What Drives Demand for Media Slant?*

Marcel Garz[†] Gaurav Sood[‡] Daniel F. Stone[§] Justin Wallace[¶]

July 28, 2017

Abstract

People tend to prefer politically congenial news—news that confirms and supports their prior beliefs and preferences. Many papers assume that this preference is driven by psychological forces: that we want news that "feels good" and/or to avoid news that "feels bad." Other papers have proposed models showing how the preference for congenial news can stem from this news having greater perceived or actual instrumental information value. We assess these theories empirically by studying how variation in congeniality in news across and within outlets affects demand for news. We exploit two types of news stories that exhibit both types of this variation: horse race stories and stories evaluating the winner of presidential debates. We use both survey-experiment data and observational web data from a variety of outlets. In the survey-experiments, we find some evidence supporting the psychology theory, particularly for right-of-center consumers. In the web data, we find that horse race news is systematically slanted in a way that makes it more congenial to an outlet's typical reader, but also that (relatively) highly congenial news is usually not more likely to make an outlet's "most viewed" list, and sometimes less likely to do so. We draw two broad conclusions: 1) unsurprisingly, but in contrast to many economic models, consumers do not make news choices to maximize instrumental value; 2) the general preference for congenial news is not strictly driven by a psychological desire to avoid uncongenial information.

*The title is a reference to Gentzkow and Shapiro (2010). We thank Isaiah West for excellent research assistance.

†Marcel is a postdoctoral researcher at Hamburg Media School, Germany, and a research fellow at the University of Hamburg's Institute for Media Economics. He can be reached at: m.garz@hamburgmediaschool.com.

‡Gaurav can be reached at: gsood07@gmail.com.

§Dan is Assistant Professor of Economics at Bowdoin College. He can be reached at dstone@bowdoin.edu

¶Justin is a Ph.D. candidate at the University of Washington's Department of Economics. He can be reached at: jtw95@uw.edu.

1 Introduction

People tend to choose news that is congenial to their pre-existing political beliefs and preferences (Iyengar and Hahn, 2009; Gentzkow and Shapiro, 2011; Guess, 2016; Sood and Lelkes, 2016). The demand for congenial news raises concerns about echo chambers and filter bubbles (e.g., Pariser, 2011; Bakshy et al., 2015; Flaxman et al., 2016; Halberstam and Knight, 2016) and has been shown to have a large impact on the market of news. For instance, partisan slant of US newspapers is driven more by consumer preferences than those of newspaper owners (Gentzkow and Shapiro, 2010).

But what drives consumer demand for (congenial) media slant? Many papers assume a psychological explanation—loosely speaking, the desire to read news that "feels good" and/or avoid news that "feels bad" (e.g., Mullainathan and Shleifer, 2005; Bernhardt et al., 2008; Iyengar and Hahn, 2009). This behavior is often attributed to the preference to avoid cognitive dissonance.[1] However, a number of papers suggest that rational consumers choose congenial news because it provides greater instrumental value leading to more informed, and better, decisions (e.g., Burke, 2008; Chan and Suen, 2008; Oliveros and Várdy, 2015; Fang, 2016). A third possibility is that consumers perceive news to be more accurate when it confirms existing beliefs. Thus, congenial news can enhance the perceived credibility and reputation of a source (Gentzkow and Shapiro, 2006), although this perception could be mistaken (Vallone et al., 1985; Stone, 2011). The distinction between these conjectures—which we refer to, for short, as the psychology, instrumental, and reputation mechanisms—has important welfare implications. If consumers choose congenial news because it actually has greater instrumental value, then proliferation of diverse media has greater social benefits, and echo chambers concerns are more likely to be overly pessimistic (Gentkzow et al., 2015).

Although this distinction between instrumental and other causes of the demand for congenial news is elementary and important, the empirical literature addressing it is thin.[2] The

[1] Stroud (2011) discusses three causes of the demand for congenial news, a.k.a. partisan selective exposure, the term she uses, that fall in the category that we broadly refer to as "psychological": cognitive dissonance, motivated reasoning, and "cognitive misers" (that processing conflicting information takes more resources than consistent information). She also discusses the possibility that congenial information is perceived as higher quality, but does not discuss the distinction between whether this perception is correct or incorrect.

[2] The literature in economics on demand for congenial information more generally (beyond political media)

present paper addresses this gap in the literature. We examine the relationship between congeniality and demand for news for two related issues: assessments of performance in US presidential campaign debates and "horse-race" stories about the presidential candidates' chances of winning the election. Both of these issues have the advantage of being reported on repeatedly by various major news outlets, with substantial variation in congeniality of reported stories both across and within outlets.

Our key identifying assumption is that holding the issue and the outlet fixed, the actual and perceived information value of news on the issue is constant. Thus, if demand is higher for more congenial stories (on a given issue, from a given outlet), this greater demand can be attributed to the psychology mechanism. This assumption is especially plausible for outlets that are generally trusted by a given group of consumers, since then it is less likely that the perceived credibility of a story will vary across stories from the outlet. We can also draw some inferences from comparisons of congeniality and demand across outlets, and from variation in the supply of stories reported over time within outlets. There are, however, certainly other factors that can muddy the interpretation of each of these analyses, which we keep in mind and discuss throughout the paper.

We use survey experiments to analyze demand for debate stories, and (observational) web data, on both all stories reported and the most popular stories for a given outlet, for horse race stories. The survey data have the advantages of being micro-level, and giving us control over the timing and menu of news options offered. The web data have the advantage of reflecting real-world choices from the entire relevant population for the outlets studied.

The surveys were conducted on Amazon's Mechanical Turk (MTurk) in the morning following each of the first three presidential election debates in 2016. Respondents were offered a choice between timely (and real) New York Times (NYT) and Fox News stories on the debate "winner" and two other stories on issues unrelated to presidential politics, and incentivized to pick the story that they were most interested in reading, and hence most likely to read in a real-world environment.

is better developed. For instance, Karlsson et al. (2009) find that investors exhibit "ostrich effect," checking their portfolios more often when markets are rising rather than falling. See Golman et al. (2017) for a review of the literature.

The web data were obtained by scraping homepage and politics section headlines, and "most viewed" headlines from six outlets in 2012 (from web.archive.org) and 2016 (from "live" websites)—two with left-of-center reputations, two with right-of-center reputations, and two with reputations for relative neutrality in each year. The outlets for 2016 were: `nytimes.com`, `washingtonpost.com`, `news.google.com`, `news.yahoo.com`, `foxnews.com`, and `wsj.com`. In 2012, the Washington Post and Google News were unavailable on the web archive, so we replaced them with `huffingtonpost.com` and `usatoday.com`. We used keyword filters and MTurk coders to determine which headlines corresponded to horse race stories, and used MTurk coders to rate the "slant" (favorability of the story to one party or the other) of each headline. The resulting measures simply refer to partisan congeniality, and not necessarily a bias in comparison to some objective truth, though we will also consider some such comparisons. From the initial set of 491,228 from 2012 and 400,905 stories from 2016, we obtained final samples with 425 and 696 horse race stories for 2012 and 2016, respectively. We provide various types of validation of our slant measures, including Krippendorff alpha calculations, showing that the mean slants are highly correlated with daily poll results, and that the slant of each outlet in each year is predictive of daily polls, even conditional on the slants of news from the other outlets that day.

The survey analysis yields some results best explained by psychological forces, and some that could be explained by either the psychology or reputation mechanisms. Survey respondents on both sides of the political spectrum were indeed more likely to choose news about debates (NYT or Fox stories, as opposed to a Yahoo story) when this news was more congenial toward their preferred candidate, supporting the psychology theory. But the debate news demand for Democrats (and Clinton supporters) for news from the NYT was not significantly sensitive to the congeniality of news, and the overall variation in their demand was driven by variation in demand for Fox News. This result could plausibly be explained by the reputation mechanism in addition to the psychology mechanism, as Democrats are relatively unlikely to trust Fox in general, and may have perceived pro-Clinton stories by Fox as more credible than other Fox stories (i.e., the increase in demand for stories from Fox that were pro-Clinton was due to a story-level reputation effect).

