bills) are available from Voteview.[15] Overall, our sample includes 1148 votes on the final passage of new legislation. For each bill, we construct a dummy variable for whether the vote of the congressional majority went against the presidential position, and then compute the rolling six-month average by topic.[16] We label an EO-topic as one of "high disagreement" if the average frequency of disagreement over the previous six months is above the median value (66.6% for periods of divided government).

2.2.3 News content

Our main source of data on TV news content is the Vanderbilt News Archive (VNA).[17] The VNA includes comprehensive information the news stories featured on the daily evening newscasts of the three main U.S. broadcast networks (ABC, CBS, NBC) since 1968, and, for CNN, since 1992. We focus on the years after 1979 for which daily data are available. For each news story the VNA reports the order, the length, the headline, and a short summary.

News coverage of executive orders. To measure news coverage of EOs, we search the VNA database for news containing the following combinations of keywords: "executive" + ("order(s)" or "action(s)" or "authority"), or "presidential" + ("order(s)" or "action(s)" or "authority"). According to this measure, the majority if EO-related airtime is concentrated on the day of the signing and on the following day (Figure 2).

Figure 2: EO News Coverage by Distance from Closest EO-Signing

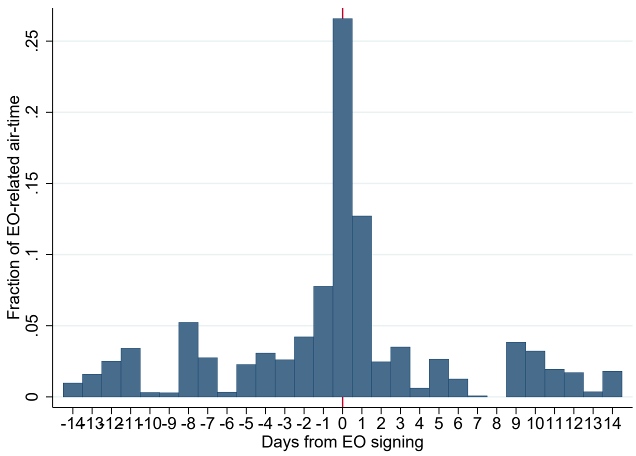

Volume of EO-related airtime in evening newscasts by day from the closest EO signing. Normalized by total EO-related airtime.

News pressure. Following previous related work (Eisensee and Strömberg 2007; Durante and Zhuravskaya 2018), we capture the availability of other news that may crowd out coverage of EOs with a measure of daily "news pressure". This is defined as the airtime devoted, on a

[15]https://voteview.com/data (accessed on March 3rd, 2021)

[16]The results we present are robust to a 12-month or a 3-month window.

[17]https://tvnews.vanderbilt.edu/ (accessed on March 3rd, 2021)

given day and given channel, to the top three news stories not related to EOs. The intuition behind this measure is that, to the extent that the top three stories capture most of the attention, and given the constraint of the 30-minute format of evening news, the more time is devoted to these stories, the less time there is to cover other news, including EOs.[18] Therefore, *ceteris paribus*, on days with higher news pressure news coverage of EOs should be lower.

To compute news pressure accurately, it is important to identify and exclude any news that may be related to an EO or to its subject matter. We therefore first exclude all news segments that explicitly mention the phrase "executive order" or synonyms. Yet, this step would omit news that discuss the policy and its consequences without explicitly mentioning that it was enacted through EO. To capture these instances, we also exclude all news segments that contain any EO-subject specific keywords and that were aired around the time an EO is signed. In our baseline specification we consider the window of -1/+1 days from the signing of the EO, but our results are robust to alternative windows.

Table A4 illustrates this approach for the example of executive order # 13505 on "Removing Barriers to Responsible Scientific Research Involving Human Stem Cells" signed by President Obama on March 9$^{\text{th}}$ 2009. In this case, our procedure excludes a story that mentions the expression "executive order", but also a story aired on the same day which, though not referring to executive order, clearly covers the same issue using words such as "stem cells" and "research".

Crucially, to be able to compare days with and without EO-related news, when excluding any news segment we adjust for the diminished total length of the newscast. This is important because, as shown by Durante and Zhuravskaya (2018), under mild assumptions the measure of news pressure adjusted for total length has no mechanical correlation with the excluded news.[19] In contrast, the un-adjusted measure has mechanically lower values on days when news about EOs are featured (and hence, on days with EOs).

Once news pressure for each network/day is computed, we take the median across all networks to derive aggregate daily news pressure.

Surprising vs. anticipated news. To investigate whether EOs are more likely to coincide with predictable news, we apply a dictionary method to decompose the news pressure variable into two components: one driven by surprising news and another driven by anticipated news. We use the NRC Word-Emotion Association Lexicon which provides a dictionary of words associated with anticipation (e.g. investigation, inauguration) and ones associated with surprise (e.g. earthquake, explosion).[20] We compute "surprise" news-pressure using the sample of segments containing strictly more "surprise" words than "anticipation" words. We compute an analogous measure of "anticipation" news pressure based on segments containing strictly more "anticipation" words. Figure 3 reports the most frequent terms in the headlines of "surprise" and "anticipation" news segments respectively.

Unpredictable newsworthy events. To validate the text-based measures of "surprise" and "anticipation " news pressure introduced above, we collect data on the occurrence of unpre-

[18] We exclude from the analysis September 11, 2001 for which news pressure is undefined because evening newscasts on that day far exceeded 30 minutes.

[19] Specifically, this is the case if, upon arrival of EO-related news, the length of other top-3 and non-top 3 news is reduced proportionately. Durante and Zhuravskaya (2018) test and confirm the validity of this assumption using the case of disaster-related news.

[20] http://saifmohammad.com/WebPages/NRC-Emotion-Lexicon.htm (accessed on March 3rd, 2021)

Figure 3: Word Clouds of News Associated with "Surprise" and "Anticipation"

(a) Surprise (b) Anticipation

Fifty most frequent words (excl. names of people and places) in the headlines of TV segments classified as associated with surprise (panel a) or with anticipation (panel b).

dictable newsworthy events, i.e., major US mass shootings[21], worldwide earthquakes[22], and worldwide terrorist attacks[23]. To ensure that we look at events that are newsworthy from the standpoint of U.S. media, we focus on U.S.-based events in which at least 10 people were killed or injured, and on foreign-based events in which at least 50 people were killed or injured. We consider all countries for earthquakes and the U.S. and Western Europe for terrorist attacks.[24] Overall, our sample includes 48 shootings, 130 earthquakes, and 113 terror attacks, for a total of 286 days with at least one event.

Associated Press coverage of EOs. As a proxy for newsworthiness, we construct a measure for whether the Associated Press (AP) news wire released EO-related news on the date of signing of each EO. To the extent that AP has a constant presence in the White House and since, compared to 30-minute TV newscasts, it faces fewer constraints on the volume of news it can cover, EOs that are not covered by AP are arguably less important and less likely to be featured on national TV. To identify AP coverage of EOs, we apply the same keyword search queries used for the VNA to the Dow Jones Factiva database restricted to "Associated Press Newswires"[25]. These data are available from 1988 on-wards. We label an EO as *not* covered by AP if no wire articles matching our search criterion was found on the day the EO was signed. This is the case for about 35% of EOs.

2.2.4 Public reactions to EOs

Google trends. To gauge how news coverage of EOs influences public awareness and interest, we collect data on the volume of Google searches related to EOs from Google trends.

[21]FBI's Supplementary Homicide Reports: https://ucr.fbi.gov/nibrs/addendum-for-submitting-cargo-theft-data (accessed on March 3rd, 2021). Available from 1982 on-wards.

[22]EM-DAT: https://www.emdat.be/. Available up to 2013.

[23]Global Terrorism Database: https://www.start.umd.edu/gtd/ (accessed on March 3rd, 2021). Available up to 2015.

[24]Indeed, attacks in other countries do not generate enough interest by U.S. media to significantly increase news pressure.

[25]https://www.dowjones.com/products/factiva/ (accessed on March 3rd, 2021)

These data are available from 2004 on-wards. We focus on the daily volume of searches for the topic "executive order" as defined by Google, which aggregates several related queries.

Presidential approval ratings. To assess how EOs affect the president's popularity, we use data on presidential approval ratings collected by Gallup and available from the American Presidency Project. Gallup conducts periodic multi-day polls asking the following question: *"Do you approve or disapprove of the way [president name] is handling his job as president?"*. Poll are carried out with on average weekly frequency (daily in more recent years) and each poll spans 1 to 4 days. We convert the share of respondents to a given poll who disapprove of the president's performance to a daily time series by assigning the reported poll-level average to the days over which the poll was conducted, and taking the mean in the case of overlap between polls.

2.3 Preliminary Evidence

Before analysing the relationship between news pressure and timing of EOs, we discuss some preliminary evidence related to the news coverage of EOs and verify the premise that the publicity of president's unilateral actions is lower on days with high news pressure.

Table 1: News Coverage of EOs: News Pressure and Google Searches

	All days		Days with EO in t or t-1		2004-2016	
	(1) Any EO-news	(2) Length EO-news	(3) Any EO-news	(4) Length EO-news	(5) Log Google searches for 'EO'	(6) Log Google searches for 'EO'
EO in t or (t-1)	0.014***	3.093***				
	(0.004)	(0.424)				
NP (t)			-0.014	-2.757**		
			(0.015)	(1.262)		
EO news (t or t-1)					1.023***	
					(0.237)	
Length of EO news (t or t-1)						0.002***
						(0.000)
EO topic in t or (t-1)	No	No	Yes	Yes	Yes	Yes
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes
Year, Month, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13880	13880	2600	2600	4685	4685
(Pseudo) R-Squared	0.014	0.018	0.042	0.077	0.267	0.278

Full sample in columns (1) and (2), sample restricted to days with EO signing in t or (t-1) in columns (3) and (4), and sample post-2004 in columns (5) and (6). Dependent variable: indicator for, and length of, EO-related news in columns (1) through (4), and log Google trends volume in columns (5) and (6). OLS in columns (1), (3), (5) and (6), maximum likelihood negative binomial in columns (2) and (4). Standard errors clustered by month $\times$ year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

We first document that EOs can make the news. In the first column of Table 1 we consider our entire sample period and regress a dummy variable for whether stories about EOs are featured in the news on a given day on a dummy for whether any EO was signed on the same or the previous day. The result indicates that about 1.4% of all EOs get covered on the same or following day. On such days, the airtime devoted to EO-related stories increases twenty-fold relative to days with no EOs (when EO-related airtime is just 2 seconds). In Table A8 we show that TV coverage is substantially greater for relatively more politically significant or contentious categories of EOs, i.e., ones of higher significance, on topics other than government operations, on topics of disagreement between president and Congress, and those covered in the Associated Press wire.

In columns 3 and 4 of table 1 we test whether news pressure crowds out news on EOs. In this case, we restrict the sample to days with EO-signing in the same or previous day, and examine the relationship between news pressure and the presence and length of EO-related news, conditional on fixed effects for EO-topic. While for the indicator for any EO-related news. We find a negative, though imprecisely estimated coefficient for the indicator for any EO-related news and a negative coefficient significant at the 5% level for the length of EO-related news – one standard deviation increase in news pressure ($\approx$ 2.5min) reduces the time devoted to EO-related news by 50%.

Furthermore, the news coverage of EOs appears to draw public attention, proxied by the daily volume of EO-related Google searches. Such searches increase two-fold if news about EOs are aired on the same or previous day or by 12% for an additional minute of coverage, controlling for the occurrence of EO signing and for EO-topic fixed effects (columns 5 and 6).

Finally, in Table 2 we examine the association between news coverage of EOs and president's popularity, measured by (dis)approval ratings in Gallup polls. While we find no relationship when government is unified (columns 1 and 2), in periods of divided government the presence of EO-related news or an additional minute of EO-related coverage is associated with a significant increase in disapproval of the president's performance of 0.7 percentage points or 0.06 percentage points respectively, controlling for EO-topic fixed effects and lagged approval (column 3 and 4).[26]

While only correlational, these patterns are in line with previous findings by Christenson and Kriner (2017b) and Reeves and Rogowski (2018) showing that, if given publicity, EOs can be politically costly for the president, especially in the presence of a hostile Congress.

2.4 Empirical Strategy and Results

2.4.1 Empirical Strategy

To examine the relationship between the timing of EOs and the presence of potentially distracting news, we conduct a time-series analysis at daily frequency, regressing an indicator for the signing of at least one EO on leads and lags of news pressure. We control for various dimensions of seasonality which are relevant for the political cycle and for the news cycle.

The following equation summarizes our econometric strategy:

$$EO_t = \alpha_0 NP_t + \beta_0 NP_{t+1} + \sum_{\tau=1}^{7} \alpha_\tau NP_{t-\tau} + \sum_{\tau=2}^{7} \beta_\tau NP_{t+\tau} + \gamma W_t + \eta_{d_t} + \psi_{m_t} + \nu_{y_t} + \epsilon_t, \quad (2.1)$$

EO_t is a dummy variable for whether at least one EO is signed on day t; NP_t indicates news pressure on day t; W_t is the number of weeks since the start of the presidential term; η_{d_t}, ψ_{m_t} and ν_{y_t} are day-of-week, calendar month, and year fixed effects respectively.

[26]The increase in disapproval is stronger for EOs on topics on which president and Congress disagree (0.84 percentage points for an additional minute of coverage). These results are available upon request.

Table 2: News Coverage of EOs: Association with Approval Ratings

	Unified Gov.		Divided Gov.	
	(1) Gallup Disapproval	(2) Gallup Disapproval	(3) Gallup Disapproval	(4) Gallup Disapproval
EO news (t or t-1)	-0.001		0.660**	
	(0.561)		(0.321)	
Length of EO news (t or t-1)		-0.000		0.001***
		(0.001)		(0.000)
Disapproval past 30 days	0.870***	0.870***	0.943***	0.944***
	(0.034)	(0.034)	(0.028)	(0.028)
EO topic in t or (t-1)	Yes	Yes	Yes	Yes
Weeks in office	Yes	Yes	Yes	Yes
Year, Month, DOW FEs	Yes	Yes	Yes	Yes
Observations	1444	1444	4318	4318
R-Squared	0.943	0.943	0.971	0.971
Mean dependent variable	42.8	42.8	40.8	40.8

Sample: unified government in columns (1) and (2), divided government in columns (3) and (4). Dependent variable: percent of Gallup respondents who report that they disapprove of the performance of the incumbent president. OLS in all columns. Standard errors clustered by month $\times$ year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

There are two possible sources of endogeneity in this regression: i) EOs may generate news that increase news pressure (reverse causality), and 2) EOs may be related to other events that generate news and increase news pressure (omitted variable bias). As explained in detail in section 2.2, we address both of these concerns by focusing on variation in news pressure that is unrelated to the direct coverage of EOs or to the subject matter of recent and forthcoming EOs.

In our baseline analysis we assume a linear probability model for the indicator of any EO (though the results are robust to Probit). Alternatively, we consider the number of EOs signed and estimate maximum likelihood negative binomial regressions. To account for serial correlation in EO signings and news pressure we cluster standard errors by month $\times$ year, but obtain similar results applying the Newey-West estimator or more aggregate clusters.

2.4.2 Baseline Results

We start by estimating equation 2.1 for the entire sample period 1979-2016. In column 1 of Table 3 we regress a dummy for the signing of at least one EO on a given day on same-day and next-day news pressure, controlling for weeks in office and calendar fixed effects. In the following specifications we gradually include 7 lags of news pressure (column 2) and 7 leads of news pressure (column 3).[27] The results indicate a positive, though only marginally significant relationship between next-day news pressure and the likelihood of EO signing.

As discussed above, presidents should arguably have a stronger incentive to time EOs strategically when facing a hostile Congress than a friendly one. To test this hypothesis, in column 4 we interact same-day and next-day news pressure, as well as all other lags and leads, with a dummy for periods of divided government. The coefficient on the interaction between next-day news pressure and divided government is positive, large, and statistically significant (at the 1% level), while all other interaction terms are insignificant. These results suggest that presidents

[27]The number of observations changes between columns due to missing new pressure for September 11 2001 and its respective leads and lags.

are more likely to sign EOs on the eve of days with high news pressure, but only when Congress is not politically aligned with them.

Table 3: News Pressure and the Timing of EOs: Divided vs. Unified Government

	Full Sample			
	(1) EO	(2) EO	(3) EO	(4) EO
NP	0.007 (0.012)	0.004 (0.013)	0.005 (0.013)	0.003 (0.024)
NP (t+1)	0.023* (0.013)	0.023* (0.013)	0.024 (0.014)	-0.038 (0.026)
NP (t-1)		-0.000 (0.014)	-0.001 (0.014)	0.023 (0.027)
NP × Divided				0.002 (0.029)
NP (t+1) × Divided				0.084*** (0.031)
NP (t-1) × Divided				-0.033 (0.032)
7 lags of NP	No	Yes	Yes	Yes
7 leads of NP	No	No	Yes	Yes
Weeks in office; Year, Month, DOW FEs	Yes	Yes	Yes	Yes
7 leads and lags of NP × Divided	No	No	No	Yes
Observations	13875	13854	13836	13836
R-Squared	0.042	0.042	0.042	0.043
Mean dependent variable	0.100	0.099	0.099	0.099

	Divided Government			Unified Government		
	(5) EO	(6) EO	(7) EO	(8) EO	(9) EO	(10) EO
NP	-0.002 (0.014)	0.003 (0.016)	0.004 (0.016)	0.028 (0.021)	0.003 (0.023)	0.004 (0.024)
NP (t+1)	0.042*** (0.015)	0.045*** (0.015)	0.045*** (0.017)	-0.029 (0.023)	-0.039 (0.025)	-0.037 (0.027)
NP (t-1)		-0.010 (0.016)	-0.011 (0.016)		0.026 (0.028)	0.024 (0.028)
7 lags of NP	No	Yes	Yes	No	Yes	Yes
7 leads of NP	No	No	Yes	No	No	Yes
Weeks in office; Year, Month, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes
7 leads and lags of NP × Divided	No	No	No	No	No	No
Observations	10133	10126	10114	3742	3728	3722
R-Squared	0.042	0.042	0.042	0.047	0.048	0.049
Mean dependent variable	0.098	0.098	0.097	0.105	0.104	0.105

Full sample in columns (1)-(4), divided government in columns (5)-(7), unified government in columns (8)-(10). Dependent variable: indicator for the signing of an EO. OLS regressions in all columns. Standard errors clustered by month × year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

To corroborate the key distinction between divided and unified government, we re-estimate the first three columns separately for these different samples (columns 5-7 and 8-10, respectively). The results are consistent with those from the interacted model. The coefficient on next-day news pressure in the divided government sample suggests that a standard deviation increase ($\approx$ 2.5min), is associated with a 1.1 percentage point increase in the likelihood of an

EO signing, or an 11% increase relative to the mean likelihood of 10%. It is robust to controlling for lags and leads, and statistically different from the coefficient estimated for unified government at the 5% level. We also find a significant increase in the *number* of EOs issued, in the order of 16% (Table **??**), which suggests that strategic timing may also affect presidents' decisions on the intensive margin.

In Figure 4 we plot the coefficients for different leads and lags of news pressure estimated either simultaneously (left panel) or one by one (right panel), separately for divided government (top) and unified government (bottom). For divided government, the coefficient on news pressure at *t+1* (i.e. next-day) is larger than the ones on other lags and leads and is the only statistically significant one. Instead, no clear pattern emerges for unified government.[28]

Figure 4: Leads and Lags of News Pressure and the Timing of EOs

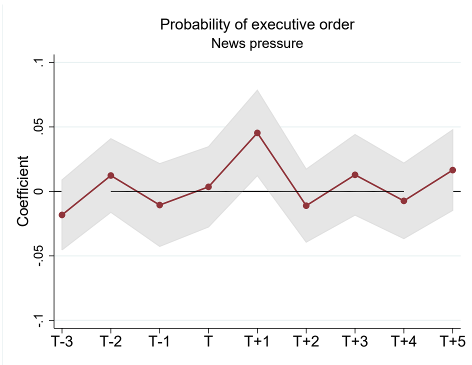

(a) Divided Government Sample: Simultaneous

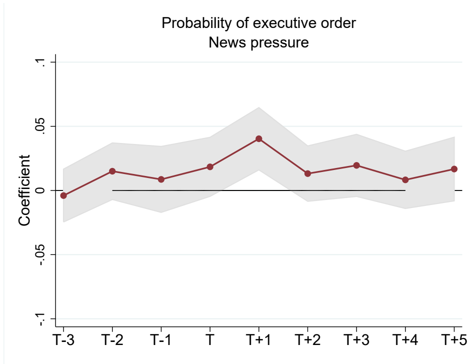

(b) Divided Government Sample: One by One

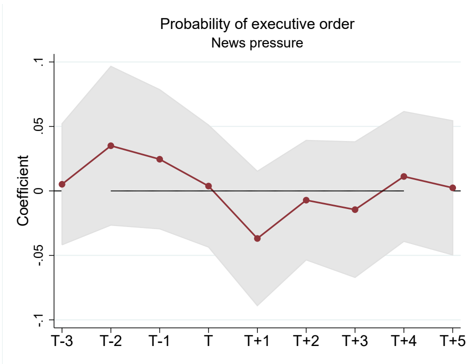

(c) Unified Government Sample: Simultaneous

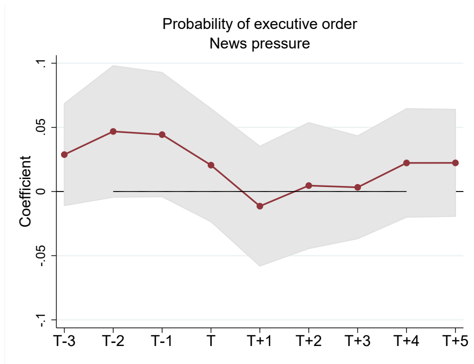

(d) Unified Government Sample: One by One

Coefficients on leads and lags of news pressure in the samples of unified vs divided government, estimated either simultaneously (corresponding to columns (7) and (10) of table 3 respectively.), or one-by-one.

In Table A9 we explore how the association of next-day news pressure with the timing of EOs varies with the *degree* of political misalignment between Congress and the president.

[28]We obtain consistent results using the full sample and plotting the interaction of each lead/ lag of news pressure with an indicator for divided government (Figure 5).

In line with the view that presidents are more likely to act strategically when facing a hostile Congress, we find that the effect is generally more pronounced when the party opposing the president controls both chambers of Congress rather than one (though the difference is not statistically significant).

Taken together, these results suggest a pattern in the timing of EOs that is in line with targeting of distracting newsworthy events. Crucially, and also in line with our hypothesis, this only applies to periods of divided government. In light of this finding, in the remainder of the analysis we will restrict our focus to periods of divided government.

Figure 5: Leads and Lags of News Pressure × Divided Government

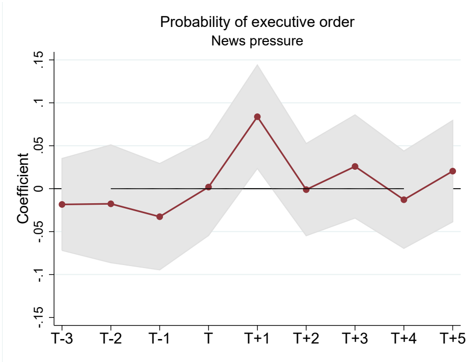

(a) Simultaneous

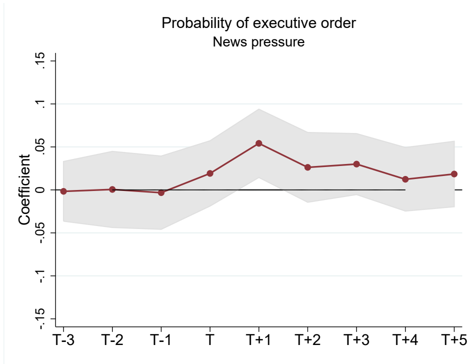

(b) One by One

Coefficients on the interaction of divided government with leads and lags of news pressure, estimated either simultaneously (corresponding to column 4 in Table 3), or one-by-one.

2.4.3 Robustness

Alternative specifications and controls. Next we show that the findings presented above are robust to various alternative specifications, estimation models and controls.

First, the relationship between next-day news pressure and the timing of EOs holds in the raw data and is not driven by any particular functional form assumption. In Figure 6 we report the share of days with EO signings by quintile of next-day news pressure (panel a), as well as a non-parametric locally weighted regression (panel b).

Second, the results are not sensitive to alternative estimation of the standard errors or to the inclusion of alternative controls. Table A5 reports results using the Newey-West estimator to adjust for serial correlation, and with clustering at the more aggregate level of calendar year or Congressional term (columns 1 to 3).[29] We also find similar results controlling more flexibly for number of weeks in office, i.e. allowing the effect to vary by President, or including number of week-in-office fixed effects (columns 3 and 4).

In Table A6 we report the following additional checks: i) probit instead of a linear probability model, ii) omitting any calendar controls, iii) controlling for year × month fixed effects, iv) controlling for lagged EO signings, and v) controlling for federal holidays[30] and days of presidential foreign visits.[31].

[29]This comes with the caveat of a substantially smaller number of clusters – 31 in the case of calendar years and 16 in the case of Congressional terms.