For Republicans, congeniality of news was not significantly related to demand. However, a relatively low percentage of Republicans supported their party's candidate, Donald Trump (69.5% of Republicans, versus 82.4% of Democrats who supported Clinton in our sample). The demand of Trump supporters specifically for debate news, rather than Republicans more generally, was positively affected by the congeniality of the news. In contrast to Democrats and Clinton supporters, the congeniality effect for Trump supporters was stronger for their preferred outlet, Fox. It is less plausible that more pro-Trump stories were perceived as particularly credible Fox stories, and so this result provides relatively clean evidence in support of the psychology mechanism.

Moving to the web data, we find that the slants of horse race stories for the different outlets varied over the course of a campaign in similar ways. Both left- and right-wing outlets tended to report bumps for each candidate at the same time, consistent with Barberá and Sood (2014)'s finding that intra-media content is perhaps surprisingly diverse. However, the mean slants vary across the outlets substantially, in ways mostly consistent with their reputations. These results are consistent with those reported by Tremayne (2015) and Searles et al. (2016), but differ in substance from the findings of Budak et al. (2016), who conclude that the political slants of "news outlets are considerably more similar than generally believed." The variation we find in mean slant across outlets on an issue which, in theory, is an objective fact—which candidate is more likely to win at any given time, to what degree, and how the candidates' chances have changed recently—likely reflects psychologically driven demand for congenial news. However, this result could also be due to the reputation mechanism. It is possible that consumer priors favored the candidate the consumers preferred to win, and the outlets pandered to consumer priors to appear more credible.

Next, we analyzed the relationship between congeniality and the probability of a story making an outlet's list of most viewed stories. We consider this the paper's most significant contribution, and, frankly, expected a priori to find numerous positive relationships. Instead, we found only one—Fox News in 2012—and even that result was not robust. Moreover, there is more robust evidence of a *negative* relationship between congeniality and popularity for both Fox and the NYT in 2016. These results may have been due to the reputation mechanism (readers may have found these stories particularly credible due to their non-standard slants) or perhaps because the outlets became too extreme for their readers' tastes in 2016 for supply-side reasons. Fox readers may have preferred more right-slanted news in 2012, as compared to 2016, because a higher fraction of readers supported the Republican candidate in 2012. The bottom line, however, is that the relative popularity of a horse race story for a given outlet in our sample was mostly not driven by the story providing a stereotypical reader of that outlet with especially congenial news.[3]

Last, we analyze whether outlets reported more horse race stories on days in which the news was more congenial for those outlets' typical readers, and find some evidence indicating that Fox did this in both years, and that the NYT did this in 2016. We also present informal evidence that the NYT and Fox produced more debate news stories following debates in which their preferred candidate was considered to have "won" in 2016 but not in 2012. These results may provide additional support for the psychology mechanism, but may also have been due to the NYT and Fox becoming more partisan than their readers, so to speak, in 2016.

In summary, we obtain some results providing relatively clean support for the psychology mechanism, particularly for consumers who are right-of-center, and in the survey data. But the fact that detailed data is required to detect these results shows that their aggregate magnitudes are not that large. We also find that both left and right-of-center consumers get skewed news on the horse race on average, a result likely driven by psychology, reputation, and supply-side forces. There is accurate variation in the slants of these stories across time for all of the outlets we study, and stories uncongenial to each outlet's typical reader were usually at least as popular as congenial stories. This implies that consumers do, to a significant extent, wish to be accurately informed, and not just told what they want to hear, at least by trusted outlets (Miller and Krosnick, 2000). However, the robust variation in mean slants across outlets, among other results, implies consumers do not actually make choices to be

[3]If total website traffic was affected by the congeniality of horse race news at a given point in time, the relatively popularity of a story at that time could be misleading. We present some data indicating that this should not be a major concern for the interpretation of our results.

optimally informed.

The remainder of the paper is organized as follows. In Section 2, we discuss relevant theories of demand for congenial news and develop a theoretical framework to guide the empirical analysis. In Sections 3 and 4 we discuss the design of the surveys and the analysis of the survey data, respectively. Sections 5 and 6 discuss the web data and analysis of this data, respectively. We conclude in Section 7 with brief additional remarks.

2 Theoretical framework

More slanted news can, in theory, increase the instrumental value of news. That is, more slanted news can increase a consumer's expected utility from a choice the news provides information about. In general, models making this point (such as Chan and Suen, 2008, and others noted in Section 1) can be interpreted as examples of the broader literature on delegation and advice. Just as a decision-maker is better off being advised by, or delegating a decision to, an agent who shares the decision-maker's preferences, a news media consumer may be better off getting news from a media outlet with a similar political viewpoint (Gentkzow et al., 2015).

This idea is potentially relevant to the first context we study: choosing which, if any, outlet to get debate news from, since interpretation of debate performance could depend on a subjective viewpoint. For example, if the Democratic candidate focused on security in a debate and the Republican focused on growth, it would be reasonable for security-concerned outlets to assess the Democrat's performance more generously than other outlets, and for security-concerned citizens to get news from these outlets.

But since an outlet's viewpoint is likely stable, at least within the time period in which the debates occur in a given year, instrumental information seeking by consumers with different preferences is unlikely to explain within-outlet variation in the demand for debate news (across time, within that year). And since horse race news is about an objective, albeit unobserved, fact (which candidate is more likely to win the election), and not an issue that consumers with different political viewpoints could reasonably interpret in different ways, instrumental information seeking also very likely does not explain the choice of where to get horse race news in general. Optimal delegation seems even less likely to explain across time variation in demand for horse race news from the same outlet.

Models of slant and reputation for accuracy and credibility, however, may explain variation in demand across outlets. Consumers may demand congenial news because they perceive it to be more informative due to Bayesian inference, though this perception could be incorrect if based on an objectively incorrect prior. Gentzkow and Shapiro (2006) is the first paper to develop a formal model demonstrating this point. For example, media consumers who had an inflated belief about the chance of Clinton winning the 2016 election may have inferred that outlets that reported that Clinton was very likely to win were relatively credible outlets. These consumers in fact may have overestimated the credibility of such outlets.

It is also possible that the reputation mechanism explains a within-outlet relationship between slant and demand, for both debate and horse race stories, at least for outlets that are not trusted by a given group of consumers. If the perceived accuracy of articles varies within an outlet, and the slant of a headline affects an article's perceived accuracy, then this could make consumers more or less likely to click on slanted articles. We think it is useful here to draw a distinction between headlines whose slants are consistent and inconsistent with the outlet's typical slant. It seems reasonable to think that if an outlet is known for usually reporting news favoring party A, and then reports news favoring party B, the story must be especially compelling. Thus, the reputation mechanism could increase demand for "atypical" news—uncongenial news from a trusted outlet and congenial news from an uncongenial one. Another reason demand could be higher for atypical news is that it is more surprising (Ely et al., 2015). In the typical case (congenial news from a trusted outlet), consumers likely do not perceive the news as particularly accurate and credible. Therefore, greater demand for this kind of news must be due to psychological factors.[4]

[4]Chan and Suen (2008), and other later papers, develop a similar logic for why congenial news would increase the instrumental value of news. However, the application here is a bit different. In our context, readers draw inferences about the credibility of a particular story, conditional on outlet, from its headline. In Chan and Suen (2008) and related models, readers know the equilibrium strategy of each outlet. Some outlets' strategies result in congenial news on average, implying that when these outlets do report uncongenial news, this is especially precise and informative.

In addition to slants varying across and within outlets, the relationship between slant and the quantity of stories produced could shed light on the mechanism underlying demand for congenial news. While it is very plausible that psychological factors could lead to demand for a greater quantity of congenial stories, it seems unlikely that either of the other mechanisms would lead to this outcome. Even if more or less congenial stories from a given outlet are considered more credible, this would likely not make readers want to read multiple stories making the same point for the sake of being better informed. But if the stories are being read because they "feel good," then we can see why readers would want to read more stories when they are more congenial.