[30]Available from https://www.calendar-365.com/2019-calendar.html (accessed on March 3rd, 2021).

[31]Available from https://history.state.gov/departmenthistory/travels/president (accessed on March 3rd, 2021).

Figure 6: Timing of Executive Orders: Non-Parametric Estimation

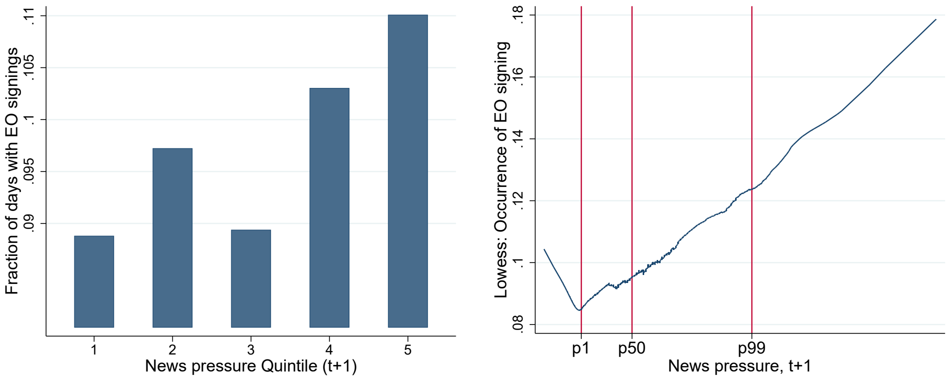

(a) Frequency of EO signings by quintile of next-day news pressure.

(b) Local linear regression of EO signing on next-day news pressure.

Sample: divided government. Panel (a): Average fraction of days with at least one EO signing, by quintile of the next-day news pressure distribution. Panel (b): Nonparametric locally weighted regression of an indicator for EO-signing on next-day news pressure. Vertical lines indicate the median, the 1st, and the 99th percentile of the news pressure distribution.

Alternative measures of news pressure. As discussed in section 2.2, our preferred measure of news pressure is computed in two steps. First, we exclude any news segments that mentions the phrase "executive order" or synonyms and correct for their length. Second, we exclude any news aired in proximity to an EO-signing that mention EO-specific keywords and correct for their length. In the left hand side panel of Table A7 we estimate our baseline specification with news pressure computed following only the first step, without correction for length of the excluded segments (column 1), and with correction (column 2). In column (3) we add the second step, thus obtaining our baseline result. The fact that both the magnitude and precision of the coefficient increase in this step confirms the importance of capturing news that, despite not mentioning EOs explicitly, talk about their subject matter. This also suggests that the observed association is likely driven by news that are entirely unrelated to EOs.

In the right-hand side of Table A7 we show that the results are robust to alternative versions of news pressure: i) using the top three news stories ranked by length, rather by order of appearance (column 4), ii) using keywords from automated text-analysis rather than human-coded ones (column 5), iii) excluding any keywords within +7/-7 days from EO-signing rather than within -1/+1 days (column 6).

2.5 Heterogeneity

Up to now we have documented a strong empirical relationship between the timing of EOs and next-day news pressure in periods of divided government. In what follows, we investigate what type of EOs and what type of news are driving this relationship.

2.5.1 Types of Executive Orders

We hypothesize that the incentive for strategic timing is more pronounced for EOs that are i) politically significant, ii) *ex ante* more likely to generate criticism, and iii) *ex ante* more likely to be covered in the news.

To test for heterogeneity with respect to these characteristics, we estimate a series of multinomial logit regressions comparing the association between next-day news pressure and the probability of issuance of an EO of one type vs. the opposite type, relative to the likelihood of no EO (Figure 7). Alternatively, we estimate a series of linear probability regressions where the dependent variable an indicator equal to one if EOs of a particular type are issued on a given day, and equal to zero for days with EOs of the opposite type or no EOs (Table 4).

Looking at various proxies for the above characteristics, we find results consistent with our predictions.[32] Regarding political significance, we find that the association with news-pressure is driven by EOs on topics other than government operations and EOs of high significance (panels a and b of Figure 7) – one standard deviation increase in next-day news pressure increases the likelihood of such EOs by 13% and 18% respectively. The likelihood of an EO that is likely to generate criticism – i.e. on a topic of high disagreement between president and Congress – increases by 22% (panel c). The likelihood of a newsworthy EO – one covered by the AP – increases by 14% (panel d). We find no correlation with news pressure when looking at EOs of the opposite, less contentious types.

Additionally, we examine two smaller categories of EOs for which we would *not* expect the president to have the incentive or ability to act strategically. First, some EOs are discussed in the news prior to their signing, which is unlikely to happen if the administration aims to "conceal" them. We find no effect EOs issued when EO-related news were aired in the prior 7 days, which are about 6% of the total (panel e). Second, some EOs are signed in response to emergency situations which call for swift presidential action – their timing is hence likely dictated by urgency rather than media considerations. We find no effect for EOs whose description contains the keyword "emergency", which are about 5% of the total (panel f).

2.5.2 Predictable *vs.* Unpredictable News

The hypothesis of *forward-looking* strategic timing implies that EO-signing should only coincide with news that can be anticipated. This prediction is reinforced by the result that only *next-day* news pressure (*t+1*) exhibits a significant correlation with the probability of EO signing.

To test this prediction, in Table 5 we conduct a placebo exercise exploiting the timing of arguably unpredictable events - earthquakes, terror attacks and mass shootings. We document that each of these events significantly increases news pressure in the day of its occurrence (columns 1-4). However, we find no significant relationship between the occurrence of an unpredictable event on the following day and EO signing (columns 9-12). Furthermore, using next-day unpredictable events as an instrument, we find no evidence that the corresponding unexpected increase in next-day news-pressure is related to EO signing (columns 5-8). Hence, the variation in news pressure generated by unpredictable news does not seem to be what is driving our result.

As a more comprehensive test, we use a dictionary-based text analysis procedure to classify all news segments in our sample into two mutually exclusive categories: those associated with

[32] As detailed in section 2.2, most of these measures are defined based on the date in which the EO was signed, its topic, or a combination of the two. An exception is the measure of significance from Chiou and Rothenberg (2014) which is EO-specific, but has the limitation that one of its components is potentially endogenous news coverage.

Figure 7: Heterogeneity by Type of EO

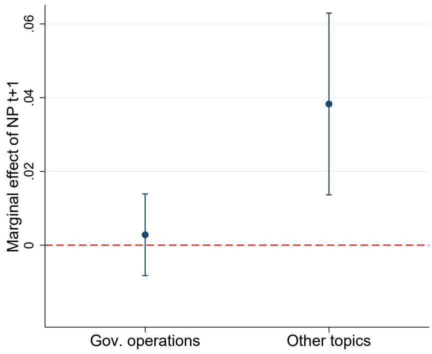

(a) Government operations vs other topics

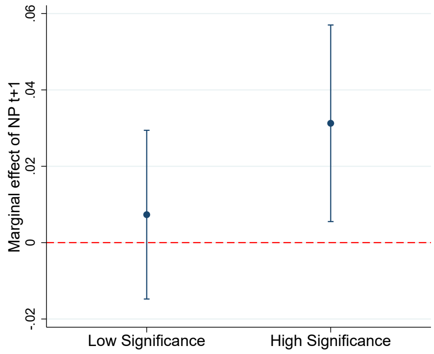

(b) EO significance

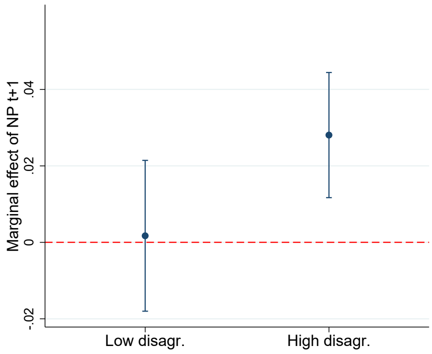

(c) President-Congress diagr.

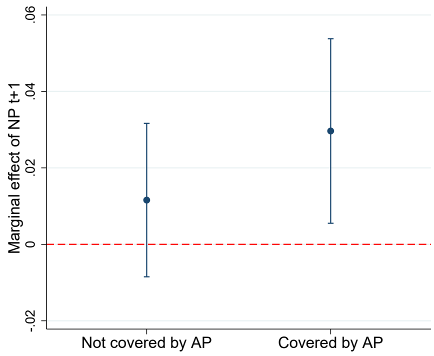

(d) Covered by AP vs not

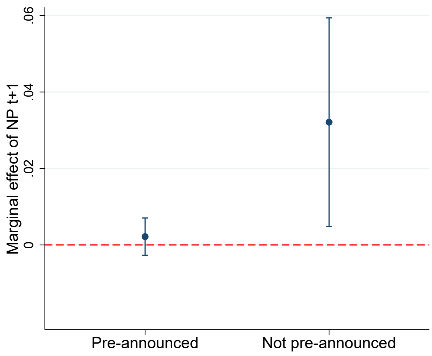

(e) Pre-announced vs not

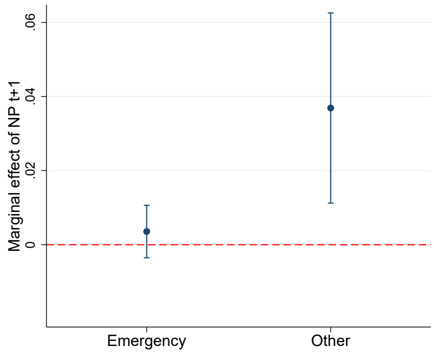

(f) Emergency vs other

Sample: divided government. Marginal effects (along with 95% confidence intervals) of a change in next-day news pressure on the probability of signing of an EO of a certain type. Estimated from a multinomial logit regression conditional on baseline controls. Standard errors clustered by month $\times$ year.

surprise and those associated with anticipation. To validate this approach, and to relate it to our previous exercise, in Figure 8 we document that the news pressure associated with surprise increases in coincidence with major unpredictable events while the news pressure associated with anticipation does not.

Exploiting this decomposition, in Table 6 we examine what type of news drives the relationship with EO-signings. Looking separately at the two components of next-day news pressure,

Table 4: Timing by Type of EO

	(1) EO Not gov. operations	(2) EO High Signif.	(3) EO Covered by AP	(4) EO High Disagr.
NP	0.006	0.004	0.007	-0.002
	(0.015)	(0.015)	(0.014)	(0.010)
NP (t+1)	0.042***	0.035**	0.032**	0.030***
	(0.014)	(0.015)	(0.015)	(0.011)
NP (t-1)	-0.005	-0.014	-0.011	-0.005
	(0.015)	(0.016)	(0.015)	(0.012)
Weeks in office	Yes	Yes	Yes	Yes
Year, Month, DOW FEs	Yes	Yes	Yes	Yes
7 lags of NP	Yes	Yes	Yes	Yes
Observations	10126	7189	7581	7954
R2	0.034	0.023	0.045	0.029
Mean dep. var.	0.081	0.050	0.057	0.034
Mean dep. var. if EO=1	0.835	0.483	0.616	0.529

	(1) EO Gov. operations	(2) EO Low Signif.	(3) EO Not covered by AP	(4) EO Low Disagr.
NP	-0.002	-0.002	0.013	0.012
	(0.007)	(0.013)	(0.011)	(0.011)
NP (t+1)	0.003	0.007	0.010	0.004
	(0.006)	(0.012)	(0.011)	(0.010)
NP (t-1)	-0.005	-0.011	-0.003	-0.014
	(0.008)	(0.013)	(0.012)	(0.011)
Weeks in office	Yes	Yes	Yes	Yes
Year, Month, DOW FEs	Yes	Yes	Yes	Yes
7 lags of NP	Yes	Yes	Yes	Yes
Observations	10126	7189	7581	7954
R2	0.013	0.032	0.053	0.025
Mean dep. var.	0.016	0.054	0.035	0.030
Mean dep. var. if EO=1	0.165	0.517	0.384	0.471

Sample: divided government. Dependent variable: indicator equal to one if an EO of a certain type was signed in the respective day, and zero if not. OLS regressions in all columns. Standard errors clustered by month $\times$ year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

we find, if anything, a negative correlation with the surprise component (columns 1 to 3), while the positive baseline relationship is driven entirely by the anticipation component (columns 4 to 6).[33] In columns 7 through 9 we include lags and leads of both variables simultaneously and confirm that only the news pressure related to anticipation is associated with the timing of EOs. Interestingly, when focusing on the relevant dimension of news pressure, i.e., that driven by predictable news, the coefficient on same-day news pressure also becomes statistically significant, though generally smaller and less precisely estimated than the one on next-day news pressure. A standard deviation increase in the anticipation component of next-day news pressure ($\approx$ 2.6min) is associated with a 12% increase in the likelihood that an EO is signed that day, and a 7% increase in the likelihood that an EO is signed in the next day.

In Figure 9 we plot the coefficients on all the lags and leads of the two news pressure components (corresponding to column (9) in Table 6). It is clear that, when focusing on news

[33]Since each news-segment is classified into either the surprise or anticipation category (or neither), the two components of news pressure are mechanically negatively correlated. This likely explains the negative coefficient on surprise news pressure.

Table 5: Placebo: Earthquakes, Mass Shootings and Terror Attacks

	First Stage			
	(1) NP	(2) NP	(3) NP	(4) NP
Mass Shooting	0.129*			
	(0.066)			
Terrorist Attack		0.099***		
		(0.036)		
Earthquake			0.072**	
			(0.031)	
Earthquake or Shooting or Attack				0.075***
				(0.020)
Weeks in office	Yes	Yes	Yes	Yes
Y, M, DOW FEs	Yes	Yes	Yes	Yes
Observations	9411	9769	9039	8694
R-Squared	0.087	0.086	0.090	0.096

	Reduced Form				Second Stage			
	(1) EO	(2) EO	(3) EO	(4) EO	5) EO	(6) EO	(7) EO	(8) EO
Mass Shooting (t+1)	-0.036							
	(0.038)							
Terrorist Attack (t+1)		-0.039						
		(0.026)						
Earthquake (t+1)			0.003					
			(0.032)					
Earthquake or Shooting or Attack (t+1)				-0.017				
				(0.019)				
NP (t+1)					-0.280	-0.374	0.037	-0.222
					(0.307)	(0.320)	(0.436)	(0.269)
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Y, M, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
First Stage F-stat.					3.76	7.71	5.50	13.58
Observations	9412	9769	9039	8695	9411	9768	9038	8694
R-Squared	0.040	0.041	0.041	0.040	0.069	0.031	0.137	0.093

Sample: divided government. Dependent variable: indicator for EO signing. The table shows results of using an indicator for the occurrence of *unexpected* events – mass shootings, terrorist attacks and earthquakes – as instruments for news pressure. Columns (1) to (4): first stage, OLS. Columns (5) to (8): second stage, 2SLS. Columns (9) to (12): reduced form, OLS. Standard errors clustered by month $\times$ year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

related to anticipation, the estimated effect of news pressure on the timing of EOs becomes more precise.

2.5.3 Time in the Electoral Cycle and Popularity

In Table A10 we examine whether the relationship between the timing of EOs and next-day news pressure varies over the electoral cycle or depending on the president's popularity.

Interestingly, we find no evidence of strategic timing in the first 100 days of the presidential term (column 1) - a period in which EOs are commonly used to address issues raised during the campaign that the president has little incentive to conceal. The correlation with news pressure is instead more pronounced in periods of high disapproval - i.e., when the average disapproval

Figure 8: News Pressure on Days with and without Unexpected Events

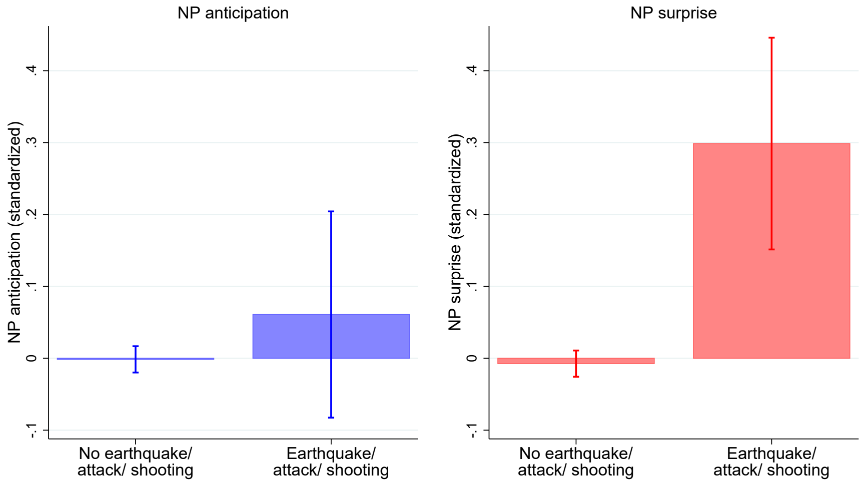

Mean levels of surprise and anticipation news pressure (standardized), along with 95% confidence intervals, on days with major unexpected events – earthquakes, terror attacks or mass shootings – vs days with no such events.

Table 6: Decomposition by News Sentiment

	NP: Surprise sentiment			NP: Anticipation sentiment			Both		
	(1) EO	(2) EO	(3) EO	(4) EO	(5) EO	(6) EO	(7) EO	(8) EO	(9) EO
NP *surpr.*	-0.008	-0.012	-0.006				-0.001	-0.004	0.001
	(0.015)	(0.015)	(0.015)				(0.015)	(0.015)	(0.016)
NP *surp.* (t+1)	-0.024	-0.030*	-0.034**				-0.011	-0.016	-0.020
	(0.015)	(0.015)	(0.016)				(0.015)	(0.016)	(0.017)
NP *surp.* (t-1)		-0.011	-0.008					-0.009	-0.007
		(0.015)	(0.016)					(0.016)	(0.016)
NP *anticip.*				0.022*	0.028**	0.031**	0.022*	0.027*	0.031**
				(0.013)	(0.014)	(0.015)	(0.013)	(0.014)	(0.015)
NP *anticip.* (t+1)				0.047***	0.049***	0.055***	0.045***	0.047***	0.052***
				(0.014)	(0.015)	(0.016)	(0.014)	(0.015)	(0.016)
NP *anticip.* (t-1)					0.005	0.001		0.002	-0.001
					(0.014)	(0.015)		(0.015)	(0.015)
7 lags of NP	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
7 leads of NP	No	No	Yes	No	No	Yes	No	No	Yes
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Y, M, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	9967	9416	9026	9967	9416	9026	9967	9416	9026
R-Squared	0.041	0.042	0.044	0.043	0.044	0.045	0.043	0.044	0.047

Sample: divided government. Dependent variable: indicator for EO signing. OLS regressions in all columns. Columns (1) to (3): Regressions on news pressure from segments associated with surprise, and its leads and lags. Columns (4) to (6): Regressions on news pressure from segments associated with anticipation, and its leads and lags. Columns (7) to (9): Regressions including both measures and their leads and lags simultaneously. Standard errors clustered by month × year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

rating over the previous month is higher than the median rating for the same president (column 4). We do not find any difference in timing depending on whether the president is a "lame-duck" (column 2), between first and second presidential terms (column 3), depending on the approval rating of Congress (column 5), or between election years non-election years (columns 5 and 6).

Figure 9: Decomposition by News Sentiment: Leads and Lags

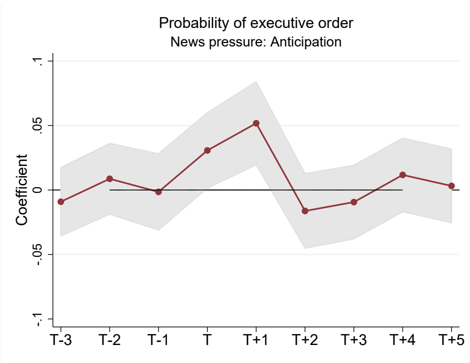

(a) NP *anticipation*

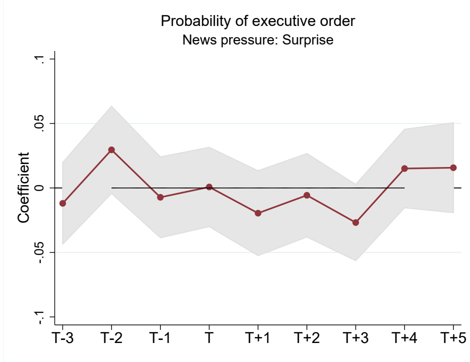

(b) NP *surprise*

Coefficients on leads and lags of the anticipation vs surprise component of news pressure – corresponding to column (9) in Table 6.

Finally, in Table **??** we estimate our baseline specification separately for different administration and for Republican and Democratic presidents. Our results indicate that no administration or party alone is driving the results.

2.6 Mechanisms

2.6.1 Same-Day vs. Next-Day News Coverage

The results discussed above indicate a significant relationship between the likelihood of EO signing and next-day news pressure, while evidence of a similar relationship with same-day news pressure is weaker. To interpret these results, it is important to better understand why presidents may be more concerned with minimizing next-day coverage of EOs rather than same-day coverage.

One potential explanation is that stories about EOs are more likely to be featured with a one-day lag due to a delay in news gathering. This hypothesis does not find support in the data – EOs receive twice as much airtime on the day they are issued than on the following day (left hand-side panel of Figure 10). Interestingly, however, conditional on EOs getting covered in the news, next-day coverage is on average longer (right hand-side panel).

An alternative explanation is that coverage of EOs may be *qualitatively* different between same and next day. For instance, shorter (though more frequent) news on the same day may provide basic information, while a lag of one day may allow reporters to produce in-depth analysis of the policy and to gather other, possibly critical, reactions.

As a simple test of this hypothesis, in Figure 11 we examine how same- and next-day news on EOs differ with respect to the frequency with which they mention reactions from Congress. We first analyze the headlines and transcripts of 84 VNA segments that contain the phrase "executive order" or synonyms and were aired on the day of, or one day after, an EO signing. We look for mentions of words with the root "Congress", or other related words such as "Senate", "House" (but not "White House"), "representative", and "speaker". With both measures the share of news segments mentioning Congress is significantly higher in next-day news (40 to 50% of segments) than in same-day news (20 to 22% of segments).

Figure 10: Media Coverage of EOs in Same- vs. Next-Day

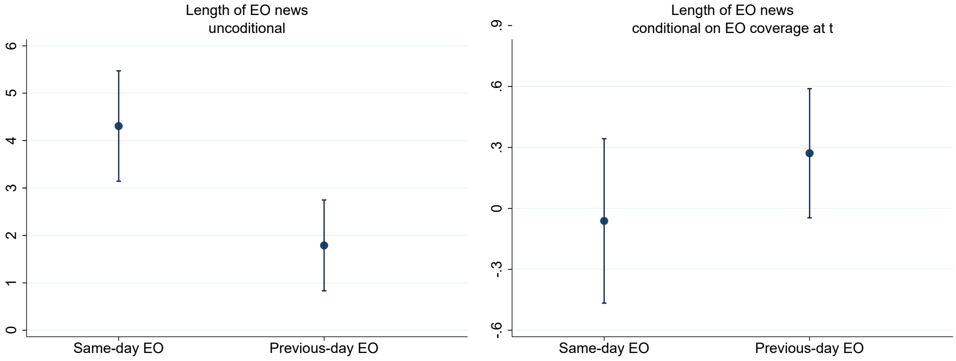

Coefficients from negative binomial maximum likelihood regressions of length of EO-news on an indicator for same-day EO and an indicator for previous-day EO, conditional on baseline controls. Standard errors clustered by year × month.

Figure 11: Mentions of Congress in the Text of Same- vs Next-Day TV News Segments

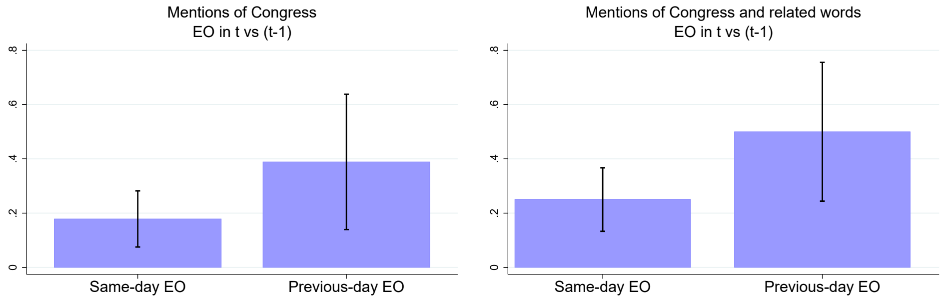

(a) Share of EO-related news segments mentioning Congress: VNA

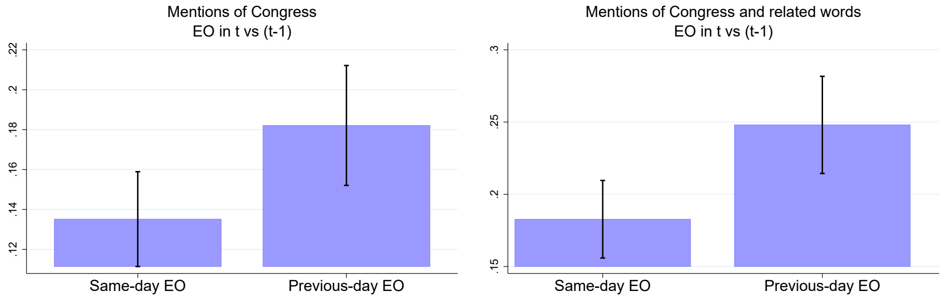

(b) Share of EO-related news segments mentioning Congress: GDELT TV Archive

Mean share of EO-related TV segments mentioning Congress (along with 95% confidence intervals). Panel (a): text of headlines and descriptions of VNA segments (covering 1979-2016). Panel (b): text of snippets of GDELT TV Archive segments (covering 2009-2016).