We summarize this discussion in Table 1.[5] There are three main take-away points. First, instrumental information seeking can drive variation in demand across outlets for debate news but not for horse race news, and cannot drive within-outlet variation in demand for either issue. Second, increased demand for stories with more congenial slant, from a given outlet (and when this congenial slant is not atypical for the outlet), is uniquely explained by psychological factors, as is an increased supply of more congenial stories from a given outlet. Third, increased demand for stories with atypical slant (that is not congenial) is uniquely explained by the reputation mechanism. These predictions are, of course, simplified, and not as unambiguous as the table implies. A particular issue worth noting is that we are completely abstracting from supply-side factors. Still, the table provides a useful distillation of testable predictions of the theories.

3 Survey-experiment data

3.1 Description

In the morning (between 9:00 AM and 10:00 AM) following each of the first three (of four) 2016 US presidential election debates, we conducted a survey on MTurk.[6] The first and third

[5]To be clear, there are other causes of news demand, such as the innate desire to be informed or the feeling that this is one's social duty, but these are less likely to cause demand for congenially slanted news.

[6]MTurk is a market for trading small services. "Workers" are paid for finishing small tasks like filling out surveys, tagging an image, etc. Research suggests that participants recruited on MTurk are "slightly more demographically diverse than are standard Internet samples and are significantly more diverse than typical

Table 1: Summary of theoretical predictions

Congenial news demand mechanism:	Instrumental information seeking	Reputation and credibility	Psychological value
Debate news			
1. ↑ demand for more congenial outlets?	✓	✓	✓
2. ↑ demand for more congenial stories, within an outlet?			✓
3. ↑ demand for stories with atypical congeniality, within an outlet?		✓	
4. ↑ demand for greater # congenial stories, within an outlet?			✓
Horse race news			
1. ↑ demand for more congenial outlets?		✓	✓
2. ↑ demand for more congenial stories, within an outlet?			✓
3. ↑ demand for stories with atypical congeniality, within an outlet?		✓	
4. ↑ demand for greater # congenial stories, within an outlet?			✓

debates were between the presidential nominees, and the second one was between the vice president candidates. In each survey, we asked a small number of demographic and party affiliation questions, and asked respondents to pick the article they were most interested in reading from four options—two articles on which candidate won the debate, one from the NYT and one from Fox, and two articles on other topics from news.yahoo.com. We use articles from the NYT and Fox because they are leading outlets with reputations and readers skewing left- and right-of-center, respectively (e.g., Pew, 2014). Respondents were also told that after choosing the article, they would be asked a question on the article's content. We noted that the difficulty of the question would be the same regardless of what article they picked, and that they would receive an additional payment if they answered the question correctly (respondents

American college samples" (Buhrmester et al., 2011). See Appendix Table A1 for sample statistics for key variables. Compared to the population, survey respondents recruited on MTurk tend to be younger, better educated, and more likely to identify with the Democratic party (Berinsky et al., 2012), though sectoral breakdown of employment is similar to more representative online surveys—the sectoral differences are no more than 7% (Huff and Tingley, 2015). Pleasingly, a broad variety of experiments done on MTurk tend to reach similar conclusions as those done on more representative samples (e.g., Mullinix et al., 2015). The two major advantages of MTurk over a survey firm are: a) MTurk is much more cost effective, allowing us to obtain a larger sample, 2) MTurk gives us control over the timing of surveys, which, as we explain in the next section, is crucial for their validity.

were paid $0.50 per survey and a bonus of $0.25 for answering the reading question correctly). The exact wording of the instructions can be found in the supplementary appendix (see Table A2).

Since respondents were presented with a choice of articles that were both timely and real, and given an incentive to actually read the article they selected, respondent choices should reflect the articles that they would be most likely to read in similar real-world situations. The most unnatural element of our design was that respondents were presented with news choices from both congenial and uncongenial outlets. For example, Democratic respondents saw both NYT and Fox headlines. This would be unrealistic for news consumers who are rarely exposed to diverse outlets (e.g., people who go straight to the NYT website to decide which story to read, or have a Facebook feed consisting of ideologically similar sources). We consider this issue in the interpretation of our results. Another issue was priming: In the first two surveys, we asked which candidate the respondents were planning to vote for last, after choosing and reading their article. We chose this order to avoid priming respondents about the election, which could influence article selection in an unnatural way. In survey 3, we randomly asked some respondents about their preferred candidate before presenting the article options to assess if the order mattered. However, (unreported) tests did not indicate significant effects on article choice.

It is worth noting that although our survey did not explicitly randomize congeniality across groups of respondents labeled treatment and control, we essentially did the same thing by varying the congeniality of news across surveys to different types of partisan respondents, while partisan congeniality remained relatively constant (neutral) for non-partisan respondents. These respondents thus constituted a control group, which is why our approach is still appropriately described as an experiment.

We ran the surveys on debate news, and not news on other topics (such as the horse race), because the timing of these stories was known in advance. This allowed us to prepare in advance to select the stories, and post the surveys shortly after the stories became available. Debate news stories also worked well for our surveys because we expected the congeniality of these articles to vary over time and across outlets in an unambiguous way—most outlets usually proclaim one candidate the winner of each debate. Indeed, both the Fox and NYT headlines that we used in the first survey stated that Clinton won the first debate, and both said that Pence won the second. The outlets only disagreed on the third debate; Fox's headline said that Trump won, while the NYT's headline did not declare a winner. All the article headlines are provided in Table A2.

We surveyed 250 MTurkers (US residents aged 18 or older) after each debate. We discarded observations in which the respondent did not answer the reading comprehension question correctly. We also discarded 16 observations in which the respondent chose "other" in response to the partisan self-identification question as we suspected many of these respondents were partisan, potentially influencing their article choices, but did not want to reveal this. Our final sample had 638 observations, with 346 identifying as Democrats or leaning toward Democrats, 177 identifying as Republican or leaning toward Republicans, and 115 as independent.

Figures 1 and 2 foreshadow the econometric results from the next section. Across the three surveys, Democrats were most likely, and Republicans least likely, to pick an "other [non-debate] news" story after the first debate, which was the only debate that both the NYT and Fox headlines said Clinton won. However, Democrats were slightly more likely to choose the NYT story after the second and third debates, as compared to the first debate. Democrats were significantly less likely to choose the Fox News story after the second and third debates compared to the first debate. Republicans' demand for NYT news was fairly stable across the debates, while their demand for Fox news grew as the congeniality (and congeniality relative to NYT) grew. Independents were much less likely than partisans to get political news in general, but became more likely to get political news in later debates, suggesting a secular increase in interest in political news. Results are somewhat sharper when respondents are split out by preferred candidate rather than party.