To validate these findings in a larger sample, we replicate the exercise using data from the *GDELT TV Archive*, which are more detailed and include a much larger number of news, though limited to the post-2009 period (see Appendix 2.7). We perform the same automated keyword search described above on the transcripts of 1497 15-second-long segments. The results, presented in panel b, are consistent with those found for the VNA sample: next-day news are significantly more likely to mention Congress than same-day news.

To evaluate more qualitative aspects of news coverage we instruct research analysts to watch and rate the GDELTs segments following a questionnaire (Table A11). The results, described in Appendix 2.7 confirm that next-day news coverage is significantly more likely to take the perspective of Congress, rather than that of the White House. We also find that the tone of Congress reactions is on average negative (in line with our expectation given our focus on divided government). In result, the distribution of tone to skews more negative for next-day coverage compared to same-day coverage.

2.6.2 Using Exogenous Events vs. Producing Distracting News

The results presented thus far are consistent with more controversial EOs being timed strategically to newsworthy events that are exogenous from the standpoint of the policy-maker.[34] However, our findings are also consistent with an alternative hypothesis, i.e., that the distracting news may, themselves, be deliberately induced by the policy-maker. Although separating these two mechanisms is beyond the scope of this paper, in Table 7 and the corresponding Figure 7 we attempt to provide some *prima facie* evidence in this regard by splitting the anticipated component of news pressure into news that mention the incumbent president (15% of the total), and news that do not. The results suggest that EO signings are strongly correlated with both components. One standard deviation increase in the president-related component of next-day (anticipated) news pressure ($\approx$ 2.4min) is associated with a 12% increase in the likelihood of an EO signing, while a standard deviation increase in the remaining component ($\approx$ 2.2min) is associated with a 14% increase in this likelihood.

Figure 12: Decomposition by News Related to President vs Other News: Leads and Lags

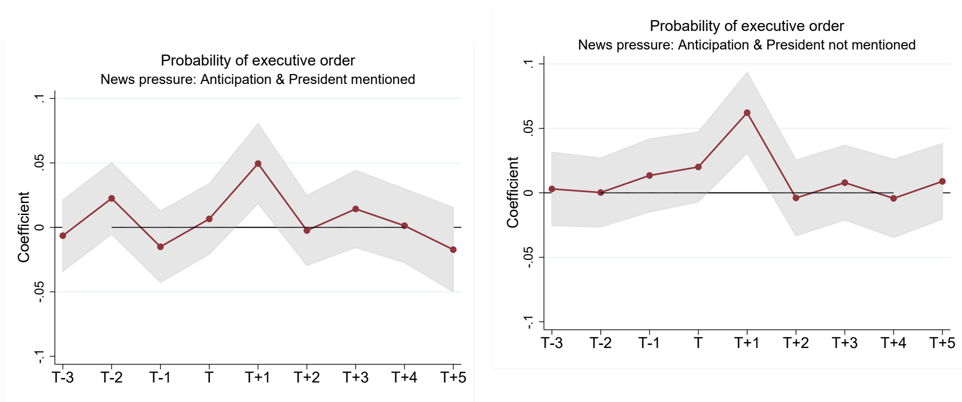

(a) NP *anticip.* & president mentioned (b) NP *anticip.* & president not mentioned

Coefficients on leads and lags of the president-related vs president-unrelated component of anticipation news pressure – corresponding to column (9) in Table 7.

2.7 Conclusion

In this paper we investigate whether politicians strategically choose to implement policies in coincidence with other important events, so as to minimize media coverage and public scrutiny

[34]This conceptual framework is analogous to that used by Durante and Zhuravskaya (2018), who assume that the Israeli army cannot influence the U.S. news cycle and take it as given when deciding on when to carry out attacks.

Table 7: News Related to President vs Other News

	NP: Anticipation & President mentioned			NP: Anticipation & President not mentioned			Both		
	(1) EO	(2) EO	(3) EO	(4) EO	(5) EO	(6) EO	(7) EO	(8) EO	(9) EO
NP *president*	0.002	0.005	0.003				0.003	0.007	0.007
	(0.013)	(0.013)	(0.014)				(0.013)	(0.014)	(0.014)
NP *president*(t+1)	0.037***	0.040***	0.038**				0.047***	0.051***	0.049***
	(0.014)	(0.014)	(0.015)				(0.014)	(0.015)	(0.016)
NP *president* (t-1)		-0.016	-0.017					-0.014	-0.015
		(0.014)	(0.014)					(0.014)	(0.014)
NP *other news*				0.019	0.018	0.018	0.018	0.020	0.020
				(0.013)	(0.014)	(0.013)	(0.014)	(0.014)	(0.014)
NP *other news* (t+1)				0.051***	0.050***	0.051***	0.061***	0.062***	0.062***
				(0.015)	(0.015)	(0.015)	(0.015)	(0.016)	(0.016)
NP *other news* (t-1)					0.018	0.017		0.014	0.013
					(0.014)	(0.014)		(0.014)	(0.014)
7 lags of NP	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
7 leads of NP	No	No	Yes	No	No	Yes	No	No	Yes
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Y, M, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10128	10121	10109	10133	10126	10114	10128	10121	10109
R-Squared	0.041	0.042	0.042	0.042	0.042	0.043	0.043	0.045	0.045

Sample: divided government. Dependent variable: indicator for EO signing. OLS in all columns. Columns (1) to (3): Regressions on news pressure from segments associated with anticipation that mention the name of the incumbent president, and its leads and lags. Columns (4) to (6): Regressions on news pressure from segments associated with anticipation that *don't* mention the name of the incumbent president, and its leads and lags. Columns (7) to (9): Regressions including both measures and their leads and lags simultaneously. Standard errors clustered by month × year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

of their actions. To shed light on this general question, we analyze the timing of the signing of executive orders by U.S. presidents over the past four decades, and its relationship with the new cycle.

We show that executive orders are disproportionately likely to be signed on the eve of days when the news cycle is dominated by other events. This relationship only holds during periods of divided government - when the presence of a hostile Congress increases the president's incentive to conceal controversial unilateral actions - and only for categories of EOs that are likely to make the news and to generate criticism. Furthermore, EO-signings tend to coincide with predictable news but not with surprising ones, and appear to be timed to minimize next-day coverage of EOs which, we document, is generally less favorable to the president. This evidence is consistent with a forward-looking PR strategy aimed at minimizing negative publicity via distraction, and suggests that, even in the presence of a free press, strategic behavior by politicians can limit public scrutiny of government policies and political accountability.

While politicians may exploit distracting events occurring outside their control, it is also possible that they may actively try to influence the media agenda through their actions or statements so as to "create" distracting news. While our analysis only provides limited evidence as to which of these scenarios is more likely, this certainly represents an interesting venue for future research.

Finally, our research documents the strategic behavior of top level elected officials characterized by a high degree of sophistication and abundant PR resources. Another question for future research is whether this type of behavior may generalize to lower level politicians, and what might be the broader implications for political accountability.

Appendix

Data and descriptive statistics

Table A1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.	N
Divided gov.	0.73	0.444	0	1	13880
EO	0.1	0.3	0	1	13880
Num. EOs	0.119	0.403	0	11	13880
Any EO news	0.012	0.11	0	1	13880
Length of EO-news (in sec)	3.356	43.72	0	1640	13880
EO on government operations	0.214	0.41	0	1	1384
EO significance	0.094	0.8	-0.965	3.198	1001
EO covered in AP	0.651	0.477	0	1	980
EO-topic disgareement President–Congress (6mo)	0.453	0.434	0	1	708
EO-topic disagreement President–Congress (12mo)	0.471	0.414	0	1	778
NP (in 10s of min)	0.816	0.253	0.114	2.95	13878
NP from segments with anticipation sentiment	0.788	0.257	0	2.95	13772
NP from segments with surprise sentiment	0.235	0.201	0	2.65	13772
Google trends "executive order"	1.043	3.18	0	100	4743
Gallup share disapproving	41.272	11.771	6	71	5767
earthquake	0.01	0.1	0	1	12784
shooting	0.004	0.062	0	1	12423
terror	0.008	0.091	0	1	13514

Table A2: Periods of divided and unified government

Congress	Years	President	Senate	House
96	1979 - 1981	D	D	D
97	**1981 - 1983**	R	R	D
98	**1983 - 1985**	R	R	D
99	**1985 - 1987**	R	R	D
100	**1987 - 1989**	R	D	D
101	**1989 - 1991**	R	D	D
102	**1991 - 1993**	R	D	D
103	1993 - 1995	D	D	D
104	**1995 - 1997**	D	R	R
105	**1997 - 1999**	D	R	R
106	**1999 - 2001**	D	R	R
107	**2001 - 2003**	R	D	R
108	2003 - 2005	R	R	R
109	2005 - 2007	R	R	R
110	**2007 - 2009**	R	D	D
111	2009 - 2011	D	D	D
112	**2011 - 2013**	D	D	R
113	**2013 - 2015**	D	D	R
114	**2015 - 2017**	D	R	R

Composition of Congress and White House control. Periods of divided government highlighted in bold.

Table A3: Coding of EO-Subject Specific Keywords

EO number	EO Description	Keyword -- tfidf	Keywords -- Manually coded
13280	responsibilities of the department of agriculture and the agency for international development with respect to faith-based and community initiatives	agricultur agenc faithbas commun faithbas commun initi agenc intern	agricultur faith commun initi
13322	adjustments of certain rates of pay	pai rate schedul statutori pai pai system	adjust rate
13323	assignment of functions relating to arrivals in and departures from the united states	departur unit relat arriv arriv departur arriv citizen unit	arriv departur
12296	president's economic policy advisory board	presid econom polici advisori econom polici advisori board board	econom polici
12723	blocking kuwaiti government property	kuwait govern kuwait kuwaiti govern block kuwaiti kuwaiti	block properti kuwait
12247	federal actions in the lake tahoe region	region lake taho taho region taho lake	lake taho
12266	food security wheat reserve	wheat secur wheat food secur wheat reserv reserv	wheat secur
12947	prohibiting transactions with terrorists who threaten to disrupt the middle east peace process	threaten disrupt peac process terrorist threaten east peac disrupt middl	prohibit transact terrorist middl east
13188	amendment to executive order 13111, extension of the advisory committee on expanding training opportunities	committe expand expand train extens advisori train opportun execut extens	technolog train
12242	synthetic fuels	synthet fuel synthet guarante rate substanti substanti term	synthet fuel fuel

Examples illustrating the coding of EO-subject specific keywords. (Stemmed) keywords coded automatically from the full text of each EO based on a tf-idf criterion are reported in the third column. (Stemmed) keywords coded manually based on EO summary reported in the fourth column.

Table A4: Construction of News Pressure: Examples

Executive Order # 13505 (**March 9 2009**) Removing Barriers to Responsible Scientific Research Involving Human Stem Cells

Keywords: **stem cells, research**.

Date	Network	N	Headline	Length (secs)	NP
8Mar2009	NBC	1	Economy: The Problems, The Politicians	200	Length of top 3 non-EO stories, adjusted to the total length of non-EO broadcast = (200+120+120) * 1200 / (1200 - 0)
8Mar2009	NBC	2	Afghanistan And Iraq Wars / Troops	120	
8Mar2009	NBC	3	Maryville, Illinois / Church Shooting	120	
8Mar2009	NBC	4	Madoff Fraud Case	150	
8Mar2009	NBC	5	Winter Weather / Storms	20	
8Mar2009	NBC	6	Airlines / Cheap Tickets	120	
8Mar2009	NBC	7	Seeking Solutions (Extended Families)	140	
8Mar2009	NBC	8	Economy: Road Work / Highway Trust	140	
8Mar2009	NBC	9	Kennedy Honors	40	
8Mar2009	NBC	10	Economy: Treasure Hunt/ Scrounging	140	
8Mar2009	NBC		Good Night	10	
			total:	1200	440

No news related to EOs or mentioning EO-keywords.

Date	Network	N	Headline	Length (secs)	NP
9Mar2009	CBS	1	**Executive Order / Stem Cell Research**	340	
9Mar2009	CBS	2	Supreme Court / Gun Companies	20	Length of top 3 non-EO stories, adjusted to the total length of non-EO broadcast = (20+120+30) * 1160 / (1160-340)
9Mar2009	CBS	3	Phoenix, Arizona / Drug War / Firearms Trafficking	120	
9Mar2009	CBS	4	Maryville, Illinois / Church Shooting	30	
9Mar2009	CBS	5	Auto Industry / Ford And Uaw / Bailout	160	
9Mar2009	CBS	6	Economy: Recession / Buffett'S Warning	20	
9Mar2009	CBS	7	Religion: Losing The Faith	130	
9Mar2009	CBS	8	China / Ships	20	
9Mar2009	CBS	9	Hitting Home (College Costs)	160	
9Mar2009	CBS	10	Barbie At 50	160	
9Mar2009	CBS		Good Night	10	
			total:	1160	240.5

News related to EOs or mentioning EO-keywords in the top 3.

Date	Network	N	Headline	Length (secs)	NP
9Mar2009	NBC	1	Economy: Global Recession / Buffett	210	Length of top 3 non-EO stories, adjusted to the total length of non-EO broadcast = (210+130+150) * 1150 / (1150-160)
9Mar2009	NBC	2	Economy: Homelessness / Sacramento, California	130	
9Mar2009	NBC	3	Japan / Auto Industry / Toyota	150	
9Mar2009	NBC	4	China-Us Relations / Us Ship	40	
9Mar2009	NBC	5	Medicine: **Stem Cell Research** / Policy	160	
9Mar2009	NBC	6	Religion Survey	140	
9Mar2009	NBC	7	Britain / Shakespeare Portrait	30	
9Mar2009	NBC	8	Medicine: Depression And Heart Disease	20	
9Mar2009	NBC	9	Medicine: Migraines	30	
9Mar2009	NBC	10	Making A Difference/Acts Of Kindness	100	
9Mar2009	NBC	11	Making A Difference (Same Café)	140	
9Mar2009	NBC		Good Night	10	
			total:	1150	569.2

News related to EOs or mentioning EO-keywords outside the top 3.

EO characteristics under divided vs unified government

Given the centrality of the distinction between unified and divided government for our analysis, it is important to shed light on how the president's use of EOs differs between these two situations. To this end, in Figure A1 we plot coefficients from uni-variate regressions of various EO-characteristics on a dummy for divided government. We test for differences in the frequency of EOs, their significance, how often they fall in the category of government operations, how frequently they get covered by AP, and how often they concern topics of prior disagreement between president and Congress. We standardize each variable to facilitate comparison of the magnitude of the differences.

Overall, EOs issued in periods of divided and unified government are largely balanced along most dimensions, particularly with regard to their frequency and their topic. The exceptions are small differences in AP coverage and significance, and a sizeable difference in the likelihood of being on a topic of prior disagreement between president and Congress – EOs issued under divided government are one standard deviation more likely to concern issues on which the president's and Congress views are not aligned.

Figure A1: EO Characteristics in Periods of Divided vs Unified Government

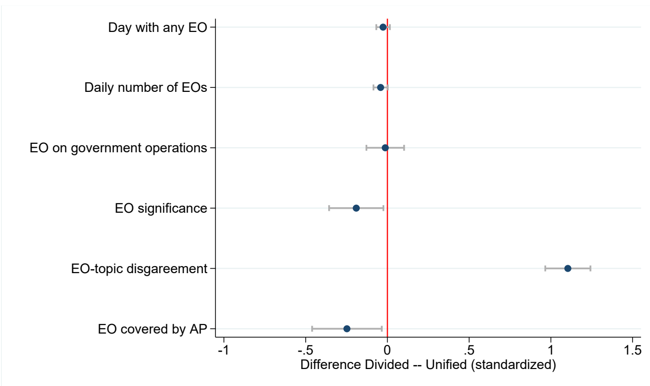

Coefficients from uni-variate regressions of standardized EO-characteristics on a dummy for divided (as opposed to unified) government. Standard errors clustered by year×month.

Robustness

Table A5: Robustness: Alternative Specifications

	(1) EO	(2) EO	(3) EO	(4) EO	(5) EO
NP	0.003	0.003	0.003	0.003	0.006
	(0.015)	(0.014)	(0.015)	(0.016)	(0.014)
NP (t+1)	0.045***	0.045***	0.045***	0.044***	0.042**
	(0.015)	(0.016)	(0.013)	(0.015)	(0.016)
NP (t-1)	-0.010	-0.010	-0.010	-0.011	-0.009
	(0.016)	(0.014)	(0.013)	(0.016)	(0.014)
7 lags of NP	Yes	Yes	Yes	Yes	Yes
Year, Month and DOW FEs	Yes	Yes	Yes	Yes	Yes
Weeks in office: linear	Yes	Yes	Yes	No	No
Weeks in office: President-specific	No	No	No	Yes	No
Weeks in office: FEs	No	No	No	No	Yes
SEs	Newey-West	CL(year)	CL(Congress)	CL(y $\times$ m)	CL(y $\times$ m)
Observations	10126	10126	10126	10126	10126
R-Squared	-	0.042	0.042	0.044	0.068

Sample: divided government. Dependent variable: indicator for EO signing. OLS regressions in all columns. Each column replicates our baseline specification (column 6 of Table 3), with the following modifications. Column (1): Newey-West standard errors. Column (2): Standard errors clustered by year instead of year $\times$ month. Column (3): Standard errors clustered by Congressional term instead if year $\times$ month. Column (4): Controlling for week-in-office $\times$ President FEs. Column (5): Controlling for week-in-office FEs. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Robustness: Alternative Specifications, Continued

	(1) EO	(2) EO	(3) EO	(4) EO	(5) EO
NP	0.004	0.018	0.001	0.004	0.004
	(0.015)	(0.016)	(0.016)	(0.016)	(0.016)
NP (t+1)	0.042***	0.032**	0.042***	0.045***	0.046***
	(0.014)	(0.015)	(0.016)	(0.015)	(0.015)
NP (t-1)	-0.008	0.000	-0.012	-0.010	-0.009
	(0.016)	(0.016)	(0.017)	(0.016)	(0.016)
Year $\times$ Month FEs	No	No	Yes	No	No
7 lags of EO	No	No	No	Yes	No
Holidays, Days Abroad	No	No	No	No	Yes
7 lags of NP	Yes	Yes	Yes	Yes	Yes
FEs & Weeks in office	Yes	No	Yes	Yes	Yes
Model	Probit	OLS	OLS	OLS	OLS
Observations	10124	10126	10126	10126	10126
(Pseudo) R-Squared	0.082	0.002	0.065	0.043	0.047

Sample: divided government. Dependent variable: indicator for EO signing. OLS regressions in all columns except for column (1).

Each column replicates our baseline specification (column 6 of Table 3), with the following modifications. Column (1): Probit instead of linear probability model. Column (2): No calendar controls. Column (3): Controlling for year $\times$ month FEs. Column (4): Controlling for lagged of EO signings. Column (5): Controlling for holidays and Presidential foreign visits.

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Robustness: Alternative Definitions of News Pressure

	Steps in NP Construction			Other Variants of NP		
	(1)	(2)	(3)	(4)	(5)	(6)
	Uncorr. excl. EO-news	Excl. EO-news	Excl. EO-news + kw's	Longest segments	Kw's from tf-idf	Excl. kw's in +/-7 days
	EO	EO	EO	EO	EO	EO
NP	0.009	0.010	0.003	0.034*	-0.004	0.002
	(0.016)	(0.016)	(0.016)	(0.018)	(0.015)	(0.015)
NP (t+1)	0.028*	0.030**	0.045***	0.073***	0.038**	0.036**
	(0.015)	(0.015)	(0.015)	(0.018)	(0.015)	(0.015)
NP (t-1)	-0.020	-0.021	-0.010	0.004	-0.006	-0.017
	(0.016)	(0.016)	(0.016)	(0.019)	(0.016)	(0.016)
7 lags of NP	Yes	Yes	Yes	Yes	Yes	Yes
FEs & Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10126	10126	10126	10117	10117	10117
R-Squared	0.041	0.041	0.042	0.045	0.041	0.042

Sample: divided government. Dependent variable: indicator for EO signing.
Each column replicates our baseline specification (column 6 of Table 3), introducing one step of our procedure for the construction of news pressure at a time (columns 1 to 3), or modifying news-pressure (columns 4 to 6).
Column (1): NP calculated excluding only segments that refer to EOs explicitly, without adjustment for total length of the newscast. Column (2): adding the step of adjustment for total length of the newscast. Column (3): adding the step of excluding and adjusting for segments containing EO-subject specific keywords, thus obtaining our baseline measure.
Column (4): NP calculated using top 3 news segments ranked by length rather than order. Column (5): NP calculated excluding EO-subject specific keywords coded automatically based on a tf-idf criterion, rather than manually coded. Column (6): NP calculated excluding segments containing EO-subject specific keywords aired with +/-7 days from signing of the respective EO, rather than aired within +/-1 day.
Standard errors clustered by month $\times$ year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Additional heterogeneity results

Table A8: News Coverage by Type of EO

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Any EO news				Length EO news			
EO in t or (r t-1)	0.005	0.000	0.003	0.004	1.019	0.300	0.535	4.055***
	(0.006)	(0.003)	(0.003)	(0.004)	(0.782)	(0.692)	(0.886)	(0.709)
EO in t or (t-1)								
× Not gov. operations	0.010				2.322***			
	(0.006)				(0.800)			
EO in t or (t-1)								
× High significance		0.025***				4.995***		
		(0.006)				(1.113)		
EO in t or (t-1)								
× High disagreement			0.022**				4.518***	
			(0.009)				(1.199)	
EO in t or (t-1)								
× Covered by AP				0.020***				-0.141
				(0.007)				(0.745)
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year, Month, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	13880	9131	10929	10602	13880	9131	10929	10602
(Pseudo)-R2	0.014	0.012	0.009	0.019	0.018	0.019	0.022	0.025

Dependent variable: indicator for any EO-related news in columns (1) to (4), and length of EO-related airtime in columns (5) to (8). Each column presents an interaction of an indicator for EO signed on day t or (t-1), with an indicator for whether this EO (or at least one in case of multiple EOs) is of a certain type. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: News Pressure and the Timing of EOs: One vs Both Chambers of Congress Against President

	One chamber against president			Both chambers against president		
	(1)	(2)	(3)	(4)	(5)	(6)
	EO	EO	EO	EO	EO	EO
NP	-0.015	-0.020	-0.019	0.009	0.019	0.018
	(0.021)	(0.024)	(0.024)	(0.019)	(0.021)	(0.021)
NP (t+1)	0.035	0.035	0.038*	0.048**	0.053**	0.049**
	(0.021)	(0.021)	(0.022)	(0.020)	(0.022)	(0.024)
NP (t-1)		0.026	0.025		-0.032	-0.033
		(0.025)	(0.024)		(0.022)	(0.022)
7 lags of NP	No	Yes	Yes	No	Yes	Yes
7 leads of NP	No	No	Yes	No	No	Yes
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes
Year, Month, DOW FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4363	4356	4350	5894	5894	5888
R-Squared	0.047	0.048	0.050	0.041	0.042	0.042
Mean dependent variable	0.100	0.100	0.100	0.096	0.096	0.096

Sample: divided government with both chambers against president columns (1) to (3), divided government with one chamber against president in columns (4) to (6). Dependent variable: indicator for the signing of an EO. OLS regressions in all columns. Standard errors clustered by month $\times$ year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A10: Interactions with the Electoral Cycle and Popularity

	(1) EO	(2) EO	(3) EO	(4) EO	(5) EO	(6) EO
NP (t+1)	0.051***	0.047***	0.028	0.016	0.049**	0.048**
	(0.016)	(0.015)	(0.020)	(0.020)	(0.024)	(0.019)
NP(t+1) × First 100 days	-0.157**					
	(0.063)					
First 100 days	0.116**					
	(0.051)					
NP(t+1) × Lame-duck		-0.032				
		(0.066)				
Lame-duck		0.073				
		(0.050)				
NP(t+1) × 2nd term			0.034			
			(0.025)			
2nd term			0.054*			
			(0.028)			
NP(t+1) × Disapproval > median				0.052**		
				(0.025)		
Disapproval > median				-0.036*		
				(0.022)		
NP(t+1) × Disapproval Congress > median					-0.009	
					(0.029)	
Disapproval Congress > median					0.007	
					(0.025)	
NP(t+1) × Presidential election year						-0.017
						(0.034)
Presidential election year						0.013
						(0.031)
NP(t+1) × Midterm election						0.003
						(0.031)
Midterm election year						-0.027
						(0.028)
NP and 7 lags of NP	Yes	Yes	Yes	Yes	Yes	Yes
Weeks in office	Yes	Yes	Yes	Yes	Yes	Yes
Year, Month, Day-of-Week FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	10126	10126	10126	10098	6847	10126
R-Squared	0.043	0.042	0.043	0.042	0.042	0.042

Sample: divided government. Dependent variable: indicator for EO signing. The table shows the coefficients on interactions of news pressure with various indicators related to the electoral cycle. Standard errors clustered by month × year. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Content of EO-related news coverage

In this section we use the GDELT TV Television explorer as a complementary source of data on news coverage to that provided by VNA.[35] Though these data are only available starting 2009, they have at least three important advantages: i) they cover a broader set of networks, ii) they cover all news-related shows, not just evening news, iii) they include the full transcripts of newscasts, not just summaries. We focus on the main news networks operating in and after 2009, i.e., ABC, CBS, NBC, CNN, MSNBC and Fox News, and on the prime time + fringe time slots, i.e., between 4pm and 12am. We assess the presence and length of EO-related news using the same procedure described for the VNA data. The GDELT TV data are organized in segments of 15-seconds; overall, our sample includes 1,497 EO-related segments.

To evaluate the tone of EO-related news in periods of divided government, we ask research analysts to watch each GDELT segment in the broader context of the newscast and to code its content along several dimensions following a questionnaire. We ask whether the news segment covers a specific EO signed on the same or previous day, whether it features statements and reactions from various actors, including the president, Congress, the judiciary, NGOs or citizens, and, finally, to assess the overall tone of the segment towards the president. Table A11 presents the full questionnaire and summary statistics for the responses.

Out of the 1324 videos aired under divided government, the analysts deemed 353 to be directly related to a specific EO signed on the same or the previous day, separating them from other news discussing EOs or presidential powers in general or talking about EOs signed further in the past or planned for the future. The results of the video analysis are summarized in Figure **??**. First, same-day coverage is significantly more likely to cover the perspective of the president - featuring the signing ceremony or official statements by the White House (panels a and b), while next-day coverage is significantly more likely to feature the reaction of Congress (panel c). No significant difference emerges between same-day and next-day news with respect to the probability of reporting the reactions of NGOs or ordinary citizens (panel d), or in any of the other dimension captured in our questionnaire. That the results based on the video analysis of newscasts are consistent with - and generally stronger than - those based on the analysis of the transcripts is reassuring since the latter approach is more prone to measurement error and the risk of misclassification.

Finally, we analyze differences between same-day and next day coverage of EOs with respect to the tone. To this end, we asked analysts to code, for each relevant news segment, the general tone used towards the president (on a five-point scale from very praising to very critical), and specifically the tone of Congressional reactions to EOs (as positive, negative or neutral). Figure **??** summarizes the results. Panel a reports the distribution of news segments by overall tone, separately for same-day and next-day news. A clear pattern emerges: while on average the tone of coverage is rather neutral, next-day news are characterized by less praise and more criticism of the president's actions. Regarding the tone of Congressional reactions, presented in panel b, we find that they are on average negative, which has to be expected under divided government, with a mean rating of 2.4 on a 1 to 3 negativity scale. Interestingly, conditional on Congress reactions being covered, we find virtually no difference in the tone of Congress' reactions between next-day and same-day news. This suggests that the difference in overall tone towards the president documented in panel a may be driven by the fact that next-day news more often features Congress reactions (and less often features the White House perspective), rather than by a difference in the nature of these reactions. In Table **??** we further test for the differences in tone between same-day and next-day coverage estimating OLS

[35] https://api.gdeltproject.org/api/v2/summary/summary?DATASET=IATV

regressions controlling for network fixed effects finding consistent effects.

Table A11: Questionnaire on EO-News Content

#	Question	Percent "Yes"
1	Does the newscast focus on a particular executive order? *(Proceed if "Yes")*	27%
2	Is the content of the executive clearly summarized?	77%
3	Was the executive order signed on the day of the newscast?	70%
4	Was the executive order signed on the day before the newscast?	22%
5	Does the newscast show footage from an executive order signing ceremony?	22%
6	Does the newscast include an interview with/ a statement by the President or a White House representative?	50%
7	Does the newscast discuss the reaction of Congress to the executive order?	18%
8	Does the newscast discuss the reaction of members of the judiciary to the executive order?	0%
9	Does the newscast discuss the reaction of any other government officials to the executive order (aside from Congress/Judiciary)?	11%
10	Does the newscast discuss the reaction of citizens/ non-governmental organizations to the executive order?	11%
11	Does the newscast question whether the executive order is within the constitutional authority of the President?	2%
12	Does the newscast mention past attempts of the President to pass legislation on the same issue through Congress?	4%
13	Overall, how praising/ critical of the President is the newscast, on a scale from 1 (very praising) to 5 (very critical)?	mean = 3.1

Chapter 3

THE IMPACT OF ONLINE COMPETITION ON LOCAL NEWSPAPERS

Joint with Ruben Durante (ICREA, UPF, Barcelona GSE, IPEG and CEPR) and Gregory J. Martin (Stanford GSB)[1]

3.1 Introduction

The Internet has profoundly changed the environment in which newspapers operate. Competition from online platforms has contributed to the sharp decline in newspapers' advertising revenues over the last two decades, forcing many news outlets to drastically rethink their business model and organization.[2] These changes, some worry, 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 the key role played by local newspapers in informing citizens about local politics (Gentzkow et al. 2011; Snyder and Strömberg 2010), they may also have important implications for electoral politics.

Despite the potentially grave consequences of these transformations for the future of journalism, rigorous evidence on the impact of online competition 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.

To overcome this limitation, in this paper we investigate the impact of the introduction of *Craigslist* (henceforth CL), the world's largest online platform for classified ads, on newspapers in the US. CL disrupted the market for classified ads, a formerly lucrative niche for newspapers (**?**Kroft and Pope 2014).[3] Tracking the expansion of CL across US counties between 1995 and 2009, we examine how the entry of a local CL website affected the organization and editorial

[1]David Ampudia Vicente, Marco Lo Faso, and Elliot Motte provided excellent research assistance. Irena Djourelova, Anna Griñó, Adriana Oliveres and Silvia Terzulli provided help with digitizing the Editor and Publisher data.

[2]According to data form the News Media Alliance (formerly Newspaper Association of America), US newspapers' advertising revenues fell from US$49 billion in 2000 to US$26 billion in 2010.

[3]As of 2000, classified advertising accounted for 40% of US newspapers' total advertising revenues. In that year advertising revenues amounted to US$49 billion, compared to US$11 billion of circulation revenues.

decisions of local newspapers, and ultimately, the electoral choices of local voters.

The expansion of CL across 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 each local CL website, 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. In addition, CL's narrow focus on classified ads provides an important source of heterogeneity, since the entry of CL should disproportionately affect newspapers that relied more heavily on classified ads *ex ante*. Finally, with 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[4] alleviates concerns that the timing of CL's entry might have been driven by strategic considerations related to the conditions of local newspaper markets. Indeed, 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.

Our empirical strategy compares the evolution of outcomes of interest between areas with and without access to a local CL website, before and after such a website is introduced. Throughout the analysis we control for the quality of local broadband Internet and for a range of demographic and socio-economic covariates. In addition, we exploit variation 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 prior to the entry of CL.[5]

To analyze the impact of CL's entry on local newspapers, we exploit new comprehensive data on the organization of over 1,500 newspapers, covering the period from 1995 to 2010. We find that while CL's entry does not significantly affect the number of active local newspapers, it does lead to substantial downsizing. After CL's entry into a county, newspapers headquartered there cut about 1.2 editor-level jobs on average. This effect is driven by newspapers that relied more heavily on classified ads at baseline: for this group the effect amounts to a 14% decline in the number of jobs relative to the mean.[6] Staff cuts affect both managerial and editorial positions, but editorial cuts appear to disproportionately affect editors responsible for the coverage of politics, leaving other areas such as sports and entertainment largely unaffected.

We then test to what extent these organizational changes affected newspapers' editorial priorities, with particular regard to their coverage of politics. First, applying keyword searches to the entire corpus of articles published in over 800 newspapers, we compute the number of mentions of local representatives in Congress. We document that, following the entry of CL in a given area, news coverage of local representatives in affected papers declined significantly, by around 30%. We complement this approach with a semi-supervised topic model, estimated on a random sample of two million articles. We find evidence of a decline in the prevalence of topics related to politics, while we find no significant effect on other topics such as sports, entertainment, or crime.

Next we examine how readers responded to these changes in content. First, we document

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

[5]To corroborate the validity of this proxy, we document that: i) the presence of classified ads managers is strongly correlated with the share of pages devoted to classified ads prior to the entry of CL, and ii) the space devoted to classified ads decreased disproportionately for papers with classified managers after the entry of CL (see Appendix Tables A1 and **??**.)

[6]We find no significant change in the number of pages published by affected newspapers, which suggests that this downsizing is likely associated with greater workload per staff member.

that, in the years after the entry of CL, local newspapers experienced a sharp decline in circulation. We further explore this readership decline using data from two large-scale surveys on media consumption. The results confirm that, following the entry of CL, local respondents are less likely to report reading a local newspaper. Interestingly, this trend is primarily driven by readers that are relatively *less* likely to be interested in classifieds, and therefore cannot be merely due to lower demand for print classified ads.[7] Evidence from both survey and browsing data suggests that the decline in newspaper readership is not compensated by increased news consumption online or through other sources (i.e., national papers, radio or TV), and is hence likely to result in an overall decline in political information among the public.[8]

Finally, we study how reduced news coverage of politics affected the behavior of local voters, with particular regard to electoral participation and ideological polarization. For electoral participation we find that the entry of CL has a negative though not robust effect on turnout in presidential elections and no effect in mid-term elections. Regarding ideological polarization, we find robust evidence that the entry of CL favored the rise of ideologically extreme candidates, and reduced the probability that voters support different parties in concurrent elections (split-ticket voting). These findings are consistent with a greater tendency to rely on national partisan cues when less information about local politicians is available (Darr et al. 2018; Moskowitz 2021; Trussler 2020).

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 potentially detrimental effect for ideological polarization. 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 builds upon and expands prior evidence on the impact of CL on US newspapers. A study by **?** shows that, after the entry of CL, local newspapers experienced lower classified-ad rates and circulation, and higher display-ad rates and subscription prices. Our paper expands these findings by documenting the entire "chain of events" triggered by the entry of CL, and its profound implications for newspapers' staff, editorial priorities, news content, and, ultimately, political outcomes.[9]

Second, our results add to previous evidence on the impact of the introduction of new advertising technologies, or new media, on incumbent media outlets. In a long-term 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 paper studies the flip side of this question: what happens to newspapers' political coverage when advertising profits are competed away by new technology. Regarding the impact of new media, Angelucci et al. (2020) argue that the introduction of commercial TV in the US had a significant negative economic impact on newspapers, leading to lower coverage of local politics and a decline in split-ticket voting between national and local elections. Bhuller et al. (2020) document large declines in circulation and shifts in editorial priorities for newspapers in Norway in response to the roll-out of broadband Internet. Our empirical setting is different in that it allows us to separate the effect of a specific shock to the advertising market from the

[7] The readership decline is also unlikely to be due to newspapers charging higher subscription prices, as we find no significant effect of CL's entry on that margin.

[8] In particular, we find that the number of visits to popular national news websites by local users is not affected by CL's entry.

[9] Other work exploits the experiment created by CL's expansion across the US or particular design features of the platform to study questions related to matching efficiency in labor and housing markets (Kroft and Pope 2014), and the impact of online personal/erotic ads on sexually transmitted diseases and violence against women (Cunningham et al. 2019; Chan and Ghose 2014).

demand-side changes brought about by new communication technology.

More broadly, our paper relates to previous work on the impact of the Internet on political participation, electoral outcomes, and public policies (Falck et al. 2014; Campante et al. 2018; Gavazza et al. 2019; Larcinese and Miner 2018; Manacorda and Tesei 2020). Exploiting variation in access to broadband or wireless technology, these studies gauge the net effect of the various changes brought about by the Internet on both the demand and supply side of the political market. In this respect, the novelty of our paper is to isolate the effect of the Internet — and specifically of digital advertising platforms — on legacy media and its implications for content quality, readership, and political outcomes.

Finally, our paper is generally related to the literature on the link between media, information, voter participation, and political accountability. This includes the seminal studies by Besley and Burgess (2002a) and Snyder and Strömberg (2010) which document that media coverage of local politicians make votersâ more informed and participative and politicians more responsive to the demands of their constituents.[10] Another set of studies (Lelkes et al. 2017; Allcott et al. 2020; Levy 2021, e.g.,) examines the influence of media consumption on ideological polarization. Our paper contributes to this literature by demonstrating how news media support the functioning of elections as selection mechanisms. The magnitude of the electoral penalty that extreme candidates face (Hall 2015) 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 voters' ability to discriminate between candidates and thus weaken the ideologically moderating force that elections apply.

The remainder of the paper is organized as follows. Section 3.1 provides some background information about Craigslist and its expansion. Section 3.2 describes the data used in the analysis, while section 3.3 discusses the empirical strategy. Section 3.4 presents the results. Section 3.5 concludes.

Background

Craiglist.org (CL) is the world's largest online platform for classified ads. It was founded in San Francisco in 1995, and served only the Bay Area until 2000, when it began to gradually expand to other U.S. locations. CL initially opened new local websites in big cities such as Boston, New York, and Chicago. Over time, it expanded to smaller markets covering 115 locations in 2005, 331 in 2008, and 416 today.

Consistently ranked among the top 20 U.S. websites by traffic,[11] the CL platform has a simple layout which has remained largely unchanged over time (see Figure A3). Ads on CL websites are organized into sections including housing, jobs, items for sale, professional services, and personals. CL websites only host classified ads and do not include any display ads or news content. Ads on CL are generally posted free of charge, with a few exceptions for brokered apartment rentals in New York and job posts for employers in some major cities.[12] CL's business model reflects the unconventional views of its founder, Craig Newmark, who prior-

[10] Also related to our work is a recent study by Gao et al. (2020) which documents the negative impact of the closure of local newspapers on local public finances. The authors interpret these findings as evidence of the importance of local newspapers in monitoring local government.

[11] `https://www.similarweb.com/top-websites/united-states/`

[12] A full list of the exceptions as of 2010, the end of our sample period, is available at: `https://web.archive.org/web/20100706030043/https://www.craigslist.org/about/help/posting_fees`

itized providing a useful service to local communities over profit maximization, and always opposed listing the company on the stock market.[13]

Being cheaper and more efficient than traditional newspaper ads, CL became very popular among users, rapidly disrupting the lucrative market for classified ads that many local newspapers had relied upon. Indeed, Figure A1, which plots the evolution of U.S. newspapers' revenues over time, shows that revenues from classified ads declined earlier on and faster than other sources, and that this decline started precisely around the time CL expanded.

Our analysis exploits variation in the timing of the introduction of CL across U.S. counties. In addition, we also exploit differences in how heavily newspapers relied on revenues from classified ads at baseline, which captures how much they were affected by the entry of a new competitor in that market.

3.2 Data

Our analysis combines data on: i) Craigslist'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 behavior outcomes including turnout, vote choices and campaign contributions, and vi) additional covariates.

3.2.1 Craigslist's expansion

Our analysis exploits the staggered introduction of CL across U.S. counties. 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 current websites. For a subset of these, the information is directly available on the CL's "about" webpage (`https://www.craigslist.org/about/expansion`). For the others, we assigned the date of the first snapshot recorded by the Internet Archive (`https://archive.org`). Figure 1 shows the evolution of the number of CL websites between 1995 and 2010, along with the evolution of a proxy for Internet quality - the average number of Internet Service providers by zipcode.[14] Figure A2 depicts the geographic distribution of CL websites in 2000, 2005, and 2010 respectively.

Mapping websites to counties is not straightforward since the area actually served by the website depends on user behavior. To address this issue, we identify the relevant market to a CL website in two ways. In our baseline approach, we assume that CL websites serve primarily the county (or set of counties) containing the place indicated in the headline of the respective local website. The place is usually a single city or town, but can also be a combination of two or three nearby cities, a region, or, in some cases, an entire state. We refer to this definition as the "core" CL market. As an alternative, we identify all the counties that account for a non-negligible share of the ads posted on the website. We describe this approach, which we use for purposes of robustness, in Appendix 3.5.

3.2.2 Newspaper characteristics and outcomes

We collect comprehensive data on a range of relevant newspaper characteristics and outcomes from a series of Newspaper Yearbooks published by Editor & Publisher for each year between

[13]For a profile of Craig Newmark and his business strategy see `https://www.theguardian.com/technology/2006/feb/19/news.theobserver1`.

[14]For more detail on this proxy, see section 3.2.5.

Figure 1: Craigslist's roll-out over time

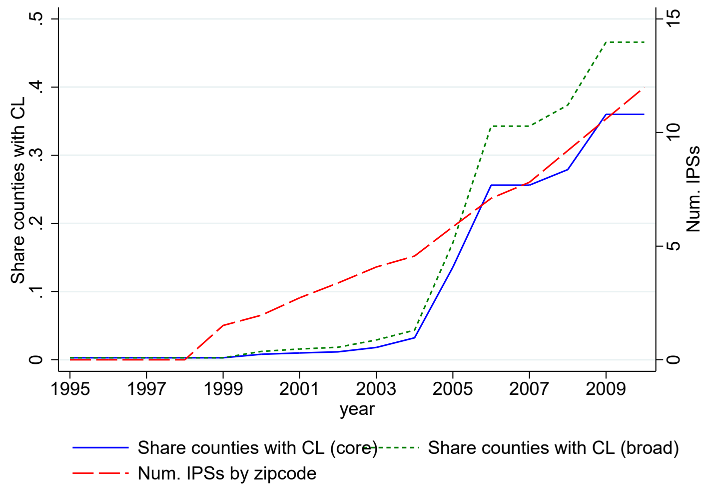

Share of US counties with a local CL website (left axis) and average number of Internet service providers by zipcode (right axis).

1995 and 2010. We accessed print copies the yearbooks and digitized the information contained in them using OCR software. The yearbooks contain detailed information for over 1,500 U.S. daily newspapers, including: address of the headquarters (HQ), circulation, subscription prices, number of pages published, as well as the list of staff members with names, broad job categories, and job titles. Figure A5 shows how the information is reported in the yearbooks for two example newspapers.

To understand whether and when a newspaper is affected by the entry of a CL website, we first need to match newspapers to counties. One way to do so is to assume that the newspaper market coincides with the county where its HQ is located. Since this a good approximation for the median newspaper,[15] this approach is common in this literature (Gentzkow and Shapiro 2010b; **?**). We obtain the HQ county for each newspaper from the address reported in the E&P data. In Appendix 3.5 we describe an alternative approach based on the use of zipcode-level circulation data available for a subset of newspapers from the Alliance of Audited Media (AAM), which we use for purposes of robustness.

Whether and how much a newspaper is affected by the entry of CL depends on how heavily it initially relied on revenues from classified ads. To measure this baseline difference across papers, we use information on the presence in a newspaper's staff of one or more classified ads managers prior to CL's entry (i.e. between 1995 and 2000), available from E&P. To validate this measure, we collect information from `https://www.newspapers.com`, an online newspaper archive, on the share of pages devoted to classified ads in a subset of 262 newspapers. The results, shown in Appendix 3.5, indicate that, prior to the entry of CL into their market, newspapers with classified ad managers devoted a significantly larger page share to classified ads than other papers, about 20% more than non-classified manager papers with similar circulation.

[15]Our disaggregated subscription data confirms that the median paper in our sample has about 85% of its subscribers in the HQ county; see Appendix 3.5.

3.2.3 Newspaper content

We are also interested in examining how changes in newspapers' organization triggered by the entry of CL influenced editorial decisions and news content. To this end, we use information from NewsBank, a database containing the text and metadata of articles published in over 800 newspapers in our sample, beginning in 1999. We use the data in two ways. First, we extract a random sample of 2 million articles and apply a semi-supervised topic model to the text of the lead paragraph to assess changes in the the coverage devoted to various topics. Second, 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 complementary measure of the prominence of political issues, both national and local. Details of the procedures used to construct these variables are reported in Appendix 3.5.

3.2.4 Political outcomes

Finally, we examine how the entry of CL, and the subsequent changes in newspapers organization and content, influenced citizens' political behavior. To this end, we collect data on a variety of electoral outcomes measured at the county, the Congressional district, or the county-by-Congressional district level.

First we look at electoral turnout, a standard measure of political participation. We use county-level turnout data from David Leip's Atlas of American Elections, covering both midterm and presidential elections between 1996 and 2010.

Second, we examine a set of outcomes capturing the electoral performance of ideologically extreme candidates. We expect the quality of political information available to voters should have strong influence on their ability to distinguish candidates on the ideological dimension. Following Hall (2015) and Dorn et al. (2020), we classify candidates based on their position in the distribution of campaign-finance-based ideological scores (CFScores) developed by Bonica (2014). Based on the 25th and 75th percentiles of the distribution of scores for all House candidates in 2000, we define thresholds that we use to separate "extremists" from "moderates." Our outcomes of interest are the presence of an "extremist" in the general election (related to the outcome studied by Hall (2015)), the shares of general election votes and individual contributions going to "extremist" candidates (as in Dorn et al. (2020)), and the absolute value of the CFscore of the winning candidate. Data on contributions and CFscores are from Bonica's Database on Ideology, Money in Politics, and Elections (DIME, 2016).

Third, we examine split-ticket voting, a measure of partisanship in voting that captures individual candidates' ability to differentiate themselves from the national party brand. Following Darr et al. (2018), we measure split-ticket voting as the absolute value of the difference between the Republican candidate vote share in the presidential election and the Republican candidate vote share in House elections in the same county and year.[16]

[16] Darr et al. (2018) use the Senate-President difference, but since only a third of Senate seats are contested in each election cycle, using House races expands the number of observations available. We measure vote shares for each office at the county level; House shares are computed by aggregating across all House votes cast by voters in the county (which for larger counties may be split across multiple representatives).

3.2.5 Additional data sources

Browsing data To measure CL take-up over time, as well as the downstream impact of CL on online news consumption, we draw on data from *Comscore*[17]. These data track the browsing behavior of a large sample of US Internet users, shown to be representative of US online buyers (De los Santos et al. 2012). The data cover the following years in our sample period: 2002, 2004, 2006 and 2007-2010. We aggregate the number of visits of the domain `craigslist.org` as well as total visits recorded by Comscore by county and year. To capture online news consumption, we rely on Comscore's classification of website categories available in the 2002 wave. We aggregate the number of visits by county for the 100 domains classified as news-related, and repeat this procedure for the following waves for the same set of domains.

Survey data To explore the impact of changes in newspapers' content on readers' news consumption habits, we use individual data from two large scale surveys. Our first source is the National Annenberg Electoral Survey (NAES), a nationally-representative rolling cross-sectional survey that was conducted in the lead-up to the 2000, 2004, and 2008 presidential elections. In particular, we use information on respondents' Internet access and self-reported media consumption in the week prior to the interview.

Our second source is the Survey of the American Consumer conducted by GfK Mediamark Research & Intelligence (GfK-MRI). The survey includes an extensive battery of questions about media consumption. In particular, we use information on respondents' self-reported readership of newspapers (in print and online, national and local), and news consumption on radio, TV, and on the Internet. We also exploit a question on what sections of the newspaper respondents usually read.[18] We use these data for the period 1999-2010.

Number of Internet Service Providers To separate the impact of CL entry from a generic Internet effect, we control for a measure of the quality of local broadband Internet. In the absence of disaggregated data on Internet subscribers for this period, we follow Larcinese and Miner (2018) and **?** and use the number of Internet service providers (ISPs) registered by zipcode as a proxy.

These data are available for the period 1998-2008 from the Federal Communication Commission (FCC) 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 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.

This measure has been shown to be a strong predictor of the number of broadband subscribers at the state level, as well as at the county level in later periods when such disaggregated data are available (Larcinese and Miner 2018). To further validate the number of ISPs as a proxy of local Internet penetration, we examine its correlation with self-reported Internet access from both the NAES and GfK-MRI surveys. Figure A9 confirms a strong positive relationship.

Other county characteristics Throughout our analysis we also use data on the following county-level variables: population (from the National Center for Health Statistics), income per capita, share of the population in urban areas, share of the population with college education,

[17] `https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/comscore/`

[18] Figure A6 shows the fractions of GfK respondents reporting that they read each section.

share of the population who rent housing, racial composition and median age (all from the 2000 Census), as well as unemployment rate (from the Bureau of Labor Statistics).

3.3 Empirical Strategy

3.3.1 Determinants of CL entry

To implement our empirical strategy, it is necessary to first understand what factors drove the timing of CL's staggered rollout. Anecdotal evidence suggests that CL prioritized larger markets and areas with good access to broadband Internet, which was crucial for the user to take full advantage of the platform. Importantly for our purposes, CL is not in the news business, and thus is not likely to have considered demand-side factors in the news market in determining where to enter. Also importantly, CL is privately held and has always operated as a mixture of profit-making business and community service. The fact that CL did not maximize profits[19] allowed more flexibility for the idiosyncracies of its influential founder and its early user base to determine the timing of the rollout, rather than systematic factors in the underlying local news markets.