3.2 Analysis

We focus on the debate news predictions of Table 1 for which the theories vary, and which we have the data to address: #2 (within outlet variation in congeniality increases demand

Figure 1: News choices by survey and party

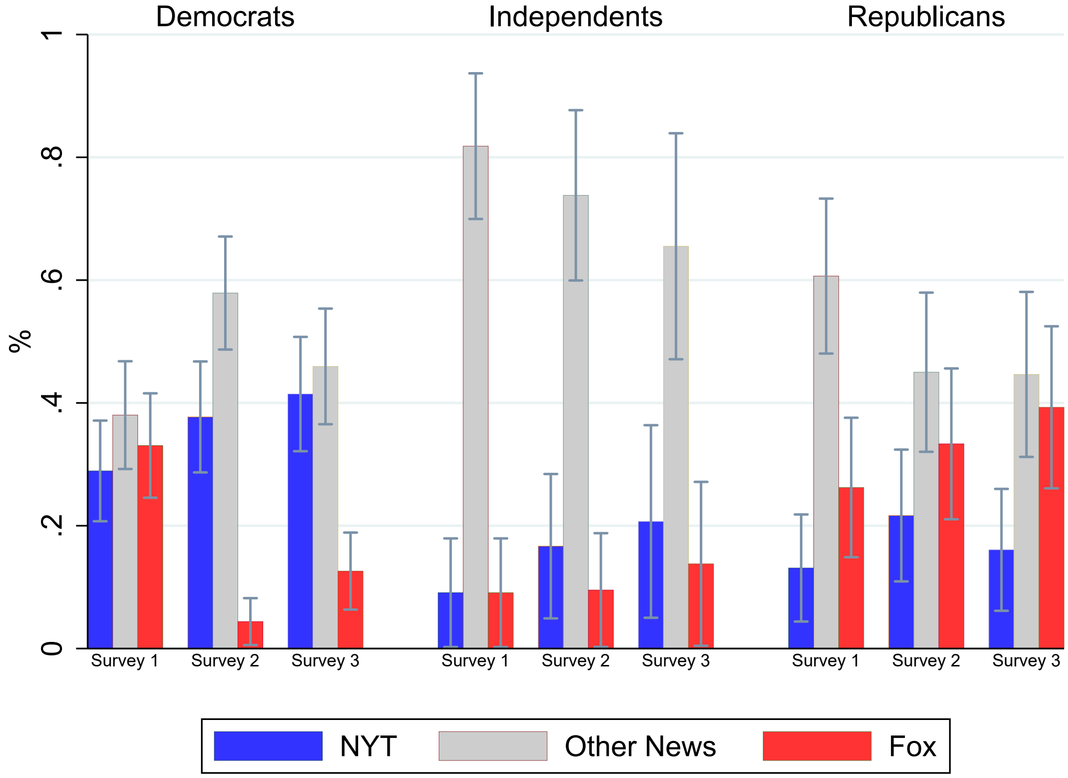

Note: 1) Both NYT and Fox survey 1 headlines said Clinton (Democrat) won first debate, 2) Both NYT and Fox survey 2 headlines said Pence (Republican, Trump's VP) won the second debate, and 3) Fox survey 3 headline said Trump won third debate while NYT survey 3 headline was ambiguous.

for psychological reasons) and #3 (atypical congeniality increases demand due to the reputation mechanism). At the end of the section, we present and discuss a limited amount of data relevant to prediction #4 (that outlets supplied more stories when the news was more congenial).

To analyze survey respondent choice, since respondents chose among four unordered alternatives, a multinomial model is ideal for analysis. However, in the interest of simplicity and transparency, we relegate the multinomial analysis to the appendix (see Table A3), and present here results of linear probability models predicting a binary outcome equal to: 1) Y^{NYT} (= 1 if the respondent chose the NYT article); 2) Y^{Fox} (defined analogously); 3) whether either type of debate news story is chosen ($Y^{NYT} + Y^{Fox}$). We run two sets of regressions, one using party identity as a measure of the respondent's politics and one using the respondent's

Figure 2: News choices by preferred candidate and survey

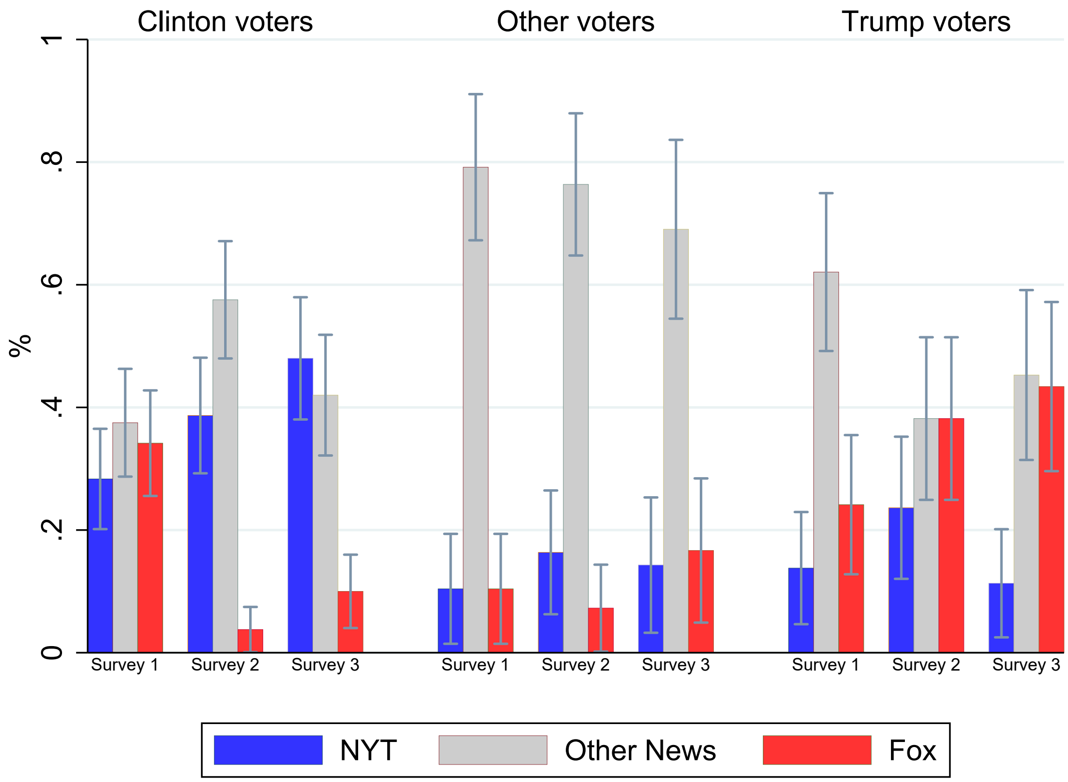

Note: 1) Both NYT and Fox survey 1 headlines said Clinton (Democrat) won first debate, 2) Both NYT and Fox survey 2 headlines said Pence (Republican, Trump's VP) won the second debate, and 3) Fox survey 3 headline said Trump won third debate while NYT survey 3 headline was ambiguous.

preferred candidate for this, each with the following structure:

$$Y_i = \beta_L D_i^L + \beta_R D_i^R + \beta_{S^2} S_i^2 + \beta_{S^3} S_i^3 + \beta_{L,2} D_i^L S_i^2 + \beta_{L,3} D_i^L S_i^3 + \beta_{R,2} D_i^R S_i^2 + + \beta_{R,3} D_i^R S_i^3 + \beta_X X_i. \tag{1}$$

D_i^L is a dummy for respondent i being "type L" (a Democrat in one set of regressions, or Clinton-voter in the other set), D_i^R is analogous (independents are the omitted group); S_i^t is a dummy for respondent i taking survey t (the survey 1 dummy is omitted); X_i is a vector of demographic and other controls. Some respondents took more than one survey, but including fixed effects for these respondents has very little effect on the results.

The parameters of interest are the politics-survey interactions: $\beta_{L,2}$ and $\beta_{L,3}$ can be interpreted as mean changes in type L demand for surveys 2 and 3, respectively, as compared to survey 1, and $\beta_{R,2}$ and $\beta_{R,3}$, have analogous interpretations for type R respondents. Recall that the NYT's first headline was most (least) congenial, and the second headline least (most) congenial, to L (R) types. Fox's first headline was most (least) congenial to L (R) types. The psychology mechanism thus implies that for Y^{NYT}, $\beta_{L,3} < \beta_{L,2} < 0$. For Y^{Fox}, $\beta_{L,3} < 0$ and $\beta_{L,2} < 0$, but $\beta_{L,3}$ and $\beta_{L,2}$ cannot be compared since the headlines for those surveys are both uncongenial. For Republicans and Trump supporters, the signs of these predictions are reversed. The reputation mechanism implies that for Y^{NYT}, $\beta_{t,2} > \beta_{t,3} > 0$, and for Y^{Fox}, $\beta_{t,3} < 0$ and $\beta_{t,2} < 0$ (these predictions are the same for respondents of type i = L or R). Note that the effects of the psychology and reputation mechanisms go in opposite directions for Y^{NYT}, and the same direction for Y^{Fox} for L types, and vice versa for R types.