To validate the assumption of conditional exogeneity of CL's timing of entry, in Figure 3.3.1 we examine the set of counties with headquartered newspapers in 2000 and plot the correlations of the year of CL's entry into a given county with its cross-sectional characteristics. The first two coefficients confirm anecdotal accounts that population and the quality of the local Internet connection, as measured by number of ISPs, were important considerations in CL's entry decisions. The magnitudes of both coefficients 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, our analysis will control for log population and number of ISPs throughout.

The other coefficients are from separate, univariate regressions on the time of CL's entry residualized by log population and number of ISPs. Conditional on these two variables, several demographic and political variables, including share urban population, racial composition, unemployment rate, age, electoral turnout and split-vote, appear unrelated to the timing of CL's entry. There are three exceptions: consistent with the profile of CL's early user base, counties with higher share of college educated population, higher rental share of housing and higher income per capita experience significantly earlier entry (3 to 5 months per standard deviation). We therefore also consider specifications controlling for each of the variables listed above.

Finally, and crucially for the purpose of our analysis, we find no relationship between timing of entry and the state of local newspapers as measured by circulation, number of jobs or the presence of classified managers. In other words, we find no evidence of CL targeting particular newspaper characteristics in its entry decisions.

[19] The judge in a 2010 civil action against Newmark by eBay 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. 3705-CC, decision dated 2010-09-09, `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.

Figure 2: Correlates of CL entry

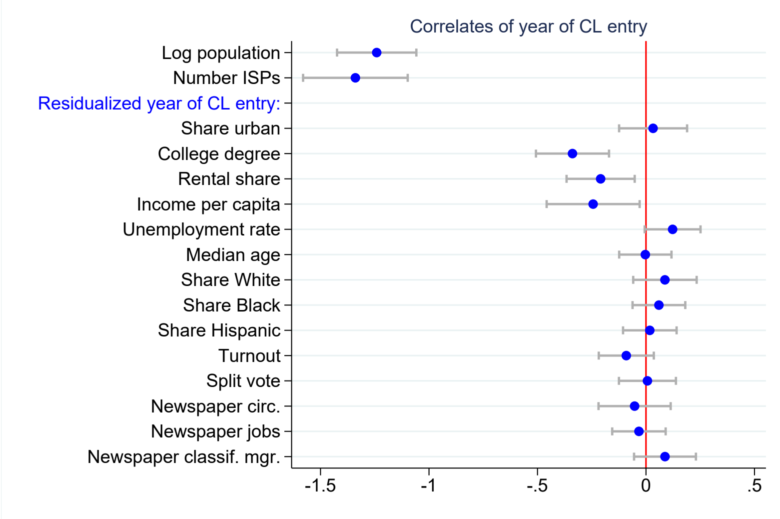

Coefficients and 95% confidence intervals representing the univariate correlations of (residualized) year of CL entry with various county characteristics. All variables are measured in the year 2000 and standardized to facilitate the comparison of magnitudes. The sample consists of all counties with newspaper HQs in the year 2000.

3.3.2 Newspaper-level regressions

To estimate the effect of CL entry on newspaper-level outcomes, we employ a difference-in-differences strategy exploiting CL's staggered introduction across US counties, combined with differences across newspapers in ex-ante reliance on classified advertising. The sample consists of the all newspapers covered by E&P, excluding ones with national circulation - i.e. the New York Times, USA Today and the Wall Street Journal.

The following equations summarize our approach:

$$Outcome_{nct} = \beta PostCL_{ct} + \phi_n + \psi_t + \rho' X_{ct} + \upsilon_t' Z_{c0} + \epsilon_{nct}, \tag{3.1}$$

$$Outcome_{nct} = \beta PostCL_{ct} + \gamma PostCL_{ct} \times ClassifMgr_{n0} + \phi_n + \psi_t + \rho' X_{ct} + \upsilon_t' Z_{c0} + \epsilon_{nct} \tag{3.2}$$

$Outcome_{nct}$ is one of the outcomes of interest for newspaper n, headquartered in county c, at time t. $PostCL_{ct}$ is a an indicator variable equal to one for years after the entry of CL in county c and zero otherwise. In the baseline analysis we follow the "core" definition of CL markets (based on location indicated in the website name) and newspaper markets (based on county of newspaper HQ).[20] ϕ_n and ψ_t are newspaper and year fixed effects, respectively. The vector X_{ct} includes controls for log population and number of ISPs, and the vector Z_c0 includes additional county-level controls from the 2000 census which we interact with year fixed effects. In alternative specifications we also control for state $\times$ year fixed effects or for DMA $\times$ year fixed effects, thus restricting the comparison to newspapers operating in the same state or in the same media market. Finally, $ClassifMgr_{n0}$ is an indicator variable equal to one if the newspaper had a classified manager at baseline. 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

[20]We obtain similar results - reported in Appendix **??** - using the respective "broad" definitions. In that case $post_C L$ is defined as 1 if there is an overlap between the market of CL website (defined based on the location of posted ads) with the newspapers' market (defined based on disaggregated circulation).

newspapers never affected by CL, by county.

Equation 3.1 estimates the impact of CL's entry under the assumption that the timing of entry is conditionally uncorrelated with pre-existing trends in these outcomes. To evaluate the plausibility of this assumption, for each outcome we will also estimate an event-study version of specification 3.1. In light of recent work showing that two-way fixed effect estimates can be biased in settings where treatments effects are heterogeneous over time or across groups, we present event-studies based on the time-corrected Wald estimator proposed by de Chaisemartin and D'Haultfoeuille (2020).

3.3.3 County-level regressions

To estimate the impact of CL entry on outcomes measured at the county level, we estimate versions of equations 3.1 and 3.2 aggregated by county, i.e.:

$$Outcome_{ct} = \beta PostCL_{ct} + \phi_c + \psi_t + \rho' X_{ct} + \upsilon_t' Z_{c0} + \epsilon_{ct}, \tag{3.3}$$

$$Outcome_{ct} = \beta PostCL_{ct} + \gamma PostCL_{ct} \times ClassifMgr_{c0} + \phi_c + \psi_t + \rho' X_{ct} + \upsilon_t' Z_{c0} + \epsilon_{ct} \tag{3.4}$$

When estimating this specification we focus on the sample of all counties where at least one newspaper was located at baseline, and compute $ClassifMgr_{c0}$ as the circulation-weighted average across newspapers headquartered in county c.

3.3.4 Congressional district-level regressions

For outcomes measured at the level of congressional districts, we estimate equations 3.3 and 3.4 at the level of county $\times$ district cells. In other words, we duplicate outcome observations and assign one duplicate to each county contained in the district. We weight observations by the share of the voting-age population of the district accounted for by the respective county, and cluster standard errors by district. To absorb variation due to changing congressional district boundaries, we include district by redistricting regime fixed effects in all regressions.[21] These fixed effects thus ensure that comparisons in the regressions are within fixed district boundaries.

3.4 Results

3.4.1 Craigslist Take-up

We first confirm that Internet users in a given area were more likely to visit CL's URL following the opening of a local website. To this end, we use the data on web browsing behavior from Comscore, described in section 3.2. In table 1 we estimate versions of our baseline specification (equation 3.1) using as dependent variable the IHS-transformed number of visits of the domain `craigslist.org` in a county in a given year. In all regressions we control for the total number of visits to any website recorded by Comscore. The results indicate that, after the entry of a local CL website, the number of visits to CL's URL increase significantly, by between 16 and 40%. The result is robust to expanding the set of county controls (col. 2), to controlling

[21] 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.

for state × year fixed effects (col. 3) or DMA × year fixed effects (col. 4), restricting the comparison to different counties within the same media market.

To get a sense of the timing of CL's take-up, in Figure 3 we plot the dynamic effects estimated following the method proposed by de Chaisemartin and D'Haultfoeuille (2020). The graph indicates no clear pre-trend, which alleviates concerns about the (conditional) exogeneity of CL's entry. It also suggests that the effect was immediate but further intensified over the following years. This pattern is consistent with network effects in the local adoption of the website: a larger number of users (and hence, a higher volume of local ads) likely increases the value of the platform and attracts yet more users.

Table 1: `craigslist.org` visits

	Dependent variable: CL visits (IHS)			
	(1)	(2)	(3)	(4)
Post-CL	0.440***	0.155*	0.200***	0.163***
	(0.129)	(0.084)	(0.073)	(0.060)
Total Comscore visits (IHS)	Yes	Yes	Yes	Yes
Log population, # ISPs	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	No	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes
Observations	19896	19896	19896	19501
Number of counties	3053	3053	3053	2995
R^2	0.80	0.83	0.84	0.85
Mean dependent variable	2.42	2.42	2.42	2.40

Regressions of number of visits of the domain `craigslist.org` by county and year (IHS-transformed) on an indicator for the availability of a local Craigslist website. All specifications control for total visits recorded in Comscore (IHS-transformed). County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.4.2 Number of newspapers

Next we examine how the entry of CL in a given market affected the number of active newspapers. One possibility is that the drop in revenues due to competition from CL might have been so extreme as to force some local papers to close or merge. In Table 2 we estimate our baseline specifications using as dependent variable the number of newspapers that, according to the E&P data, were operating in a given county in a given year. In this case our unit of analysis is a county-year, and the sample includes all counties where at least one newspaper was headquartered in 2000 (i.e., before CL's major roll-out). We find no evidence that CL entry affected the number of newspapers in a county.CL does not seem to affect the number of newspapers even in those counties where newspapers relied more heavily on classified ads *ex ante*; indeed, in columns 5 through 8, the coefficient on the interaction between *Post-CL* and *Classif.*

Figure 3: `craigslist.org` visits – Event Study

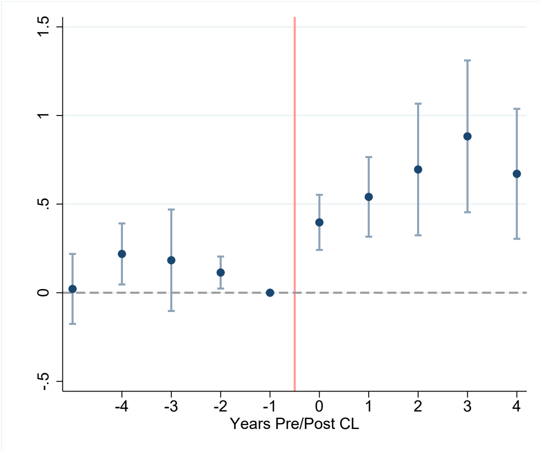

Dynamic effect of the availability of a local CL website on the number of `craigslist.org` visits (IHS-transformed) by county and year. Coefficients and 95% confidence intervals based on the time-corrected Wald estimator proposed in de Chaisemartin and D'Haultfoeuille (2020). Controls include total Comscore visits (IHS transformed), log population and number of Internet service providers. Standard errors clustered by CL-area.

Manager is small and never statistically significant. The same picture emerges when looking at the event-study graph depicted in Figure 4: all the coefficients for the years after CL entry are close to zero and statistically insignificant.[22] Taken together, these findings do not support the view that CL contributed to the disappearance of local newspapers. Yet, it is possible that lower advertising revenues affected the organization, functioning, and content of newspapers in other ways, a hypothesis which we explore below.

Figure 4: Number of newspapers – Event Study

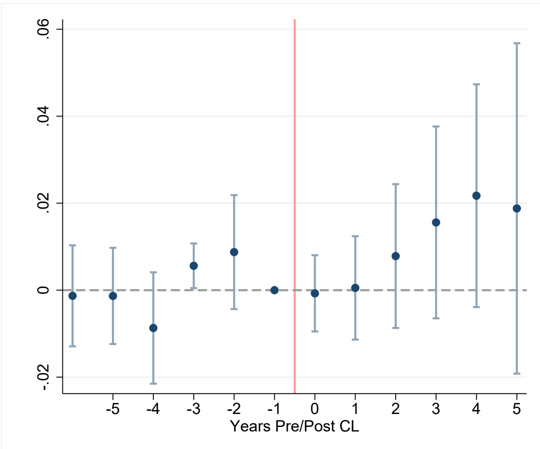

Dynamic effect of the availability of a local CL website on the number of newspapers headquartered by county and year. Coefficients and 95% confidence intervals based on the time-corrected Wald estimator proposed in de Chaisemartin and D'Haultfoeuille (2020). Controls include log population and number of Internet service providers. Standard errors clustered by CL-area.

[22]Given the lack of evidence that CL affected newspapers' exit and entry, in the remainder of our analysis we use the full unbalanced panel as baseline. We obtain similar results for the balanced panel of newspapers that remained in the sample for the entire period of analysis.

Table 2: Number of newspapers

	Dependent variable: Number of newspapers HQ-ed in county							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	0.012	0.017	0.009	0.010	0.027**	0.020	0.014	0.008
	(0.011)	(0.011)	(0.012)	(0.013)	(0.013)	(0.013)	(0.014)	(0.015)
Post-CL $\times$ Classified Mgr.					-0.028	-0.007	-0.010	0.002
					(0.018)	(0.018)	(0.018)	(0.020)
Log population, # ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. $\times$ Year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Year FEs	No	No	Yes	No	No	No	Yes	No
DMA $\times$ Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	19424	19424	19424	18672	19360	19360	19360	18608
Number of counties	1214	1214	1214	1167	1210	1210	1210	1163
R^2	0.94	0.95	0.95	0.95	0.94	0.95	0.95	0.95
Mean dependent variable	1.18	1.18	1.18	1.18	1.18	1.18	1.18	1.18

Regressions of the number of newspapers headquartered by county and year on an indicator for the availability of a local Craigslist website and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.4.3 Newspaper staff size and composition

In this section we examine to what extent the entry of CL local websites influenced the organization of local newspapers. The first outcome we focus on is staff size. In Table 3 we estimate our baseline specifications using as dependent variable the number of jobs reported in the E&P yearbooks by newspaper and year. In the first four columns we examine the effect of being in the *Post-CL* period. Results in columns 1-4 indicate that following the entry of CL, the number of jobs in local papers decreased significantly. The effect amounts to 1 to 1.2 fewer jobs, or about 5% of the mean, and is robust to the inclusion of additional controls and finer fixed effects. The results in the last four columns confirm that the effect is mainly driven by newspapers that relied more on classified ads prior to CL, for which the effect, looking at the most demanding specification (column 8), is over 12% relative to the mean. Figure 5 displays the associated event-study, which shows the absence of any significant pre-trend, and that the effect, which is significant in the first year, increases considerably over several years.

This result is robust to several additional checks. First, we verify that the indicator for classified manager does not simply capture newspaper size, as newspapers with larger circulation are more likely to have a dedicated classified manager (correlation = 0.26 in our sample). In Table B1 we show that the results are robust to controlling for baseline circulation interacted with $post\_CL$, and that this interaction has no significant effect on jobs-count. Second, we consider as dependent variable the number of (unique) employees instead of the number of job titles. These two variables can differ if, for example, financial difficulties push a newspaper to assign to the same person multiple jobs that were previously carried out by different people. Table B2 shows that we obtain very similar results.

One important question that our data allow us to examine is what categories of workers were most affected by staff cuts. In fact, based on the job title we can determine whether a worker holds a managerial or an editorial position, and, for editorial staff, we can identify the corresponding topical area (e.g., politics, sports, entertainment, etc.).[23]

In the first two columns of Table 4, we estimate our main specification with the *Post-CL*× *Classif. Manager* interaction separately for managerial staff and editorial staff. The results indicate that, for the newspapers most affected by the entry of CL, staff cuts concerned both types of positions, with the effect on managerial positions being larger (i.e., 19% vs. 7% of the mean). In the following columns we examine what topic areas were most affected by cuts in editorial staff, looking in particular at politics, sports, and entertainment. The evidence indicates that newspapers most affected by CL were significantly less likely to have dedicated political editors and reporters after CL's entry, while the same was not the case for sports or entertainment. One interpretation of this result is that, when facing financial difficulties, newspapers affected by CL opted to cut staff especially in areas - like local politics - that readers do not value as much and for which producing quality content is more costly.

On the other hand, despite this significant downsizing, we find no effect on the number of pages published by affected newspapers (Appendix Table B3). Together, these results imply increased workload per staff member and may have implications for the distribution of editorial priorities. We examine this issue further in the next section where we look at how the entry of CL affected the evolution of news content.

Table 3: Number of Jobs

	Dependent variable: Newspaper number of jobs							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	-2.030***	-1.084***	-1.242***	-1.262***	-0.103	0.453	0.274	0.163
	(0.398)	(0.390)	(0.418)	(0.404)	(0.412)	(0.411)	(0.432)	(0.442)
Post-CL × Classified Mgr.					-3.623***	-2.975***	-2.911***	-2.802***
					(0.575)	(0.548)	(0.549)	(0.594)
Log population, #ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	22724	22724	22723	21968	22516	22516	22515	21762
Num. of newspapers	1540	1540	1540	1492	1505	1505	1505	1459
R^2	0.90	0.91	0.91	0.92	0.90	0.91	0.91	0.92
Mean dep. var.	21.31	21.31	21.31	21.14	21.38	21.38	21.39	21.23

Regressions of number of jobs by newspaper and year on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

[23]Specifically, we classify as managerial positions the jobs listed in the E&P section "Corporate/ General Management", and as editorial positions the jobs listed in the sections "News Executives" and "Editorial Management" (see figure A5). Regarding topic areas, we classify jobs containing the keyword "sports" as sports related, ones containing the keywords "entertainment"/ "movie" / "music" / "theater" / "travel" / "lifestyle" as entertainment-related, and ones containing the keywords "politics", "government", "Washington" as politics-related.

Figure 5: Number of Jobs – Event Study

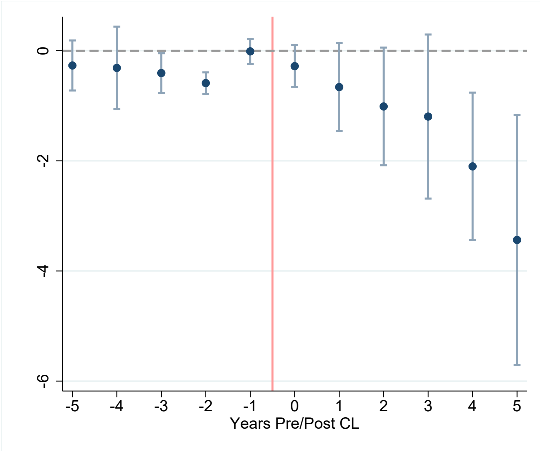

Dynamic effect of the availability of a local CL website on number of jobs by newspaper and year. Coefficients and 95% confidence intervals based on the time-corrected Wald estimator proposed in de Chaisemartin and D'Haultfoeuille (2020). Controls include log population and number of Internet service providers. Standard errors clustered by CL-area.

Table 4: Jobs by type and topic

	(1) Num. Managers	(2) Num. Editors	(3) Editor Politics	(4) Editor Sports	(5) Editor Entertainment
Post-CL	0.050	0.138	0.022	-0.019	0.001
	(0.086)	(0.306)	(0.019)	(0.018)	(0.001)
Post-CL × Classified Mgr.	-0.667***	-0.782*	-0.050**	-0.003	-0.001
	(0.122)	(0.444)	(0.024)	(0.021)	(0.002)
Log population, share urban, # ISPs	Yes	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes
Observations	21834	22155	21454	21558	21353
Number of newspapers	1503	1504	1504	1504	1504
R^2	0.80	0.89	0.69	0.51	0.08
Mean dependent variable	3.46	10.60	0.40	0.84	0.00

Regressions of number of jobs by type and indicators for the presence of an editor dedicated to a specific topic on an indicator the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.4.4 Newspaper content

Next we explore how the transformations in newspapers' organization documented above translate into changes in news content, in general, and news coverage of politics, in particular. This is a question of paramount importance considering that, for most citizens, local newspapers still represent the main source of political information which allows them to monitor and keep politicians accountable (Mahone et al. 2019; Snyder and Strömberg 2010). To analyse changes in news content we apply text analysis techniques to large corpus of news articles published in a subset of the newspapers in our original sample. Our analysis is composed of three parts.

First, we study how competition from CL affected the volume of coverage devoted to different topics. To do so, we estimate a topic model on a corpus consisting of the first paragraphs of over 2 million randomly drawn articles published in over 800 newspapers between 2000 and 2010, available from NewsBank. We use the Correlation Explanation (CorEx) model of Gallagher et al. (2017a). This model has the advantage that it tends to produce coherent topics for corpora consisting of short texts, and also allows us to define anchor words to generate topics corresponding to specific areas. Specifically, we seed the model to produce topics associated with: i) President, ii) Congress, iii) local politics, and iv) foreign policy.[24] Figure A10 reports the most frequent words associated with each of the ten topics generated by the algorithm, with the anchor words underlined. For each article in the corpus, we obtain a probability associated with each of the 10 topics, and aggregate the probability weights by newspaper and year. Figure 6 reports the coefficients on the interaction $Post - CL \times Classified.Mgr.$ obtained from regressions of the form specified in equation **??** with dependent variables corresponding to the probability of each topic.[25] The results indicate a general decline in news coverage of politics. For newspapers most affected by CL, the drop is significant for presidential, congressional, and foreign politics, and marginally insignificant for local politicsFor the remaining un-anchored topics (i.e., accidents, events, traffic, obituaries, sports, crime) we find mixed results, with insignificant coefficients on the interaction *Post-CL* $\times$ *Classif. Manager*.

Figure 6: Content: Distribution of Topics

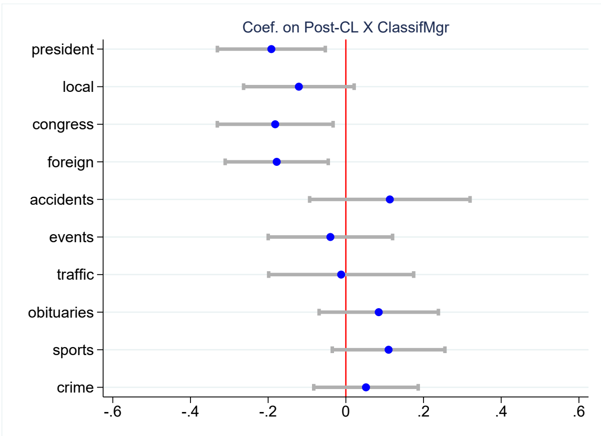

Effect of CL-entry on the (standardized)) probability that a randomly sampled article covers a given topic. The graph presents the coefficients and 95% confidence intervals on the interaction of an indicator for the availability of a local Craigslist website with an indicator for the presence of a classified manager at baseline. Controls include log population, number of Internet service providers, and county characteristics in 2000 interacted with time. Standard errors clustered by CL-area.

To further explore changes in political coverage, we examine how frequently local newspapers report about local representatives and national politicians. Using keyword searches in the NewsBank archive, we first identify all articles published in a given newspaper and year

[24]Additional details about the procedure are reported in Appendix 3.5.

[25]Tables B5 and B6 report the corresponding full regression results.

that contain the names of the House representative elected in the district corresponding to the newspaper HQ or senate representatives elected in the respective state.[26]

Table 5 suggests that, following the entry of CL, coverage of local Congressional representatives declines significantly, particularly for newspapers that relied more heavily on classified ads. In that case the differential decline in representatives' coverage amounts to about 30% of the mean. In Table B7 we perform the same exercise for national party leaders: the President and the leadership of both parties in each chamber of Congress. We find no significant effect on coverage of these national politicians. A possible explanation for this pattern is that coverage of national politicians can be sourced from wire services, whereas coverage of the district's representatives is original content produced by in-house reporting staff. It is exactly these staff, per the results in Table 4, whose jobs were most likely to be cut following CL's entry.

Table 5: Mentions of local Congressional representatives

	Dependent variable: Articles mentioning local House/Senate representives (IHS)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	0.001	-0.054	-0.028	0.041	0.142	0.097	0.078	0.208
	(0.074)	(0.079)	(0.079)	(0.099)	(0.100)	(0.103)	(0.104)	(0.127)
Post-CL × Classified Mgr.					-0.258**	-0.292**	-0.218*	-0.317**
					(0.121)	(0.122)	(0.125)	(0.157)
Total articles (IHS)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Log population, # ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	5453	5453	5426	5036	5410	5410	5383	4985
Number of newspapers	821	821	818	771	814	814	811	763
R^2	0.82	0.82	0.85	0.88	0.82	0.82	0.85	0.88
Mean dependent variable	2.88	2.88	2.87	2.86	2.89	2.89	2.88	2.87

Regressions of the (IHS-tranformed) number of articles mentioning the name of a local House or Senate representative by newspaper and year on an indicator for the availability of a local Craigslist website, and its interaction with an indicator for the presence of a classified manager at baseline. All specifications control the (IHS-transformed) total number of articles recorded by Newsbank by newspaper and year. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.4.5 Newspaper Readership

How did readers react to the changes newspaper organization and content documented above? One way to tackle this question is by looking at the evolution of the total number of copies sold, i.e., newspaper circulation, available from E&P for all newspapers in our sample. In Table 6 we estimate our standard set of regressions using as dependent variable yearly circulation. The

[26]Details of the keyword searches are reported in Appendix 3.5. We restrict the sample to the period after decennial redistricting (i.e. post-2003), to maintain stable district boundaries.

coefficient on Post-CL is negative although it loses significance when including state or DMA fixed effects interacted with year dummies. The effect is instead negative, large, and highly significant for the newspapers most affected by CL. The size of the coefficient is fairly stable across specifications and has a magnitude of about 9% of the sample mean when all county controls are included. The corresponding event study graph, reported in Figure 7, shows the absence of pre-existing trends and a steady increase in the effect within a few years of CL's entry.

Table 6: Circulation

	Dependent variable: Newspaper circulation in thousands							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	-2.712***	-1.468**	-1.194	-0.892	0.012	0.108	0.397	0.700
	(0.599)	(0.645)	(0.765)	(0.939)	(0.752)	(0.874)	(0.953)	(1.231)
Post-CL $\times$ Classified Mgr.					-5.148***	-2.991***	-2.896***	-2.921**
					(1.297)	(1.096)	(1.055)	(1.351)
Log population, #ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. $\times$ year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Year FEs	No	No	Yes	No	No	No	Yes	No
DMA $\times$ Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	22950	22950	22948	22195	22633	22633	22631	21880
Number of newspapers	1555	1555	1555	1507	1506	1506	1506	1460
R^2	0.98	0.98	0.98	0.98	0.98	0.98	0.98	0.99
Mean dependent variable	33.34	33.34	33.34	32.89	33.40	33.40	33.39	32.93

Regressions of circulation by newspaper and year, measured in thousands of copies, on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

We can verify this decline in circulation using self-reported newspaper readership, available from the NAES and GfK-MRI surveys. These data include several questions related to media consumption. While we have 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*, *USA Today* or *the Wall Street Journal*, and the rest. We rely on residents' county of residence to match them to locally headquartered newspapers, and re-define the variable *Classif. Mgr.* as the circulation-weighted average across newspapers based in that county.

In Table 7 we report the results for self-reported readership of non-national newspapers and its frequency from individual-level regressions. In all specifications we control for a range of respondent characteristics including age, race and education. The results based on NAES data - reported in columns 1-4 - suggest a significant decline readership in areas where newspapers were most exposed to CL-competition. The effect is quite sizeable, corresponding to about 7% of the sample mean for both outcomes. Results are qualitatively similar, though somewhat smaller in size ($\approx 4\%$), when estimated using GfK data. Since in both surveys respondents are asked about reading either the print or the online version of the newspaper, these results indicate that the decline in circulation documented in Table 6 does not merely reflect substitution of print editions with online editions.

Figure 7: Circulation – Event Study

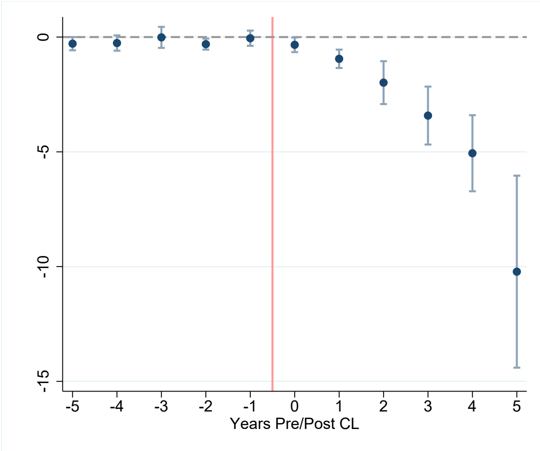

Dynamic effect of the availability of a local CL website on circulation (in thousands of copies) by newspaper and year. Coefficients and 95% confidence intervals based on the time-corrected Wald estimator proposed in de Chaisemartin and D'Haultfoeuille (2020). Controls include log population and number of Internet service providers. Standard errors clustered by CL-area.

Table 7: Self-reported newspaper readership

	NAES				GfK-MRI	
	Read newspaper Dummy		Read newspaper Days per wk		Read newspaper Dummy	
	(1)	(2)	(3)	(4)	(5)	(6)
Post-CL	-0.023	0.019	-0.108	0.118	-0.016***	-0.003
	(0.017)	(0.015)	(0.090)	(0.108)	(0.005)	(0.008)
Post-CL × Classified Mgr.		-0.053**		-0.284*		-0.018*
		(0.164)		(0.027)		(0.009)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes
County controls	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	108332	107503	108332	107503	253513	251442
Number of counties	1207	1203	1207	1203	792	790
R^2	0.06	0.06	0.14	0.14	0.09	0.09
Mean dependent variable	0.75	0.75	3.71	3.71	0.42	0.42

Regressions of self-reported newspaper readership on an indicator for the availability of a local Craigslist 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. County controls include contemporaneous log population and number of Internet service providers, as well as share urban population, pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share - all measured in 2000 and interacted with year FEs. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The decline in readership documented above 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. We test for this mechanism in Appendix Table B4, looking at the impact of CL's entry on yearly subscription prices reported in the E&P yearbooks. With the exception of the last column, we obtain insignificant coefficients on both the indicator for CL-entry and its interaction with classified manager. There is hence no clear evidence of an increase in subscription prices.[27]

A second explanation may be 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 and Figure 6 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 1998-2001 waves of the GfK-MRI survey. The distribution of readers' preferences, depicted in Figure A6, 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 Classifieds as one of the sections they frequently read. It is, therefore, possible that a reduction in the value of print classifieds might be a driver of circulation declines.

To understand what types of readers drive the drop in readership of local papers following the entry of CL, we examine heterogeneity in the readership effect by propensity to read the classified versus general news sections. Using the 1998-2001 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.[28] Based on the model estimated in the 1998-2001 data, we then project two propensity scores for respondents in the later years. This procedure allows us to assign to each respondent in the post-CL surveys 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), indicating that the groups that tend to read each section are relatively distinct.

We then 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 Table 8, indicate that the decline in readership after the entry of CL is entirely driven by indi-

[27] **?** on the other hand find a significant increase in subscription prices in response to CL entry. Our analysis differs in using more complete data (i.e. covering all newspapers in the period 1995-2010) and a somewhat different empirical strategy, e.g. looking at reliance on classified ads at baseline rather than contemporaneously, to account for its likely endogeneity.

[28] The characteristics most strongly associated with general news reading are white race, having a post-graduate degree, income in the 75K-150K range, being retired, being married 25 years or more, and being age 45-49. The characteristics most associated with classifieds reading 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.

viduals with high news propensity and low classified propensity, and is again more pronounced where newspapers relied more heavily on classifieds ex ante. The effect is not significant for the high-classified-interest, low-news-interest group in any of the specifications.[29] Though newspapers which offered the most classifieds were most affected by the CL shock, the readers least interested in classifieds reduced their newspaper reading most.

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

Table 8: Self-reported newspaper readership: Heterogeneity

	(1)	(2)	(3)	(4)
	Low news propensity & High classif. propensity		High news propensity & Low classif. propensity	
Post-CL	-0.005	-0.004	-0.016**	0.009
	(0.007)	(0.012)	(0.008)	(0.012)
Post-CL × Classified Mgr.		0.001		-0.033**
		(0.013)		(0.013)
Respondent controls	Yes	Yes	Yes	Yes
County controls	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	84341	83475	84338	83854
Number of counties	784	782	781	779
R^2	0.05	0.05	0.10	0.10
Mean dependent variable	0.32	0.32	0.54	0.54

Regressions of self-reported newspaper readership on an indicator for the availability of a local Craigslist 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. County controls include contemporaneous log population and number of Internet service providers, as well as share urban population, pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share - all measured in 2000 and interacted with year FEs. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

3.4.6 Substitution to Other News Sources

To understand the implications of the documented decline in readership, a crucial question is to what extent this decline is offset by consumption of other news sources. If for instance readers merely substitute newspapers for other sources that cover similar content, such mode-switching, though detrimental for the papers, would not necessarily imply a reduction in actual news consumption.

Both NAES and the GfK survey include questions about news consumption via sources other than local newspapers, which we use to study substitution patterns. The results are reported in Tables B9, B10 and B11. In columns 1 and 2 we examine how CL affected the

[29] Alternatively, in appendix table B8 we explore the two dimensions of readership propensity separately. We find more pronounced declines in readership for respondents with high rather than low news propensity and for respondents with low rather than high classified propensity.

readership of national newspapers. In both surveys we find that the entry of CL in an area is associated with an increase in the likelihood that local respondents read national papers. Yet, the effect does not appear to be stronger in areas where newspapers relied more heavily on classified ads. The remaining columns show no significant effect on news consumption on TV, radio, and online, including in areas where newspapers were most affected by CL.[30]

The data on browsing behavior available from Comscore allows us to perform an alternative test for substation to online news sources. In Appendix table B12 and Figure **??** we examine the effect of CL's entry on the number of visits of 100 domains classified by Comscore as news-related. We detect no significant effect of CL's entry on visits of these domains.

Taken together these findings suggest that the decline in readership of local newspapers associated with the entry of CL is not fully compensated by increased news consumption online or through other media. These effects are therefore likely to translate in a net decline in exposure to political information.

In the previous sections we documented that newspapers affected by the entry of CL reported less about politics, in general, and local politics, in particular. We also found that individuals in areas affected by CL experienced a decline in readership of local papers not compensated by increased consumption of other news sources. In this section, we examine how these changes affected the behavior of local voters. 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 outcomes examined in the existing literature on media, political participation, and electoral accountability. Specifically, we investigate how the entry of CL affects: i) voters' propensity to turn out in elections (Gentzkow et al. 2011), ii) to vote for ideologically extreme candidates (Hall 2015; Dorn et al. 2020), and iii) to rely on national partisan cues when voting for local candidates, measured by the incidence of split-ticket voting (Darr et al. 2018; Moskowitz 2021; Trussler 2020).

3.4.6.1 Turnout

As described in section 3.2, electoral turnout is measured at the county level, and turnout data are available for both presidential and midterm election years. We assume that newspapers affect voters' behavior 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. Since turnout is defined at the county-level, we estimate equation 3.4 as our baseline specification. Given the considerable differences in the number of residents across counties, we weight observations by the county's total voting-age population.

In Appendix Table B13 we examine the effect of CL on voter turnout. The results provide some evidence of a negative effect of CL on turnout in presidential elections, which is concentrated in counties where newspapers relied more on classified ads. However, the coefficient on *Post-CL* $\times$ *Classif. Mgr* becomes smaller when controlling for county-level covariates and, especially, when $State \times Year$ fixed effects are included. In addition, as shown in Appendix Table B14, we find no significant effect when looking at turnout in House elections, for which one might expect the decline in news coverage of local members of Congress to be especially relevant. We conclude that there is at most weak evidence that the changes in newspaper content and readership induced by the entry of CL reduced aggregate electoral participation.

[30]The analysis of online news consumption with NAES data is not possible since the question is not asked consistently across waves.

3.4.7 Support for extreme candidates and split-ticket voting

Though changes in the information environment may not impact *whether* people vote, they may affect *how* they vote. In Table 8 we examine whether the reduction in political news coverage and readership brought about by the entry of CL favored the emergence and success of ideologically extreme candidates, focusing on House elections. The hypothesis we test is that a coarser information environment, by making it harder for voters to acquire information about candidates' ideological positioning, makes the entry of more extreme candidates more likely and improves their electoral prospects. Specifically, we look at the following outcomes: i) the probability that ideologically extreme candidates win a primary election, ii) the vote share they obtain in the general election, iii) the individual campaign contributions they attract, and iv) the ideological extremity[31] of the winning candidate.

Outcome variables i), iii) and iv) are defined at the electoral district level, while ii) varies at the county-by-district level. As explained in section 3.3, all regressions include district $\times$ redistricting regime fixed effects to absorb the effect of changes in district boundaries. Hence, we exploit variation over time within a fixed district boundary. To simplify the presentation, we report the results from the specification with the interaction term *Post-CL* $\times$ *Classif. Mgr.* and the full set of time-varying controls.

In column (1) we look at the effect on the probability that a candidate with an extreme CFScore wins a primary election.[32] The results indicate that, after the entry of CL, and in districts where newspapers were most affected, this probability increases significantly by about 7 percentage points. In columns (2) and (3) we look at the effect on general election vote share and share of campaign contributions, respectively, won by extreme candidates. Again, we find a positive, significant, and sizeable effect of CL entry on popular support for extreme candidates along both dimensions. Finally, in column (4) we examine the effect on the extremity of the CFscore of the winning candidate relative to the median among House candidates in 2000. Again the effect is positive, significant, and concentrated in areas served by the most CL-affected papers.

Finally, in Table 9 we examine the impact of the entry of CL on split-ticket voting, i.e., voters' tendency to support candidates from different parties in concurrent elections. We focus on presidential election years since split-ticket voting is defined by comparing candidates' vote shares in presidential vs. legislative elections. The results indicate that, following the entry of CL, voters become significantly less likely to split their vote between candidates of different parties. As for the other outcomes, the effect is driven by areas where newspapers were most vulnerable to CL's competition, where split-ticket voting drops by 16% of the sample mean. One interpretation of this finding is that, as local media provide less information about local candidates, voters tend to rely more heavily on partisan cues — which are shaped primarily by the national political debate — when deciding on down-ballot races. This substitution of local cues for national ones can lead to higher party alignment between different races, i.e. the "nationalization" of local elections (Moskowitz 2021; Trussler 2020), and further incite ideological divisions. This result relates to similar findings by Darr et al. (2018) regarding the impact of newspaper closures on polarization. Our results indicate that closures are not a necessary condition, and that the impoverishment of local newspapers, and the associated cuts in news coverage of local politics, can produce similar consequences.

[31] Measured by the absolute value of the difference between the candidate's CFScore and the CFScore of the median House candidate in 2000.

[32] The variable takes value of 1 if either the Republican or the Democratic candidate who advances to the general election qualifies as an "extremist" according to our definition.

Figure 8: Ideological Polarization

	Extremist in General	Vote Share of Extremists	Indiv. Contrib. Share of Extremists	Winner CFScore Dev. from 2000 Median
	(1)	(2)	(3)	(4)
Post-CL	-0.005	-0.037*	-0.011	0.002
	(0.030)	(0.021)	(0.024)	(0.018)
Post-CL × Classif. Mgr.	0.067**	0.055**	0.049*	0.046**
	(0.030)	(0.024)	(0.025)	(0.023)
Log pop., pct. urban, #ISPs	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes
District FEs	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes
Observations	12,797	12,797	12,534	11,048
R^2	0.58	0.64	0.74	0.86
Mean dependent variable	0.74	0.38	0.51	0.72

Regressions of electoral outcomes by county × district cell and year on an indicator for the availability of a local Craigslist website, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. Even-numbered years from 1996 to 2010 are included. Vote shares are computed by county × district cell; other outcomes are defined at district level only. The column (1) outcome is an indicator for the presence in the general election of a candidate with CFScore outside the central 50% interval of House candidates in 2000, following the method of Dorn et al. (2020). Columns (2) and (3) are the share of general election votes and contributions from individuals of such candidates. Column (4) is the absolute value of the CFScore (Bonica 2014) of the candidate who won the election. "District" means unique combination of state, congressional district number, and redistricting regime (either 1991 or 2001 for all states, plus 1997 for VA and NC, 2003 for TX, and 2005 for GA). County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. Observations are weighted by the share of the county in the district's voting-age population. OLS regressions in all columns. Standard errors clustered by district. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Taken together, our findings suggest that the profound transformations in the media landscape triggered by the entry of CL had a tangible impact on electoral politics in the US, and contributed to increase ideological polarization and further divisions across party lines.

3.5 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 generally to under-provision of news production 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

Table 9: Split-Ticket Voting in House vs. Presidential Elections

	Dependent variable: House-President Rep. vote share differential					
	(1)	(2)	(3)	(4)	(5)	(6)
Post-CL	-0.011	-0.013*	-0.013*	0.004	0.004	0.006
	(0.008)	(0.007)	(0.007)	(0.010)	(0.009)	(0.010)
Post-CL × Classified Mgr.				-0.020*	-0.022**	-0.022*
				(0.011)	(0.011)	(0.011)
Log population, # ISPs	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	No	Yes	Yes	No	Yes	Yes
State × Year FEs	No	No	Yes	No	No	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4807	4807	4807	4791	4791	4791
Number of counties	1214	1214	1214	1210	1210	1210
R^2	0.49	0.51	0.62	0.49	0.51	0.62
Mean dependent variable	0.13	0.13	0.13	0.13	0.13	0.13

Regressions of split-ticket voting by county and election year on an indicator for the availability of a local Craigslist website, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. Split-ticket voting is defined as the absolute value of the difference between the Republican candidate vote share in the Presidential election and the Republican candidate vote share in the House election(s). County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Observations are weighted by voting-age population. OLS regressions in all columns. Standard errors clustered by CL-area. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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 such competitor, the classified advertising platform Craigslist, had severe impacts on newspapers' staffing levels and production of news coverage relating to local politics.

The Craigslist 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-Craigslist 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 particularly reduced political coverage and especially 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 Craigslist shock saw their House elections become more nationalized, which we interpret as a consequence of thinner information about the local incumbent's behavior. Changes in the media landscape may thus be an important driver of the overall trend towards nationalization of elections in the United States (Hopkins 2018).

The change in voters' information about candidates 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.

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 Craigslist 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.

Appendix A: Background and Data

Background

Figure A1: Evolution of newspaper revenues by source

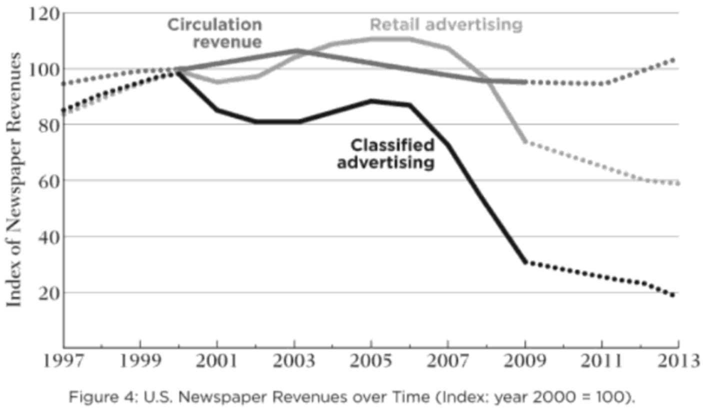

Index of newspaper revenues from circulation, retail advertising and classified advertising - 1997 to 2013. Source: Newspaper Association of America.

Figure A2: Geographic distribution of CL websites (2000-2010)

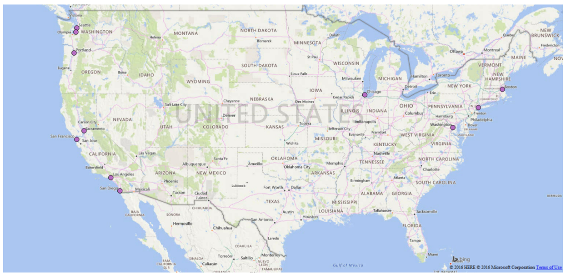

(a) 2000

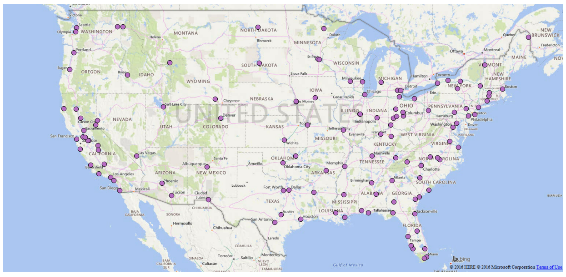

(b) 2005

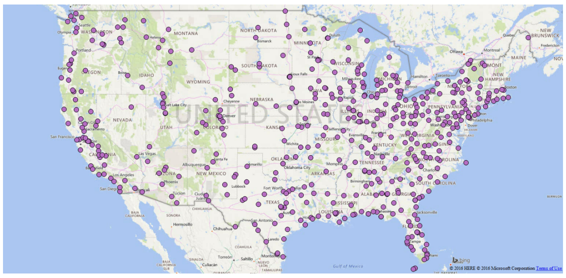

(c) 2010

Georgraphic distribution of CL websites over time.