Table 2 reports the results. Congeniality drives Democrat and Clinton supporter demand for Fox news, but not for NYT news. Democrats are nearly 30 percentage points less likely to get Fox news when it is uncongenial than when it is congenial ($\beta_{L,2} < 0$). Congeniality also drives both Democrats and Clinton supporters demand for debate news from either source. There are no significant survey effects for Republicans. For Trump supporters, there is also evidence of a congeniality effect but it is primarily for Fox: Trump supporters are around 20 percentage points more likely to get Fox news in survey 2 as compared to survey 1 ($\beta_{R,2} > 0$). It is hard to interpret the external validity of the magnitudes of these effects, however, given the limited set of options the respondents had available.

For left-of-center respondents (L types), these results are consistent with off-setting reputation and psychology mechanisms for NYT demand, or neither mechanism. For those right-of-center (R types, and Trump supporters specifically), the results are only consistent with the psychology mechanism for Fox demand (the results are not consistent with either type of effect for NYT demand). Thus, the results support the existence of psychology and/or reputation mechanisms for L types, and of the psychology mechanism for R types. The magnitudes are substantial (news demand can approximately double due to changes in these effects) but not completely dominant (these effects cannot be the only explanation for observed news demand).

Table 2: Analysis of survey data

	Y^{NYT}	Y^{Fox}	$Y^{NYT} + Y^{Fox}$
Panel A: Survey-party interactions			
Survey 2 × Democrat	0.010	-0.292***	-0.282**
	(0.097)	(0.078)	(0.111)
Survey 3 × Democrat	-0.017	-0.216**	-0.233*
	(0.110)	(0.100)	(0.127)
Survey 2 × Republican	0.007	0.076	0.083
	(0.103)	(0.102)	(0.127)
Survey 3 × Republican	-0.098	0.134	0.036
	(0.113)	(0.119)	(0.139)
Adj R^2	0.062	0.093	0.086
N	637	637	637
Panel B: Survey-preferred candidate interactions			
Survey 2 × Clinton Supporter	0.040	-0.275***	-0.235**
	(0.094)	(0.075)	(0.106)
Survey 3 × Clinton Supporter	0.139	-0.276***	-0.137
	(0.097)	(0.095)	(0.117)
Survey 2 × Trump Supporter	0.056	0.194*	0.250**
	(0.101)	(0.103)	(0.122)
Survey 3 × Trump Supporter	-0.052	0.181	0.129
	(0.099)	(0.117)	(0.132)
Adj R^2	0.080	0.121	0.090
N	637	637	637

Note: All models are estimated using OLS with robust standard errors and include survey, education, gender, age, and party identity (Democrat, lean Democrat, Republican, lean Republican, independent) fixed effects. *, **, *** denote 10%, 5%, 1% significance.

To assess prediction #4, we manually collected the number of debate-related links (both articles and videos) from web.archive on both the Fox and NYT websites in the morning following each of the four debates of both 2012 and 2016. These numbers are presented in Table 3. The table shows that the number of stories was fairly constant for both outlets in 2012, but this was not the case in 2016. Fox had the most links after the third debate that year, and the fewest links after the fourth debate. The NYT had the fewest links after the second debate. Recall that Fox claimed Trump won the third debate, and the NYT said the Republican (Pence) won the second. Fox's headlines favored Clinton after the fourth debate (e.g., "Trump winning on points until terrible mistake"). Thus, both Fox's and the NYT's numbers of debate links are correlated with the congeniality of the debate outcome for their readers, and so the table provides further support for the psychology mechanism. The table also raises the question of how the slant of the outlets may have changed over time.

Table 3: Counts of links to reports on presidential debates the morning following each debate

	2012				2016			
Debate:	#1	#2 (VP)	#3	#4	#1	#2 (VP)	#3	#4
Fox	12	7	12	13	24	25	31	17
NYT	20	19	21	19	23	14	23	25

Note: The counts are based on web.archive.org snapshots of nytimes.com and foxnews.com at approximately 10:00 AM the morning following each debate.

4 Web data

4.1 Description

The survey analysis is intended to complement our analysis of real-world demand for horse race stories—stories about which candidate, if any, is winning the "race" at any given time, and by how much. For the 2016 presidential election, we started scraping news articles on July 27, 2016, well after the presumptive nominee for each party had been decided. The websites of *Fox News* (Fox), *Google News* (Google), *New York Times* (NYT), *Washington Post* (WashPost), *Wall Street Journal* (WSJ), and *Yahoo! News* (Yahoo) were scraped three times daily until election day (November 8, 2016). We downloaded outlets' landing pages, politics sections, and most viewed lists. We chose these outlets because of their prominence, ideological diversity, and because each of them publicly reports "top," "trending," "most popular," or "most viewed" stories, which we henceforth refer to as simply "most viewed". The selection of Google's "top stories" is based on an algorithm, whereas the other five websites each use terms that explicitly or implicitly refer to stories being most frequently clicked on within a recent period of time. Given the inclusion of Fox News, it would be natural to include the two other major cable news outlets, CNN and MSNBC. However, neither of these outlets report most viewed stories.

News data for the 2012 presidential election were collected by scraping snapshots of the outlets' homepages stored by web.archive.org, also dating between July 27 and election day of that year. Since web.archive.org respects `robots.txt` advisory files, it did not collect data from Google and WashPost, and partly for Yahoo, which mandate `no follow`. We substituted for these losses with *USA Today* (USAT) and the *Huffington Post* (HuffPost) for 2012.[7] We also downloaded data from snapshots in 2016 to complement the live data, and collected some of the archived data manually (rather than by scraping), when necessary.

Both the archive and the live data include article date, time, URL, source, headline, text, author, and keywords, and, where applicable, the current rank in the most viewed list. Often, the outlets keep an article long enough to survive more than one visit of our web scraper (or more than one snapshot of web.archive.org). And sometimes the same article is posted multiple times, e.g., after updating or moving it to another section. Thus, there are many cases in which we download the same article multiple times.

Next, we processed the text data so that it could be used for statistical analysis. To do this, we first had to identify a set of articles that could potentially be horse race stories. We used a fairly broad set of keywords to make the initial set inclusive and then narrowed this set down with human coding and additional restrictions. The initial set included any article with one of the following keyword combinations: at least one of the terms from the set {Obama (Clinton in 2016), Romney (Trump in 2016), president, white house, electoral} in the headline or URL and at least one term from {win, winning, momentum, lead (and not "leader"), bounce, bump, tied, gallup} in the headline. These keywords were determined to lead to a very small fraction of false negative classifications (actual stories stories that were not classified as horse race stories) at the expense of having a high number of false positives. We erred in this direction because the cost of cutting false positives with additional steps was relatively low. We then dropped articles that included a clear indicator of being an opinion piece in the headline, due to our focus on hard news stories.[8]

We then had three "master" MTurkers rate every headline on a five point scale—very good news or good news for either the Democratic or Republican candidate's chances of winning, or

[7] We scraped the text for USAT and the HuffPost for 2016 as well but exclude them from the analysis to maintain consistency across years (i.e., to maintain an estimation sample with two outlets considered left-of-center, two considered right-of-center, and two relatively neutral). Moreover, USAT stopped reporting most viewed articles relatively early in the 2016 campaign.