Figure A3: Craigslist: Layout in 2000 and 2016

2000

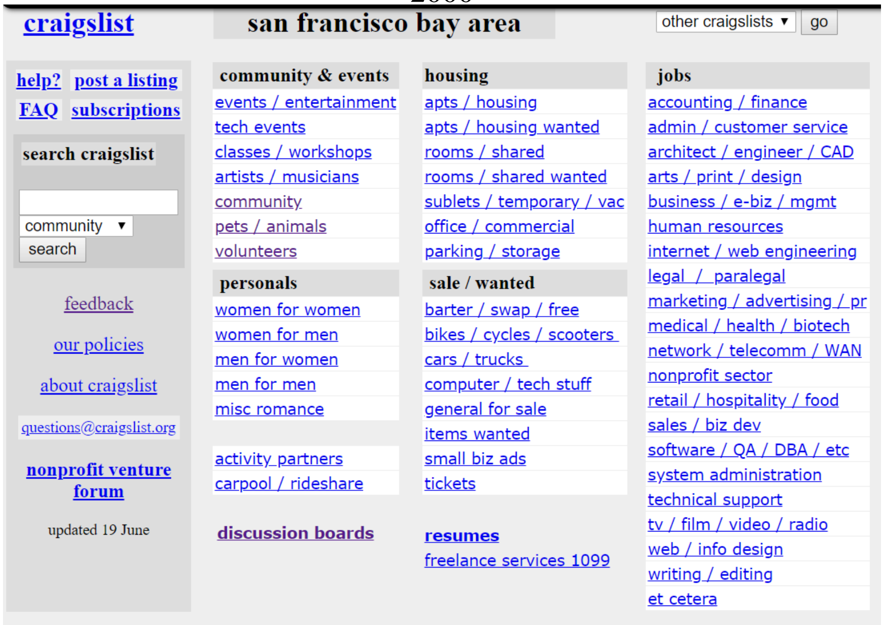

2016

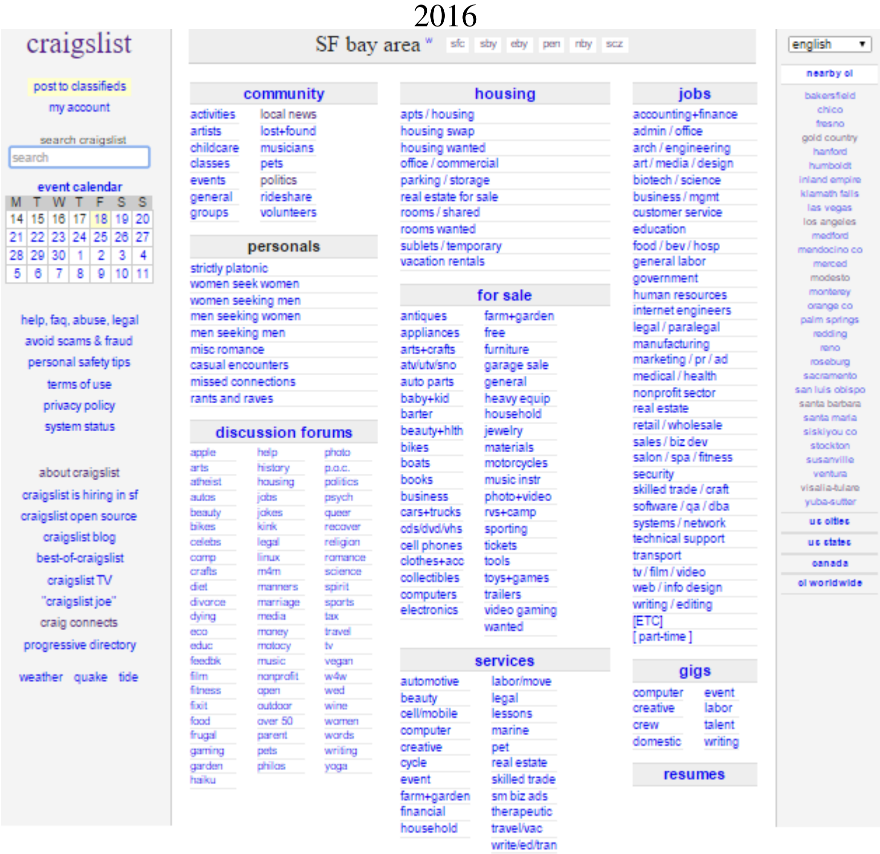

Figure A4: Craigslist: Layout of Housing Listings

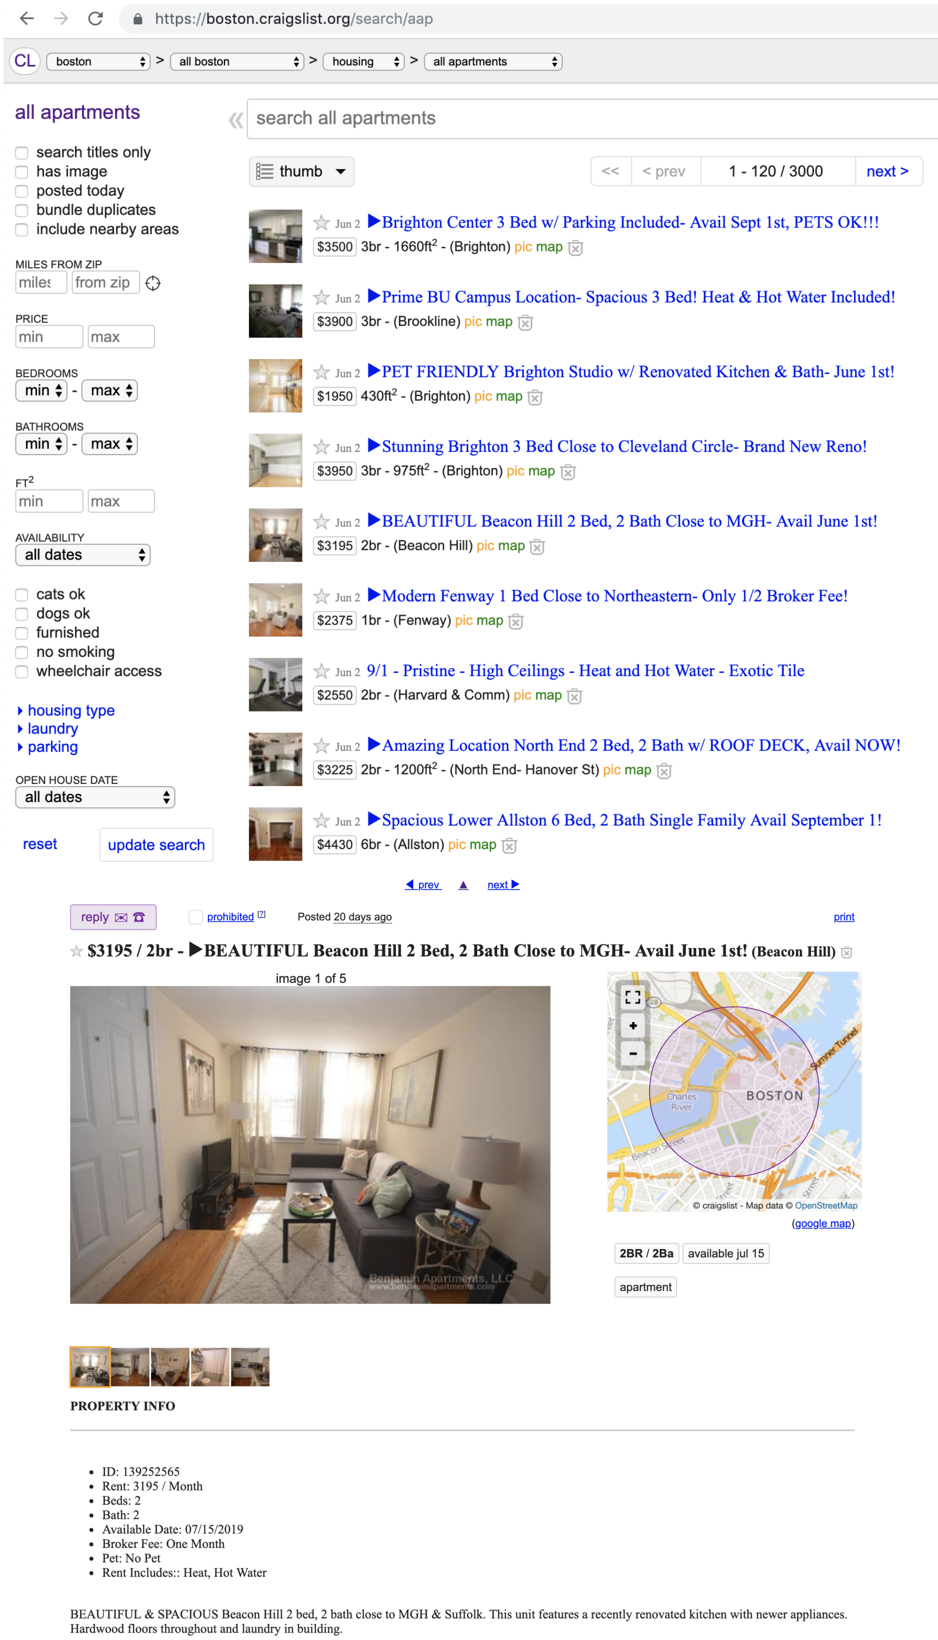

Figure A5: Extract from the Editor and Publisher Yearbooks

The Reporter
(m-mon to fri; m-sat)

The Reporter, 307 Derstine Ave.; PO Box 390, Lansdale, PA 19446; gen tel (215) 855-8440; adv tel (215) 361-8849; ed tel (215) 361-8814; gen fax (215) 855-6147; ed fax (215) 855-3432; adv email imaging@thereporteronline.com; ed email letters@thereporteronline.com; web site http://www.thereporteronline.com.
Group: Journal Register Co.
Circulation: 17,808(m); 15,590(m-sat); ABC Sept. 30, 2003.
Price: $0.50(d); $0.50(sat); $3.00/wk (carrier); $156.00/yr (carrier), $196.00/yr (mail).
Advertising: Open inch rate $33.83(m); $33.83(m-sat). **Representatives:** Landon Media Group; U.S. Suburban Press Inc.; Robert Hitchings & Co.
News Services: AP, GNS.
Politics: Independent. **Established:** 1870.

CORP. MGMT./GEN. MGMT.

Pres./Pub.	Al Frattura
Controller/Purchasing Agent	Bernard DeAngelis

ADVERTISING SALES MGMT.

Adv. Dir.	Robert Twesten
Display Adv. Mgr.	Angel Hernandez

NEWS EXECUTIVES

Exec. Ed.	Nona Breaux

EDITORIAL MGMT.

City Ed.	Monica Thompson
Lifestyles Ed.	Aixa Torregrosa
Night Ed.	Linda Doell
Page 1 Ed.	Dan Sharer
Chief Photographer	Geoff Patton
Special Sections	Kass Picozzi
Sports Ed.	Kevin Lilley

The Reporter, Landsdale PA

Chicago Tribune
(m-mon to tues; m-wed to fri; m-sat; S)

Chicago Tribune, 435 N. Michigan Ave., Chicago, IL 60611; gen tel (312) 222-3232; gen fax (312) 222-2595; gen email tribletter@tribune.com; web site http://www.chicagotribune.com.
Group: Tribune Co.
Circulation: 680,879(m); 512,455(m-mon to tues); 571,576(m-sat); 1,002,166(S); ABC Sept. 30, 2003.
Price: $0.50(d); $0.50(sat); $1.79(S); $4.40/wk; $228.80/yr.
Advertising: Open inch rate $580.00(m); $580.00(m-sat); $842.00(S). **Representatives:** Western States Associates Inc.
News Services: AP, RN, NYT, TMS, DJ, KRT.
Politics: Independent. **Established:** 1847.
Advertising not accepted: Handguns, ammunition and tobacco.

CORP. MGMT./GEN. MGMT.

Pres./Pub./CEO	Scott C. Smith
Sr. Vice Pres./Gen. Mgr.	Richard Malone
Sr. Vice Pres./Ed.	Ann Marie Lipinski
Vice Pres., Circ./Consumer Mktg.	Vincent Casanova
Vice Pres./Chief Tech. Officer	Darko Dejanovic
Vice Pres., Adv. Mktg./Sales	Ken DePaola
Vice Pres., Finance	Phil Doherty
Vice Pres., Human Resources	Janice Jacobs
Vice Pres., Devel.	Owen Youngman
Vice Pres./Dir., Opns.	Tony Hunter
Gen. Mgr., Chicago Tribune Interactive	Alison Scholly
Dir., Technical Devel.	Scott Tafelski
Dir., Technical Opns./Help Desk	Robert Trinchet
Dir., Client Servs.	Deepak Agarwal

ADVERTISING SALES MGMT.

Dir., Nat'l Adv.	Dan Dunn
Dir., Network Adv.	Ron Goldberg
Dir., Classified Adv.	Barbara Swanson
Dir., Major Accts.	Douglas Thomas
Dir., Preprint Adv.	John Wollney
Dir., Adv. Planning/Analysis	Margaret Durkin
Dir., Adv. Devel.	Kathy Manilla
Dir., Regl. Accounts	Steve Brooks
Dir., Group Sales/Mktg.	Robert Fleck
Dir., Devel.	Susan Zukrow
Dir., Devel.	Sue Klose

MARKETING MGMT.

Sr. Mgr., Multimedia Mktg.	Tom Garritano
Dir., Community Rel.	Frank Gihan
Dir., Brand Mktg.	Kelly Shannon

CIRCULATION MGMT.

Dir., Distr.	Shelia Davidson
Dir., Consumer Mktg.	Carrie Hoye
Dir., Circ. Planning/Opns.	Becky Brubaker

NEWS EXECUTIVES

Mng. Ed.	James O'Shea
Public Ed.	Don Wycliff
Deputy Mng. Ed., Features	Jim Warren
Deputy Mng. Ed., News	George de Lama
Deputy Mng. Ed., Opns.	Randy Weissman
Assoc. Mng. Ed., Electronic News	Mark Hinojosa
Assoc. Mng. Ed., Features	Mary Elson
Assoc. Mng. Ed., Financial News	Rob Karwath
Assoc. Mng. Ed., Foreign News	Tim McNulty
Assoc. Mng. Ed., Graphics/Design	Stacy Sweat
Assoc. Mng. Ed., Lifestyle	Geoff Brown
Assoc. Mng. Ed., Metropolitan News	Hanke Gratteau
Assoc. Mng. Ed., Nat'l News	Joycelynn Winnecke
Assoc. Mng. Ed., Photography	Bill Parker
Assoc. Mng. Ed., Sports	Dan McGrath
Assoc. Mng. Ed., Washington Bureau	Vicki Walton-James
Sr. Ed.	Tony Majeri
Sr. Ed., Recruiting	Sheila Solomom

EDITORIAL MGMT.

Books Ed.	Elizabeth Taylor
Editorial Page Ed.	Bruce Dold
Entertainment Ed.	Scott Powers
Foreign Ed.	Colin McMahon
Good Eating Ed.	Carol Haddix
Nat'l Ed.	Storer Rowley
Special Sections Ed.	Janet Franz
Sports Ed.	Bill Adee
Sunday Magazine Ed.	Elizabeth Taylor
Tempo Ed.	Tim Bannon
Travel Ed.	Randy Curwen
Womanews Ed.	Cassandra West

The Chicago Tribune

Figure A6: Newspaper sections by readership (1998-2001)

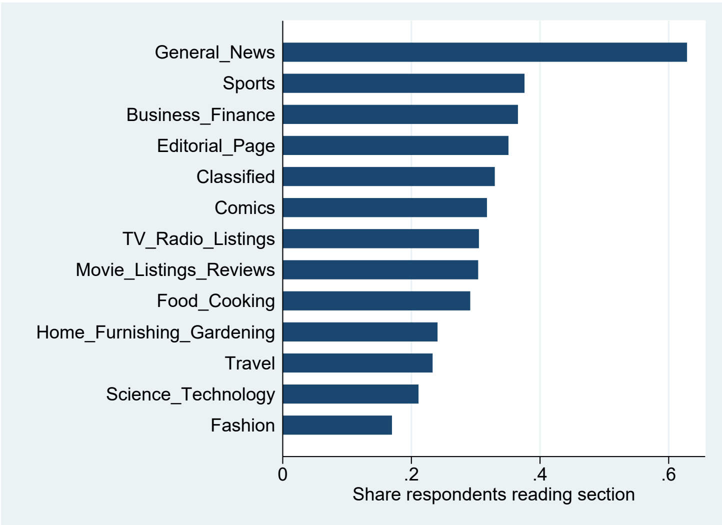

Source: GfK-MRI Survey

Craigslist markets

In our baseline analysis we assume that Craigslist markets consist of the county (or counties) containing the locality indicated in the website's url. In this section we discuss an alternative approach which relies on the locations indicated in ads posted on the respective websites.

To do so, we retrieve the snapshots of each website available from `https://archive.org//`, and code the exact location of all the ads posted on the first page of the "housing", "jobs", and "sales" sections. Here we focus on the ads post in the first two years after the entry. We then match the resulting locations to a comprehensive list of towns, cities, and counties (if the location includes the word "county") in the same or a neighboring state. Finally, we consider all counties that account for at least 5% of the ads as part of what we define as the website's "broad" market.

Figure A7 depicts the geographic distribution of ads posted on `https://brainerd.craigslist.org/` in the 1st and 2nd year after the opening of the website. In this case, the "core" market is represented by the central county (Crow Wing County) containing the city of Brainerd. This "core" county accounts for over 80% of total ads, while the "broad market" includes five additional neighboring counties. This is a typical pattern in our data: on average the "core" market accounts for 73% (median 76%) of posted ads once we exclude outliers.

Figure A7: Distribution of ads posted on `https://brainerd.craigslist.org/`

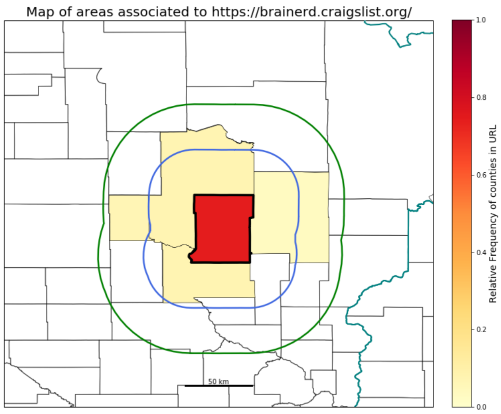

Geographic distribution of the location of ads posted in the housing, jobs and sales sections of `https://brainerd.craigslist.org/` in years 1 and 2 after the website opening. Source: Internet Archive.

Newspaper markets

In our baseline analysis we assume that local newspaper markets consist of the county in which they are headquartered. In this section we discuss an alternative approach which relies on the availability of geographically dis-aggregated circulation data.

This method consists in matching a newspaper to all counties where it is read in proportion to circulation. Zip-code-level circulation data are available from the Alliance for Audited Media (AAM) for 2002, which we use to construct a weighted measure of CL availability by newspaper-year. The weights in this "broad" measure of CL availability are the fraction of the paper's subscribers in each county.[33] However, AAM data covers only about 40% of the papers in the E&P sample. For the newspapers for which no zip-code circulation data is available, we assign 100% of the circulation to the county where the paper's HQ is located.[34]

Figure A8 shows the geographic distribution of circulation for a newspaper in our sample, the Brainerd Dispatch, with darker colors representing higher values. Crow Wing County where the newspaper's HQ is located, shown in blue, accounts for 81% of the total.[35] The median paper in our AAM data has about 85% of its total circulation in the headquarters county once we exclude outliers.

Figure A8: Distribution of circulation of the Brainerd Dispatch

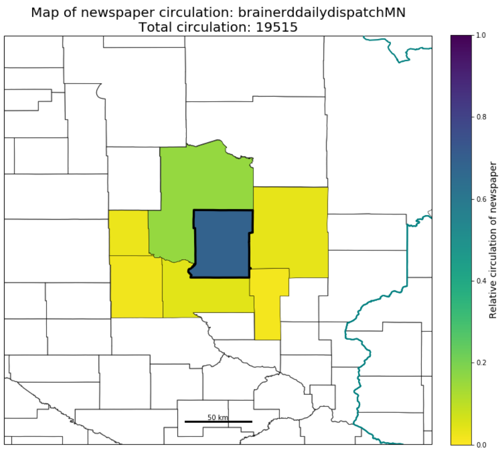

Geographic distribution of the circulation of the *Brainerd Dispatch* in 2002. Source: Alliance for Audited Media.

Validating the classified manager proxy

We validate the classified manager indicator as a proxy for classified intensity using data from the website *Newspapers.com*, which archives digitized historical copies of newspapers. We

[33]We measure geographically disaggregated circulation only once, in 2002, and hence year-to-year variation is driven entirely by changes in CL availability and not by changes in circulation patterns.

[34]The papers that are missing from AAM are generally smaller papers and, if anything, less likely to have circulation beyond the county boundaries than the papers that appear in AAM. Papers which appear in AAM had median circulation in 2002 of 67K, compared to 14K for papers not appearing AAM. Hence, we believe that assigning all circulation to the headquarters county is a good approximation for these papers.

[35]Similarly to CL ads, we exclude outlier counties that account for less than 5% of total circulation.

located 262 papers in our dataset which appear in the Newspapers.com archive. For each of these papers, we sampled the edition of the paper published on the first Sunday of each month in all years from 1995 until 2010, substituting another day when the Sunday edition was not available.

We measure classified intensity as the number of pages on which the term "Classified" appears, divided by the total number of pages in the issue. We collected this measure for a total of 43,165 issues across the 262 papers available in the Newspapers.com archive. Prior to the entry of Craigslist in a market, the median issue in our sample had 8 pages of classified advertising, or 15% of the issue's total page count.

We examine cross-sectional variation in classified intensity prior to Craigslist entry. Table A1 shows the results of regressions where the outcome is the fraction of pages in a newspaper issue that contain classified advertising. The left two columns use all days with weekday fixed effects to control for cyclical variation in classified intensity, while the last two use Sundays only (traditionally the biggest day for classified advertising). Columns (2) and (4) additionally add controls for the newspaper's circulation and the population of the county in which it operates, to test whether the classified manager dummy merely picks up larger papers.

Table A1: Share of Pages Devoted to Classified Ads in Pre-CL Period, by Existence of Classified Manager in 2000.

	Classified Page Share		Classified Page Share (Sun.)	
	(1)	(2)	(3)	(4)
Classif. Mgr.	0.029*	0.023	0.034*	0.027
	(0.015)	(0.015)	(0.018)	(0.017)
Pop / Circ controls	No	Yes	No	Yes
Year FEs	Yes	Yes	Yes	Yes
Day-of-week FEs	Yes	Yes	No	No
Observations	31,273	31,273	24,368	24,368
R^2	0.01	0.03	0.01	0.02
Mean dependent variable	0.20	0.20	0.21	0.21
Number of newspapers	262	262	220	220

Regressions of the fraction, by newspaper issue, of pages devoted to classified advertising on an indicator for the presence of a classified manager in 2000, in the pre-CL period. Data from issues published in all years prior to the year of CL entry in the newspaper's market are included. OLS regressions in all columns. Standard errors clustered by newspaper.
Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The table shows that the fraction of pages devoted to classified ads was on average between 2 and 4 percentage points higher at papers that had a classified manager in 2000, prior to Craigslist entry. Coefficient magnitudes are comparable but fall below conventional significance thresholds when population and circulation controls are added.

Number of ISPS

Throughout the paper, we use the number of ISPs registered in a county as a proxy for broadband internet penetration, a measure used in prior work (Larcinese and Miner 2018; **?**; Lelkes et al. 2017). To further validate the number of ISPs as a proxy of local Internet penetration, we examine its correlation with self-reported Internet access from both the NAES and GfK-MRI surveys. Figure A9 shows a binned scatter plot comparing the number of ISPs with the fraction of respondents in the county who report having internet access at home (in GfK-MRI) or either at home or at work (in NAES), along with the OLS line of best fit. Both datasets show a strong positive relationship between the two.

Figure A9: Number of ISPs as a proxy for local Internet penetration

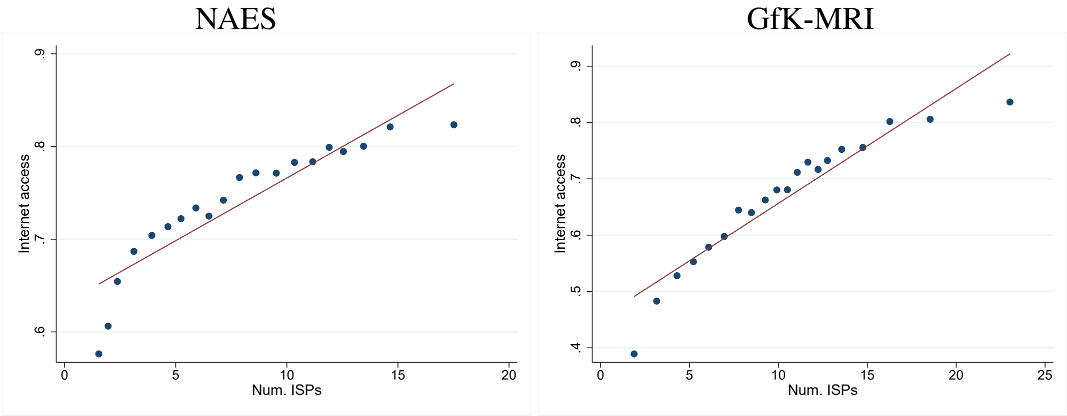

Binned scatter plot for the relationship between number of ISPs available in the county of the respondent and self-reported Internet access at home (GfK-MRI) and at home or at work (NAES).

Details on Newspaper Content Processing

This section contains details on the procedures used to process raw text content from the newspapers in our sample to construct lower-dimensional representations of the content. Source data are from the *NewsBank* database. We conduct two main kinds of processing on text data: keyword searches and topic modeling. Keyword searches use the full database containing more than 100M full-text articles, while our topic model uses a smaller random sample of about 2M articles consisting of all articles published on 10 randomly sampled dates in each newspaper-year between 1999 and 2010. The topic-modeling sample limits to the first paragraph of text, plus the headline.

Politician Names Our first set of keyword searches look for the names of representatives in the US House and Senate. We use a list of representatives from `https://github.com/unitedstates/congress-legislators`, and for each member construct a (case-insensitive) regular expression of the form `"(congres.*|rep.*) FIRSTNAME LASTNAME"` or "senat.* FIRSTNAME LASTNAME". This expression matches strings like "Rep. Adam Smith" or "Congressman Adam Smith" but not "Adam Smith" alone. We require the inclusion of the title to cut down on false positives, as many members of Congress have common names. This does introduce the possibility of false negatives, but we have found that articles covering a member usually include the title and full name at first mention before switching to a shorter form like "Mr. Smith".

We count the number of *articles*[36] in which the pattern described above appears on each newspaper-day, and then aggregate to the level of congressman by newspaper by year.

[36]I.e., each article that mentions the congressperson at all counts as 1, regardless of how many times the congressman is referenced in the article.

Figure A10: Top ten words associated with each CorEx topic

```
0: presid,feder,govern,compani,tax,washington,percent,increas,pai,billion
1: council,mayor,board,plan,student,educ,fund,commun,project,program
2: repres,senat,congress,republican,elect,democrat,vote,candid,polit,gov
3: intern,war,foreign,iraq,militari,movi,film,american,soldier,terrorist
4: man,kill,injuri,injur,accid,crash,woman,diseas,victim,suffer
5: music,art,food,festiv,featur,concert,event,artist,band,holidai
6: car,vehicl,driver,road,truck,traffic,highwai,drive,mile,street
7: di,born,funer,son,daughter,church,surviv,servic,cemeteri,obituari
8: game,team,coach,win,season,plai,victori,footbal,score,player
9: polic,charg,court,arrest,judg,investig,attornei,accus,sheriff,suspect
```

Topic model Our method for extracting the topical coverage of affected newspapers follows Gallagher et al.'s 2017b Correlation Explanation (CorEx) method. This is a semi-supervised method that allows input of a minimal set of "anchor" words, and then finds topics by searching for groups of words that co-occur with the anchors. We apply this method to the text of a random sample of 2 million articles from the NewsBank corpus.

We use the semi-supervised method rather than the more traditional unsupervised Latent Dirichlet Allocation (LDA) because it allows us to focus on specific topics of interest. We are interested in separating various dimensions of political news coverage: coverage related to local, congressional, national and foreign politics. We seed separate anchors for these 4 topics, and run the CorEx model with 10 topics in total. Figure A10 presents the resulting topics, as described by their most representative words, with our anchor-words highlighted in bold. The 10 resulting topics can be labeled as follows: local politics, congressional politics, national politics, foreign politics, entertainment, health / family, weather, crime, obituaries.

For each of the 2 million articles in the corpus, the CorEx model outputs a set of 10 unconditional probabilities for the article belonging to that a topic. Importantly, these probabilities do not necessarily sum to 1 - an article can simultaneously belong to more than one topic, or to none. To examine the effects of CL's entry, we aggregate the distribution of probabilities by newspaper and year, and estimate the standard diff-in-diff equations specified in section 3.3, with the average probability for each one of the 10 topics as dependent variable.

Appendix B: Additional results

Table B1: Number of Jobs: Robustness to controlling for newspaper size

	Dependent variable: Newspaper number of jobs							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	-1.049*	-0.523	-0.753	-0.893	0.303	0.667	0.451	0.278
	(0.595)	(0.558)	(0.596)	(0.735)	(0.482)	(0.484)	(0.508)	(0.599)
Post-CL × Circ. 2000	-0.016	-0.011	-0.009	-0.007	-0.012	-0.007	-0.006	-0.004
	(0.013)	(0.014)	(0.014)	(0.015)	(0.013)	(0.014)	(0.015)	(0.016)
Post-CL × Classified Mgr.					-3.006***	-2.657***	-2.667***	-2.623***
					(0.668)	(0.594)	(0.613)	(0.698)
Log population, #ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	22516	22516	22515	21762	22499	22499	22498	21745
Number of newspapers	1506	1506	1506	1460	1504	1504	1504	1458
R^2	0.90	0.91	0.91	0.92	0.90	0.91	0.91	0.92
Mean dependent variable	21.38	21.38	21.38	21.22	21.39	21.39	21.39	21.24

Regressions of number of jobs by newspaper and year on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B2: Number of Staff

	Dependent variable: Newspaper number of employees							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	-1.701***	-1.012***	-1.107***	-1.086***	-0.112	0.260	0.131	0.070
	(0.339)	(0.329)	(0.344)	(0.338)	(0.341)	(0.339)	(0.345)	(0.349)
Post-CL × Classified Mgr.					-3.010***	-2.490***	-2.405***	-2.299***
					(0.513)	(0.480)	(0.477)	(0.516)
Log population, #ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	22797	22797	22796	22044	22583	22583	22582	21832
Number of newspapers	1540	1540	1540	1492	1505	1505	1505	1459
R^2	0.92	0.92	0.92	0.93	0.92	0.92	0.92	0.93
Mean dependent variable	17.84	17.84	17.84	17.68	17.90	17.90	17.90	17.74

Regressions of number of employees by newspaper and year on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B3: Number of Pages

	Dependent variable:							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	0.293	0.250*	0.125	0.259*	0.140	0.100	-0.026	0.136
	(0.191)	(0.139)	(0.171)	(0.151)	(0.128)	(0.126)	(0.162)	(0.213)
Post-CL × Classified Mgr.					0.286	0.292	0.289	0.238
					(0.316)	(0.286)	(0.294)	(0.347)
Log population, #ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	21637	21637	21637	20921	21509	21509	21509	20793
Number of newspapers	1471	1471	1471	1426	1450	1450	1450	1406
R^2	0.97	0.97	0.97	0.98	0.97	0.97	0.97	0.98
Mean dependent variable	28.50	28.50	28.50	28.31	28.52	28.52	28.52	28.34

Regressions of number of pages published by newspaper and year on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B4: Subscription price

	Dependent variable: Price							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	-0.273	-0.872	-0.984	-1.103	-1.372	-2.021*	-2.021	-2.781*
	(0.896)	(0.939)	(1.033)	(1.259)	(1.075)	(1.114)	(1.352)	(1.490)
Post-CL × Classified Mgr.					1.966	2.142	1.911	3.403**
					(1.384)	(1.425)	(1.519)	(1.628)
Log population, #ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	19643	19643	19642	18934	19424	19424	19423	18720
Number of newspapers	1431	1431	1431	1387	1393	1393	1393	1349
R^2	0.93	0.94	0.94	0.95	0.93	0.93	0.94	0.94
Mean dependent variable	118.66	118.66	118.67	118.36	118.80	118.80	118.80	118.48

Regressions of yearly subscription price by newspaper and year on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B5: Topic model: Probability of political topics

	(1) president	(2) local	(3) congress	(4) foreign	(5) president	(6) local	(7) congress	(8) foreign
Post-CL	-0.009**	-0.009	-0.006**	-0.001	-0.000	-0.002	-0.001	0.005
	(0.004)	(0.006)	(0.003)	(0.003)	(0.005)	(0.007)	(0.003)	(0.004)
Post-CL × Classified Mgr.					-0.017***	-0.013*	-0.011**	-0.012***
					(0.006)	(0.008)	(0.005)	(0.004)
Full county controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7178	7178	7178	7178	7129	7129	7129	7129
Number of newspapers	869	869	869	869	862	862	862	862
R^2	0.52	0.47	0.41	0.54	0.52	0.47	0.40	0.54
Mean dependent variable	0.21	0.31	0.10	0.10	0.21	0.31	0.10	0.10

Regressions of the average probability of an article covering a particular topic by newspaper and year on an indicator for the availability of a local Craigslist website, and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B6: Topic model: Probability of other topics

	(1) accidents	(2) events	(3) traffic	(4) obituaries	(5) sports	(6) crime	(7) accidents	(8) events	(9) traffic	(10) obituaries	(11) sports	(12) crime
Post-CL	-0.000	-0.009**	-0.005*	0.009	0.013**	-0.001	-0.003	-0.007	-0.005	0.002	0.007	-0.003
	(0.003)	(0.005)	(0.003)	(0.007)	(0.006)	(0.003)	(0.004)	(0.006)	(0.004)	(0.010)	(0.007)	(0.003)
Post-CL $\times$ Classified Mgr.							0.006	-0.003	-0.001	0.013	0.012	0.003
							(0.006)	(0.007)	(0.006)	(0.012)	(0.008)	(0.004)
County controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	7178	7178	7178	7178	7178	7178	7129	7129	7129	7129	7129	7129
Number of newspapers	869	869	869	869	869	869	862	862	862	862	862	862
R^2	0.44	0.41	0.39	0.56	0.44	0.44	0.44	0.41	0.39	0.56	0.45	0.44
Mean dependent variable	0.12	0.17	0.14	0.15	0.21	0.11	0.12	0.17	0.14	0.15	0.21	0.11

Regressions of topic probabilities aggregated by newspaper and year on an indicator for the availability of a local Craigslist website and its interaction with an indicator for the presence of a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B7: Mentions of national party leaders

	Dependent variable: Articles mentioning the president and party leaders (IHS)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	0.143*	0.066	0.078	0.002	0.113	0.080	0.085	0.015
	(0.074)	(0.076)	(0.087)	(0.115)	(0.103)	(0.100)	(0.108)	(0.134)
Post-CL × Classified Mgr.					0.063	-0.037	-0.030	-0.041
					(0.129)	(0.130)	(0.127)	(0.171)
Total articles (IHS)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Log population, # ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Newspaper FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State × Year FEs	No	No	Yes	No	No	No	Yes	No
DMA × Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	6716	6716	6680	6097	6667	6667	6631	6043
Number of newspapers	825	825	822	775	818	818	815	767
R^2	0.83	0.84	0.86	0.88	0.83	0.84	0.86	0.88
Mean dependent variable	1.79	1.79	1.79	1.84	1.79	1.79	1.79	1.84

Regressions of the (IHS-transformed) number of articles mentioning the names of a set of national party leaders by newspaper and year on an indicator for the availability of a local Craigslist website, and its interaction with an indicator for the presence of a classified manager at baseline. All regressions control for the (IHS-transformed) total number of articles recorded by Newsbank by newspaper and year. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B8: Self-reported newspaper readership: Heterogeneity

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	News propensity				Classified propensity			
	Low		High		Low		High	
Post-CL	-0.011*	-0.010	-0.017***	0.007	-0.023***	-0.005	-0.010	-0.000
	(0.007)	(0.010)	(0.007)	(0.010)	(0.007)	(0.010)	(0.007)	(0.010)
Post-CL × Classified Mgr.		-0.002		-0.032***		-0.024**		-0.013
		(0.012)		(0.011)		(0.012)		(0.012)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	126747	125255	126752	126173	126747	125637	126756	125795
Number of counties	786	784	790	788	787	785	792	790
R^2	0.05	0.05	0.09	0.09	0.11	0.11	0.07	0.07
Mean dependent variable	0.32	0.32	0.52	0.52	0.47	0.47	0.38	0.38

Regressions of self-reported newspaper readership on an indicator for the availability of a local Craigslist website in the county of the respondent, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. The sample is split by the median respondent-level propensity to read the news section / classified section. Respondent controls include age bins, an indicator for college degree and race. County controls include contemporaneous log population and number of Internet service providers, as well as share urban population, pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share - all measured in 2000 and interacted with year FEs. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B9: Self-reported media consumption: GfK

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Newspaper, national		TV news		Radio news		Online news	
Post-CL	0.009**	0.007	-0.006	0.004	0.005	0.002	0.003	0.005
	(0.004)	(0.005)	(0.007)	(0.009)	(0.006)	(0.007)	(0.004)	(0.007)
Post-CL × Classified Mgr.		0.003		-0.014		0.004		-0.004
		(0.005)		(0.009)		(0.008)		(0.007)
Respondent controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	253513	251442	253513	251442	253513	251442	253513	251442
Number of counties	792	790	792	790	792	790	792	790
R^2	0.14	0.14	0.08	0.08	0.09	0.09	0.16	0.16
Mean dependent variable	0.09	0.09	0.70	0.70	0.16	0.16	0.22	0.22

Regressions of self-reported media consumption on an indicator for the availability of a local Craigslist 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. County controls include contemporaneous log population and number of Internet service providers, as well as share urban population, pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share - all measured in 2000 and interacted with year FEs. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B10: Self-reported media consumption: NAES

	(1)	(2)	(3)	(4)	(5)	(6)
	Newspaper, national		TV news		Radio news	
Post-CL	0.030**	-0.000	0.002	0.008	0.004	0.000
	(0.015)	(0.013)	(0.004)	(0.009)	(0.006)	(0.008)
Post-CL × Classified Mgr.		0.037		-0.009		0.005
		(0.024)		(0.009)		(0.008)
Respondent characteristics	Yes	Yes	Yes	Yes	Yes	Yes
County controls	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	105989	105173	107778	106981	149111	148083
Number of counties	1207	1203	1207	1203	1208	1204
R^2	0.21	0.20	0.04	0.04	0.05	0.05
Mean dependent variable	0.03	0.03	0.92	0.92	0.37	0.37

Regressions of self-reported media consumption on an indicator for the availability of a local Craigslist 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. County controls include contemporaneous log population and number of Internet service providers, as well as share urban population, pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share - all measured in 2000 and interacted with year FEs. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B11: Self-reported media consumption: NAES, continued

	(1)	(2)	(3)	(4)	(5)	(6)
	Newspaper, national days per wk		TV news days per wk		Radio news days per wk	
Post-CL	0.153**	0.004	0.043	0.094	0.024	0.003
	(0.073)	(0.069)	(0.038)	(0.074)	(0.026)	(0.039)
Post-CL $\times$ Classified Mgr.		0.186		-0.068		0.030
		(0.117)		(0.089)		(0.038)
Respondent characteristics	Yes	Yes	Yes	Yes	Yes	Yes
County controls	Yes	Yes	Yes	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	107856	107054	107939	107138	149111	148083
Number of counties	1207	1203	1207	1203	1208	1204
R^2	0.09	0.09	0.10	0.10	0.05	0.05
Mean dependent variable	0.15	0.15	5.04	5.04	1.51	1.51

Regressions of self-reported media consumption on an indicator for the availability of a local Craigslist 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. County controls include contemporaneous log population and number of Internet service providers, as well as share urban population, pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share - all measured in 2000 and interacted with year FEs. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B12: Visits of news-related web domains: Comscore

	Dependent variable: Visits of news-related web domains (IHS)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Post-CL	-0.021	-0.026	0.001	0.037	-0.035	-0.026	0.010	0.066
	(0.028)	(0.028)	(0.032)	(0.035)	(0.038)	(0.039)	(0.042)	(0.047)
Post-CL $\times$ Classified Mgr.					0.032	0.008	-0.010	-0.053
					(0.045)	(0.047)	(0.048)	(0.051)
Total Comscore visits (IHS)	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Log population, $\#$ ISPs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. $\times$ Year FEs	No	Yes	Yes	Yes	No	Yes	Yes	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State $\times$ Year FEs	No	No	Yes	No	No	No	Yes	No
DMA $\times$ Year FEs	No	No	No	Yes	No	No	No	Yes
Observations	8303	8303	8303	7975	8280	8280	8280	7952
Number of counties	1214	1214	1214	1167	1210	1210	1210	1163
R^2	0.88	0.88	0.89	0.90	0.88	0.88	0.89	0.90
Mean dependent variable	7.89	7.89	7.89	7.89	7.89	7.89	7.89	7.89

Regressions of number of visits of news-related web domains (IHS-transformed) by county and year an indicator for the availability of a local Craigslist website in the county of the respondent, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. OLS regressions in all columns. Standard errors clustered by CL-area. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B13: Turnout in Presidential Elections

	Dependent variable: Turnout in Presidential elections					
	(1)	(2)	(3)	(4)	(5)	(6)
Post-CL	-0.016***	-0.000	-0.001	-0.003	0.007**	0.002
	(0.005)	(0.003)	(0.002)	(0.005)	(0.003)	(0.002)
Post-CL $\times$ Classified Mgr.				-0.016**	-0.010**	-0.004
				(0.007)	(0.004)	(0.003)
Log population, # ISPs	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. $\times$ Year FEs	No	Yes	Yes	No	Yes	Yes
State $\times$ Year FEs	No	No	Yes	No	No	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4855	4855	4855	4839	4839	4839
Number of counties	1214	1214	1214	1210	1210	1210
R^2	0.97	0.99	0.99	0.97	0.99	0.99
Mean dependent variable	0.48	0.48	0.48	0.48	0.48	0.48

Regressions of midterm turnout county and election year on an indicator for the availability of a local Craigslist website, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. Turnout is defined as the ratio of votes cast to Census-estimated voting-age population in the county-year. Years included are 1998, 2002, 2006, and 2010. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. Observations are weighted by voting-age population. OLS regressions in all columns. Standard errors clustered by CL-area. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table B14: Turnout in House Elections

	Dependent variable: Turnout in House elections					
	(1)	(2)	(3)	(4)	(5)	(6)
Post-CL	-0.019***	-0.006	0.000	-0.024***	-0.012*	-0.006
	(0.007)	(0.006)	(0.005)	(0.008)	(0.007)	(0.005)
Post-CL × Classified Mgr.				0.008	0.008	0.008*
				(0.008)	(0.006)	(0.005)
Log population, # ISPs	Yes	Yes	Yes	Yes	Yes	Yes
2000 county char. × Year FEs	No	Yes	Yes	No	Yes	Yes
State × Year FEs	No	No	Yes	No	No	Yes
County FEs	Yes	Yes	Yes	Yes	Yes	Yes
Year FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4856	4856	4856	4840	4840	4840
Number of counties	1214	1214	1214	1210	1210	1210
R^2	0.79	0.83	0.90	0.78	0.82	0.90
Mean dependent variable	0.37	0.37	0.37	0.37	0.37	0.37

Regressions of midterm turnout county and election year on an indicator for the availability of a local Craigslist website, and its interaction with the circulation-weighted share of newspapers with a classified manager at baseline. Turnout is defined as the ratio of votes cast to Census-estimated voting-age population in the county-year. Years included are 1998, 2002, 2006, and 2010. County characteristics for the year 2000 include pct. college educated, pct. rental, median age, share white/ black/ hispanic, income per capita, unemployment rate, presidential turnout and Republican vote share. Observations are weighted by voting-age population. OLS regressions in all columns. Standard errors clustered by CL-area. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Bibliography

Alesina, A., Miano, A., and Stantcheva, S. (2020). The polarization of reality. *AEA Papers and Proceedings*, 110:324–328.

Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review*, 110(3):629–76.

Angelucci, C., Cage, J., , and Sinkinson, M. (2020). Media competition and news diets. *NBER Working Paper No.26782*.

Balles, P., Matter, U., and Stutzer, A. (2018). Special Interest Groups Versus Donors and the Political Economics of Attention.

Baum, M. (2004). How public opinion constrains the use of force: The case of operation restore hope. *Presidential Studies Quarterly*, 34:187–226.

Besley, T. and Burgess, R. (2002a). The political economy of government responsiveness: Evidence from india. *Quarterly Journal of Economics*, 117(4):1415â1451.

Besley, T. and Burgess, R. (2002b). The political economy of government responsiveness: Theory and evidence from india. *Quarterly Journal of Economics*, 117 (4):1415â1451.

Bhuller, M., Havnes, T., McCauley, J., and Mogstad, M. (2020). How the internet changed the market for print media.

Bolton, A. and Thrower, S. (2021). Legislative constraints, ideological conflict, and the timing of executive unilateralism. *Legislative Studies Quarterly*, n/a(n/a).

Bonica, A. (2014). Mapping the Ideological Marketplace. *American Journal of Political Science*, 58(2):367–386.

Bonica, A. (2016). Database on Ideology, Money in Politics, and Elections. Stanford, CA: Stanford University Libraries. https://data.stanford.edu/dime.

Cage, J., Herve, N., and Viaud, M.-L. (2020). The Production of Information in an Online World: Is Copy Right? *Review of Economic Studies*, forthcoming.

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(4):1094â1136.

Chan, J. and Ghose, A. (2014). Internet's dirty secret: assessing the impact of online intermediaries on hiv transmission. *MIS Quarterly*, 38(4).

Chiang, C.-F. and Knight, B. (2011). Media bias and influence: Evidence from newspaper endorsements. *Review of Economic Studies*, 78(3):795–820.

Chiou, F.-Y. and Rothenberg, L. S. (2014). The elusive search for presidential power. *American Journal of Political Science*, 58(3):653–668.

Christenson, D. P. and Kriner, D. K. (2017a). Constitutional qualms or politics as usual? the factors shaping public support for unilateral action. *American Journal of Political Science*, 61 (2):335–349.

Christenson, D. P. and Kriner, D. L. (2017b). Mobilizing the public against the president: Congress and the political costs of unilateral action. *American Journal of Political Science*, 61:769–785.

Christenson, D. P. and Kriner, D. L. (2017c). The specter of supreme court criticism: Public opinion and unilateral action. *Presidential Studies Quarterly*, 47(3):471–494.

Christenson, D. P. and Kriner, D. L. (2019). Does public opinion constrain presidential unilateralism. *American Political Science Review*, 113(4):1071â1077.

Couttenier, M. and Hatte, S. (2016). Mass media effects on non-governmental organizations. *Journal of Development Economics*, 123:57–72.

Cunningham, S., DeAngelo, G., and Tripp, J. (2019). Craigslist's effect on violence against women.

Darr, J. P., Hitt, M. P., and Dunaway, J. L. (2018). Newspaper closures polarize voting behavior. *Journal of Communication*, 68(6):1007–1028.

de Chaisemartin, C. and D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, forthcoming.

De los Santos, B., HortaÃ§su, A., and Wildenbeest, M. R. (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review*, 102(6):2955â80.

DellaVigna, S. and Gentzkow, M. (2010). Persuasion: Empirical evidence. *Annual Review of Economics*, 2:643–669.

DellaVigna, S. and Kaplan, E. (2007). The Fox News Effect: Median Bias and Voting. *Quarterly Journal of Economics*, 112(3):1187–1234.

DellaVigna, S. and Pollet, J. (2009). Investor inattention and friday earnings announcements. *The Journal of Finance*, 64:709–749.

Djourelova, M. (2020). Media Persuasion through Slanted Language: Evidence from the Coverage of Immigration. *BGSE Working Paper 1177*.

Djourelova, M. and Durante, R. (2020). Media Attention and Strategic Timing in Politics: Evidence from US Presidential Executive Orders. *CEPR Discussion Paper 13961*.

Djourelova, M., Durante, R., and Martin, G. J. (2021). The Impact of Online Competition on Local Newspapers: Evidence from the Introduction of Craigslist. *CEPR Discussion Paper 16130*.

Dorn, D., Hanson, G., and Majlesi, K. (2020). Importing political polarization? The electoral consequences of rising trade exposure. *American Economic Review*, 110(10):3139–83.

Durante, R. and Knight, B. (2012). Partisan control, media bias, and viewer responses: Evidence from berlusconi's italy. *Journal of the European Economic Association*, 10(3):451–481.

Durante, R. and Zhuravskaya, E. (2018). Attack when the world is not watching? u.s. media and the Israeli-Palestinian conflict. *Journal of Political Economy*, 126(3):1085–1133.

Eisensee, T. and Strömberg, D. (2007). News droughts, news floods, and US disaster relief. *Quarterly Journal of Economics*, 122(2):693–728.

Falck, O., Gold, R., and Heblich, S. (2014). E-lections: Voting behavior and the internet. *American Economic Review*, 117(4):2238–2265.

FCC (2016). Federal communications commission report fcc 16-107. Technical report.

Fenby, J. (1986). *The International News Services*. Schocken Books.

Ferraz, C. and Finan, F. (2008). Exposing corrupt politicians: The effects of brazil's publicly released audits on electoral outcomes. *Quarterly Journal of Economics*, 123(2):703–745.

Gabaix, X., Laibson, D., Moloche, G., and Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96 (4):1046–1068.

Gallagher, R. J., Reing, K., Kale, D., and Steeg, G. V. (2017a). Anchored Correlation Explanation: Topic Modeling with Minimal Domain Knowledge. *Transactions of the Association for Computational Linguistics*, 5:529–542.

Gallagher, R. J., Reing, K., Kale, D., and Steeg, G. V. (2017b). Anchored correlation explanation: Topic modeling with minimal domain knowledge. *Transactions of the Association for Computational Linguistics*, 5:529â542.

Gao, P., Lee, C., and Murphy, D. (2020). Financing dies in darkness? the impact of newspaper closures on public finance. *Journal of Financial Economics*, 135(2):445–467.

Garz, M. and Maass, S. (2020). Cartels in the european union, antitrust action, and public attention.

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(5):2092â2135.

Gentzkow, M., Kelly, B. T., and Taddy, M. (2018). Text as data. *Journal of Economic Literature*, forthcoming.

Gentzkow, M. and Shapiro, J. M. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2).

Gentzkow, M. and Shapiro, J. M. (2010a). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica*, 78:35–71.

Gentzkow, M. and Shapiro, J. M. (2010b). 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(7).

Gibson, J. (1999). Political timing: A theory of politician's timing of events. *Journal of Theoretical Politics*, 11(4):471–496.

Grimmer, J. and Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, 21(3):267–297.

Hall, A. B. (2015). What happens when extremists win primaries? *American Political Science Review*, 109(1):18–42.

Hamilton, J. (2004). *All the news that's fit to sell: How the market transforms information into news*. Princeton University Press.

Hopkins, D. J. (2018). *The increasingly United States: How and why American political behavior nationalized*. University of Chicago Press.

Howell, W. (2003). *Power without Persuasion: The Politics of Direct Presidential Action*. Princeton, NJ: Princeton University Press.

Kaplan, E., Spenkuch, J. L., and Yuan, H. (2018). Natural Disasters, Moral Hazard, and Special Iinterests in Congress.

Kling, J. R., Liebman, J. B., and Katz, L. F. (2007). Experimental Analysis of Neighborhood Effects. *Econometrica*, 75(1):83–119.

Kroft, K. and Pope, D. G. (2014). Does online search crowd out traditional search? evidence from craigslist. *Journal Labor Economics*, 32(2):259–303.

Larcinese, V. and Miner, L. (2018). Was obama elected by the internet? broadband diffusion and voters' behavior in us presidential elections.

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(1):5–20. Publisher: Wiley Online Library.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American economic review*, 111(3):831–70.

Luntz, F. (2007). *Words That Work: It's Not What You Say, It's What People Hear*.

Mahone, J., Wang, Q., Napoli, P., Weber, M., and McCullough, K. (2019). Who's Producing Local Journalism? Assessing Journalistic Output Across Different Outlet Types. Technical report.

Manacorda, M. and Tesei, A. (2020). Liberation technology: Mobile phones and political mobilization in africa. *Econometrica*, 88(2):533–567.

McChesney, R. W. and Nichols, J. (2011). *The death and life of American journalism: The media revolution that will begin the world again*. Bold Type Books.

Merolla, C. H. J. L. and Ramakrishnan, S. K. (2016). *Framing Immigrants: News Coverage, Public Opinion, and Policy*. Russell Sage Foundation.

Moe, T. M. and Howell, W. G. (1999). The presidential power of unilateral action. *Journal of Law, Economics, and Organization*, 15:132–197.

Moskowitz, D. J. (2021). Local news, information, and the nationalization of US Elections. *American Political Science Review*, 115(1):114–129.

Peterson, E. (2021). Paper Cuts: How Reporting Resources Affect Political News Coverage. *American Journal of Political Science*, 65(2):443–459.

Petrova, M. (2011). Newspapers and parties: How advertising revenues created an independent press. *American Political Science Review*, pages 790–808.

Posner, E. and Vermeule, A. (2010). *The Executive Unbound: After the Madisonian Republic*. New York: Oxford University Press.

Potter, R. A. (2017). Slow-rolling, fast-tracking, and the pace of bureaucratic decisions in rulemaking. *The Journal of Politics*, 79(3):841–855.

Potter, R. A. (2019). *Bending the Rules: Procedural Politicking in the Bureaucracy*. University of Chicago Press.

Puglisi, R. and Snyder, J. M. (2011). Newspaper coverage of political scandals. *The Journal of Politics*, 73(3):931–950.

Reeves, A. and Rogowski, J. C. (2015). Public opinion toward presidential power. *Presidential Studies Quarterly*, 45(4):742–759.

Reeves, A. and Rogowski, J. C. (2016). Unilateral powers, public opinion, and the presidency. *Journal of Politics*, 78(1):137–151.

Reeves, A. and Rogowski, J. C. (2018). The public cost of unilateral action. *American Journal of Political Science*, 62:424–440.

Reeves, A., Rogowski, J. C., Seo, M. H., and Stone, A. R. (2017). The contextual determinants of support for unilateral action. *Presidential Studies Quarterly*, 47(3):448–470.

Seamans, R. and Zhu, F. (2014). Responses to entry in multi-sided markets: The impact of craigslist on local newspapers. *Management Science*, 60(2):265–540.

Sellers, M. (2017). Gubernatorial use of executive orders: Unilateral action and policy adoption. *Journal of Public Policy*, 37(3):315–339.

Snyder, J. M. and Strömberg, D. (2010). Press Coverage and Political Accountability. *Journal of Political Economy*, 118(2):355–408.

Snyder, Jr, J. M. and Strömberg, D. (2010). Press coverage and political accountability. *Quarterly Journal of Economics*, 117 (4):1415â1451.

Starkman, D. (2014). *The watchdog that didnât bark: The financial crisis and the disappearance of investigative journalism*. Columbia University Press.

Trussler, M. (2020). Get Information or Get in Formation: The Effects of High-Information Environments on Legislative Elections. *British Journal of Political Science*, pages 1–21.