[8] We used the following headline keywords, determine by inspection, to identify articles as opinion pieces: {opinion, schoen, goodwin, rove, strassel, power play, juan williams, bias alert, gainor, reich, douthat, dan rather, whalen, starnes}.

neutral— with two additional options, ambiguous or unclear, or not relevant to the candidates' chances of winning the election. The exact instructions provided to MTurkers are in Table A4. We just used the headlines for this purpose, and not the article text, both for simplicity and because the headline is likely the key factor in determining whether or not an article lands on the most viewed list. We did not include the names of the outlets to not prime the coders when rating the headlines. There was a trade-off between incentivizing effort and accuracy on the one hand, and excessive monitoring and potential "demand effects" on the other. Thus, we kept instructions intentionally vague, offering the incentive of generously paid additional work (workers were paid $3 for each batch of 40 headlines). We monitored the coding done by these MTurkers by choosing four headlines that were relatively clear, and spot-checked each worker's ratings for these headlines, for each batch of 40 headlines. These spot-checked ratings were consistent with our expectations in all but one batch (out of dozens of batches); we manually checked the other headlines in that batch and they seemed reasonable so we continued to invite the worker to do additional work. We used the same three MTurkers for the vast majority of the coding for consistency and because presumably, as they gained experience, their ratings likely became more accurate, but we used three other MTurkers for a small number of batches (we required additional manpower at that time due to the presence of a deadline).

There were 2,025 headlines coded in total. The Krippendorff's α, a standard measure of inter-coder reliability, for all the coded headlines is 0.313. Condensing to an ordinal three-point scale—good news for the Democrats, good news for the Republicans, or neutral or ambiguous—increases α to 0.816, exceeding the standard threshold of 0.80.[9] Of the 1,177 headlines that were not rated as irrelevant by any of the coders, the α for the five-point scale is 0.405 and for the three-point scale, it is 0.859. Of the 871 headlines that were not rated as irrelevant or ambiguous by any of the coders, the α for the five-point scale is 0.440 and for the three-point scale, it is 0.900. Thus, the three-point scale appears much more valid than the five-point scale, and so we only use the three-point scale going forward, with very good or

[9] For the three-point scale, we used the ordinal method to calculate the Krippendorff α and thus coded "not relevant" as missing values. For the five-point scale, Krippendorff α values were very similar whether we used ordinal or nominal methods.

good news for Democrats coded as -1, very good or good news for Republicans coded as 1, and neutral and ambiguous news coded as 0. Restricting the sample to headlines that all three raters agreed were relevant and unambiguous increases the validity, but results in a substantial loss of observations. Thus, for transparency and to examine robustness, we consider three variations of the measure for all of our analyses: $Slant_1$ = mean slant of headlines rated as irrelevant or ambiguous by at most one coder; $Slant_2$ = mean slant of headlines not rated as irrelevant by any coder; $Slant_3$ = mean slant of headlines not rated as irrelevant or ambiguous by any coder. Headlines used for $Slant_3$ are a strict subset of those used for $Slant_1$, and this is almost the case for $Slant_2$ as well, so there are some sub-groups of headlines with more $Slant_2$ ratings than $Slant_1$ ratings.

Figures 3 and 4 present smoothed plots of daily means of each slant measure versus daily poll averages (we use R's "pollstR" library to obtain poll data from HuffPost Pollster). The plots are quite similar to one another, and to the polling average. Furthermore, Tables A5 and A6 show that the slant of each outlet's stories is predictive of that day's polls, for both years. Most of the outlets' slants are also significant predictors of polls even controlling for a date polynomial, and some slants are also significant predictors of changes in polls (sometimes with date polynomials). These results and figures provide further support for the validity of each slant measure. However, a number of results from the analyses (reported in the next section) are significant for one measure of slant but not for others, and so it could be misleading to restrict the analysis to just one of these measures. To illustrate the coding, Table A7 shows the three slant ratings for all headlines that contained the appropriate keywords for the day before the 2016 election (November 7, 2016). Most of the ratings seem very reasonable; sometimes $Slant_2$ and/or $Slant_3$ seem to appropriately drop a non-horse race headline (e.g., "Trump urges voters to deliver justice at polls"); sometimes these more restrictive versions seem to mistakenly drop a horse race headline (e.g., "polls Trump and Clinton virtually tied in key swing states"). Thus, the table also supports the use of the various *Slant* measures in the analysis.

Table 4 reports the number of unique most viewed and other articles per outlet, and

Figure 3: Mean $Slant_i$ and poll average versus day relative to election day (day 0) in 2012.

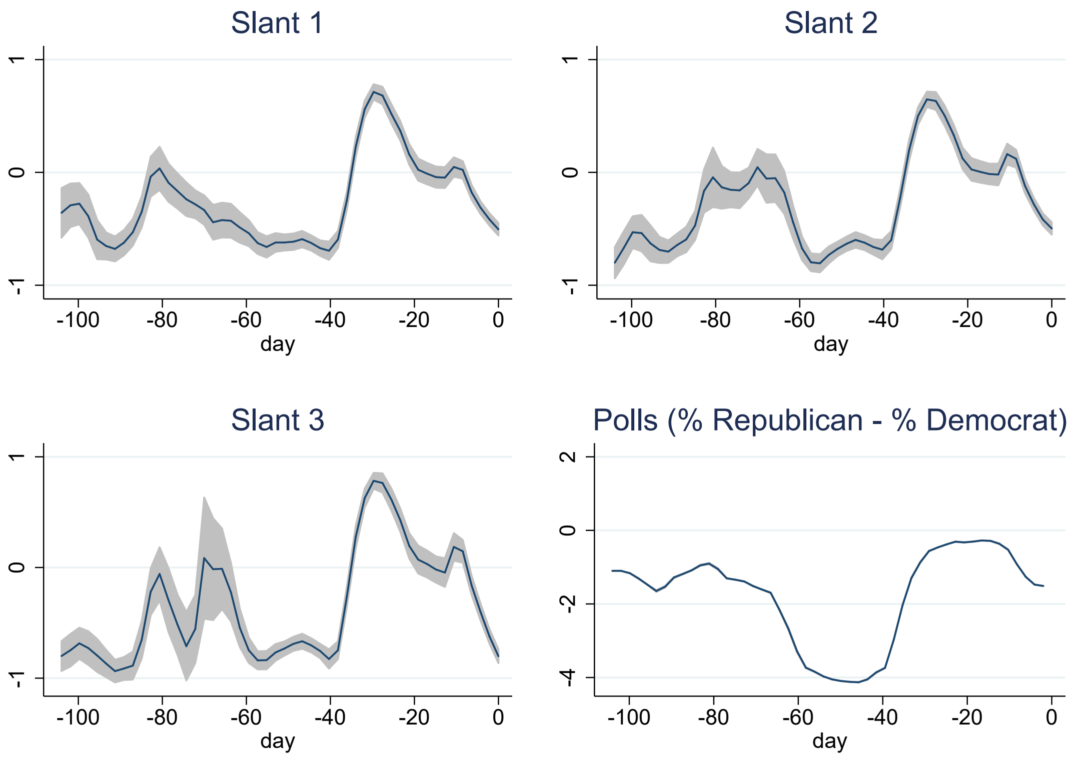

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands.

their means for various slant measures.[10] The average slants vary across outlets substantially, largely consistently with the outlets' reputations. Panel A of the table, on the 2012 data, also suggests that the most viewed stories were slanted to the left of other stories for all outlets except for Fox. Panel B (on the 2016 data) suggests that both Fox's and the NYT's most viewed stories were more neutral than the outlet's other stories.

To further examine the data before proceeding to the formal analysis, Figures 5 and 6 present smoothed polynomials with 95% confidence bands of the relationship of $Slant_1$ and

[10]The number of headlines reported in this table is less than the corresponding number referred to in the Krippendorff alpha analysis (2,025) because the sample used in this table, and for most of the subsequent analysis, for two reasons. First, even the broadest slant definition that we use for the main analysis, $Slant_1$, is restricted to headlines coded as irrelevant or ambiguous by at most one coder, and therefore excludes many of the original 2,025 headlines. Second, the story-level data set collapses headlines with slight variants in wording to a unique observation, whereas the MTurkers coded multiple variants of headlines, with wording that slightly differed, for some stories (such as "FOX NEWS POLL Clinton leads Trump by 10 points both seen as flawed presidential candidates" and "Fox News Poll Clinton Leads Trump by 10 Pts Yet Both Flawed Say Voters"). Including these variants in the α calculations should not bias results up or down since the coders are as likely to disagree on variants of headlines for a given story as they are on a single version of a headline.

Figure 4: Mean $Slant_i$ and poll average versus day relative to election day (day 0) in 2016.

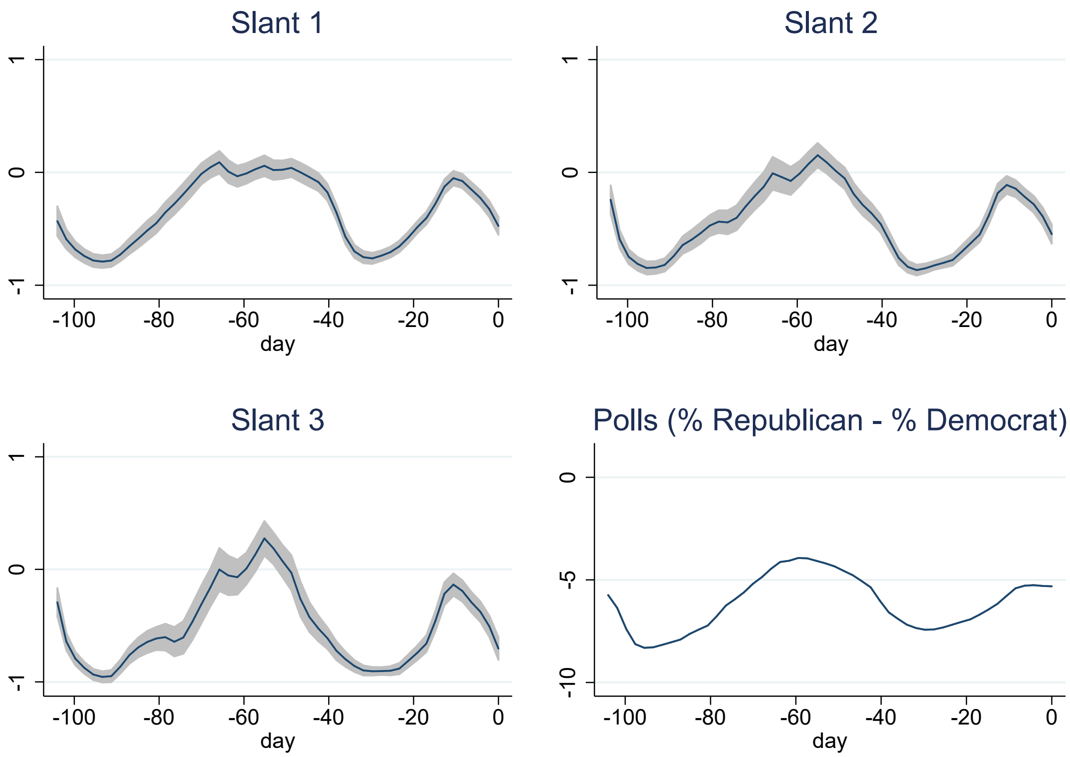

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands.

days to the election for stories that made the most viewed list that day, and for all other stories, for each outlet and year. We restrict these figures to just $Slant_1$ for two reasons: 1) presenting these figures for the other slant measures would be unwieldy and 2) the figures are not formal statistical tests. The figures use scraped headline, and not story, level data, i.e., a data set with a separate observation for each headline scraped for a given story, because here we want to include separate observations for the same story made available on different days. The confidence bands reflect variation in $Slant_1$ within an outlet within days and across days that are close together, but the bands can be misleading as they can be small (or non-existent) due to limited data. Thus, the figures should be interpreted as merely illustrative. Still, they are useful for seeing broad trends in the data, and availability of the data.

For 2012, the four outlets for which we have the most data—Fox, NYT, Yahoo, and HuffPost—all show pro-Republican bumps for "other" stories around 80 and 30 days prior to the election, consistent with Figure 3. The most viewed stories show the latter bump as well,

Table 4: Mean slants and article counts by outlet

Outlet	Type	$Slant_1$	N	$Slant_2$	N	$Slant_3$	N
Panel A: 2012							
Fox	Other	0.27	38	0.33	31	0.38	23
	Most viewed	0.37	25	0.34	27	0.42	22
WSJ	Other	-0.37	6	-0.37	6	-0.28	6
	Most viewed	-0.53	7	-0.52	7	-0.53	7
USAToday	Other	0.10	13	0.00	9	0.00	8
	Most viewed	-0.31	19	-0.24	20	-0.29	16
Yahoo	Other	-0.12	76	-0.23	50	-0.31	31
	Most viewed	-0.23	30	-0.21	32	-0.36	22
NYT	Other	-0.17	59	-0.15	58	-0.21	42
	Most viewed	-0.48	21	-0.31	24	-0.48	14
HuffPost	Other	-0.58	119	-0.50	102	-0.59	76
	Most viewed	-0.56	12	-0.37	16	-0.50	10
Panel B: 2016							
Fox	Other	0.35	23	0.50	12	0.67	8
	Most viewed	0.00	47	-0.11	33	-0.16	25
WSJ	Other	-0.65	12	-0.90	7	-1.00	6
	Most viewed	-0.67	11	-0.93	9	-0.93	9
Google	Other	-0.23	270	-0.28	205	-0.33	161
	Most viewed	-0.16	46	-0.36	34	-0.42	26
Yahoo	Other	-0.27	85	-0.42	60	-0.68	39
	Most viewed	-0.50	8	-0.33	6	-0.33	6
NYT	Other	-0.65	33	-0.79	24	-0.90	21
	Most viewed	-0.37	22	-0.35	19	-0.61	11
WashPost	Other	-0.38	96	-0.46	60	-0.61	45
	Most viewed	-0.36	43	-0.40	32	-0.51	24

while only Fox's most viewed stories reflects the earlier bump. The "most viewed" stories for HuffPost and WSJ all come from the last 30–40 days of the sample time-frame. For 2016, there is evidence of more general biases in reporting: Neither the NYT nor WashPost show the pro-Republican bumps seen in the Google and Fox most viewed stories at around 80 days prior to election day, and NYT "other" stories do not show the pro-Republican bump seen for most other outlets and around 10 days prior to the election. The NYT most viewed graph also reveals a lack of data in general and in particular, early in our time-frame; this is because their most viewed list often included a link with variants on simple headlines along the lines of "2016 election polls". These headlines were not classified as horse race stories since they did not refer to the presidential race in particular and would almost surely be coded as neutral

Figure 5: Mean $Slant_1$ versus day relative to election day (day 0) by outlet in 2012

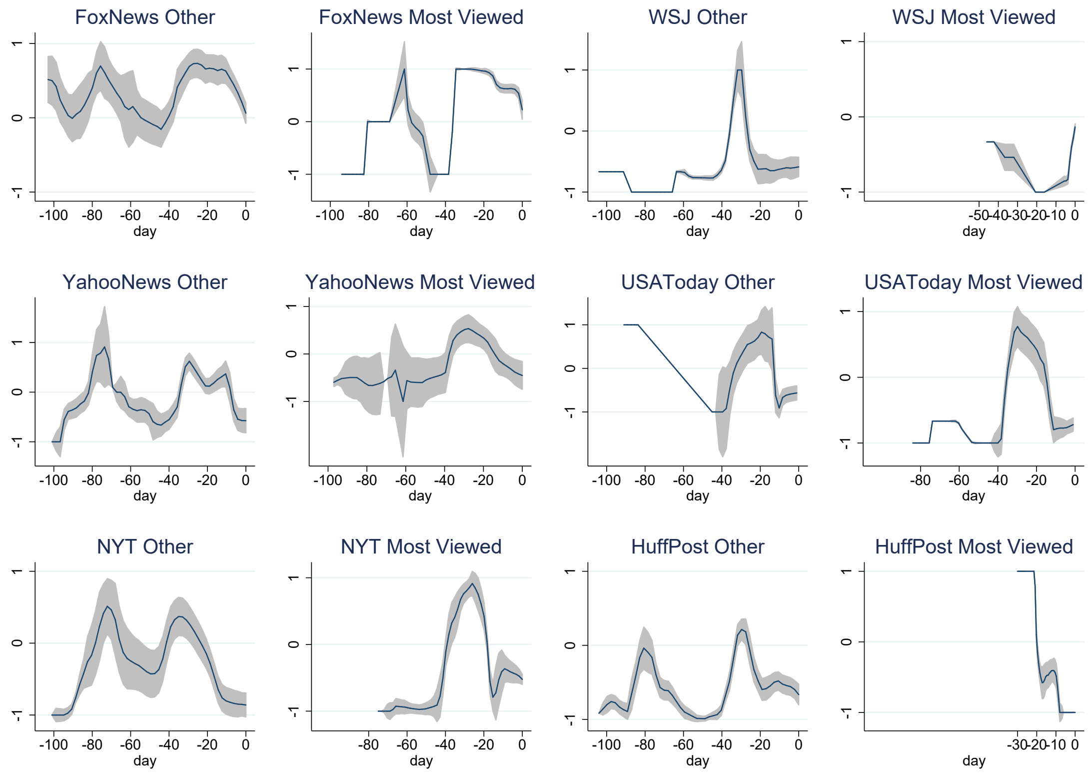

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands.

regardless.

4.2 Analysis

We first address prediction #1 of the horse race section of Table 1: that outlets may increase their supply of congenial news to satisfy consumer demand, either due to the reputation or the psychology mechanism. We do this by analyzing how the mean slant of stories varied across outlets. Table 5 presents results for simple regressions of $Slant_i$ (for each i) on outlet fixed effects, using the story-level data set. Here, we use Yahoo as the reference category, given the outlet's relative neutrality. The models also include day fixed effects to account for general trends affecting horse race story demand and supply—the content of these stories, such as the results of recent poll changes; the importance of those stories, which may be influenced by the timing relative to the election; and the importance of competing stories to fill the "news hole." The outlet coefficients thus represent the mean difference in slant for that outlet as

Figure 6: Mean $Slant_1$ versus day relative to election day (day 0) by outlet in 2016

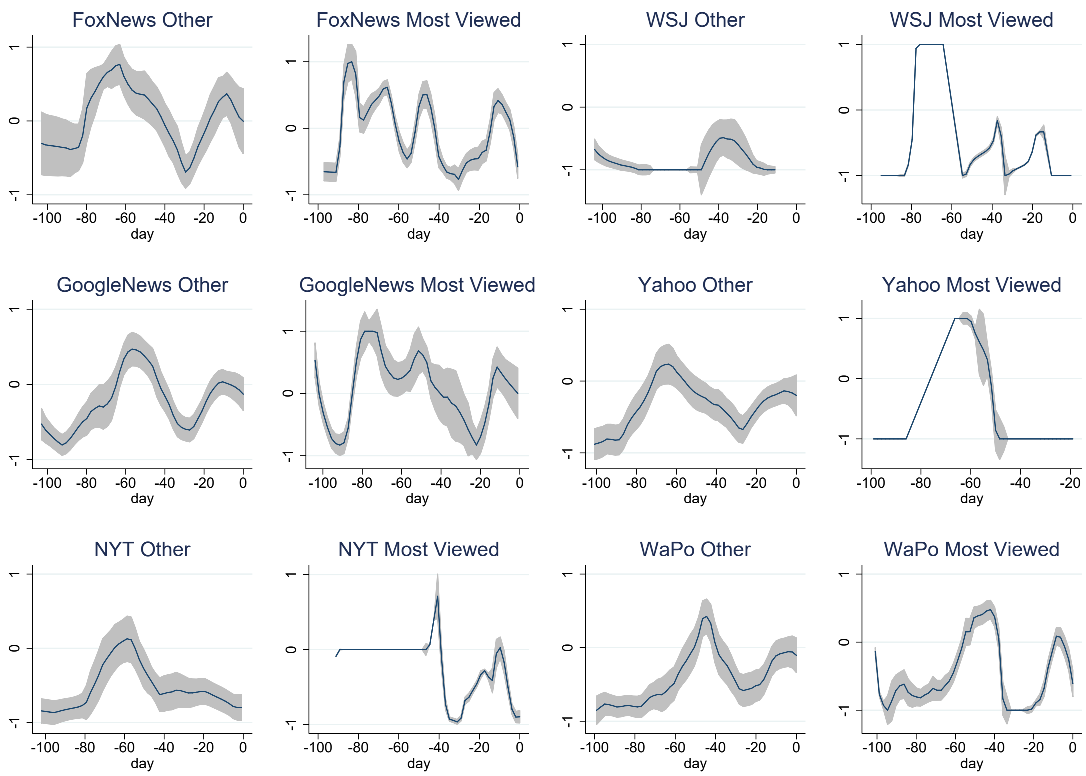

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands.

compared to Yahoo given the mean slant of stories reported on that day across outlets. In a second version of each regression we include the difference in polls for the Republican and Democratic candidate for the day the story was first reported and the average change in this difference over the previous week as additional controls. Results are very similar with various alternative controls, such as a dummy for whether or not the story made the most popular list, or interactions of this dummy with day fixed effects (not reported here).

As expected, in both years, Fox's stories were slanted significantly to the right of the other outlets. The leftmost stories on average in 2012 were from the HuffPost. In 2016, the NYT was significantly to the left of Yahoo at the 5% level for two specifications, and WSJ was significantly to the left at 10% for one specification, while Google was significantly to the right at 5% for both $Slant_3$ models. However, these Google results (and lack of significant NYT results for these models) seem to be due to Yahoo's slant skewing to the left for these specifications, as all of the other outlets' estimated slants are more conservative as compared to the other specifications.[11] The magnitudes of the effects are large: On a day on which Yahoo's expected $Slant_3$ is 0 in 2012, the expected $Slant_3$ for a Fox story is expected to be 0.795, i.e., there is a 79.5% chance the headline would be about good news for Republicans. In 2016, the expected value of $Slant_3$ for Fox for a neutral Yahoo story is 0.685.

Next, to address horse race predictions #2 and #3 from Table 1, we estimate the effect of within-outlet variation in slant on the probability of making the outlet's most viewed list. We do this by using linear probability models with day fixed effects, outlet fixed effects, and outlet-$Slant_i$ interactions. The coefficients on these interactions are the estimates of interest, as each represents the marginal effect of slant on the probability of being most viewed, for a given outlet. The outlet fixed effects account for differences in both the number of most viewed articles reported per outlet, and the total number of articles reported per outlet per day. We present results both with and without additional controls for the number of competing horse race headlines at the same time and from the same outlet.

Results are presented in Table 6. There is some evidence of a congeniality effect for Fox in 2012, but for $Slant_1$ only, implying a one unit increase caused an approximately 20 percentage point increase in being most viewed. There are no other significant results for that year. For 2016, there are a few marginally significant results for Fox, the NYT and the WSJ, but each indicates that *less* congenial stories were more likely to be most viewed. The significant effects for Fox are 13-19 percentage points per unit decrease in congenial slant, and for the NYT, 19-28 percentage points. The 2016 results only become stronger when we replace the day fixed effects with a date polynomial, while the 2012 Fox effects disappear in this case (these results are unreported). Of the explanations included in Table 1, only the reputation mechanism is consistent with the 2016 results (assuming there are enough readers of these websites that do not "generally trust" those websites, and would find stories with unusual slants particularly credible). In contrast, the 2012 Fox effects are only consistent with the psychology mechanism.

[11]When we use Google as the omitted outlet in 2016, both Fox's and the NYT's estimated slants are more stable and significant across specifications (see Table A8). In addition, we obtain results indicating that each of the other outlets aside from Fox are slanted to the left of Google. These results are precise and consistent for the WashPost, and imply a slant that is approximately half of that of the NYT. Of the six estimates for each outlet, three are significant (at at least 10%) for the WSJ, four are significant for WashPost, and two are significant for Yahoo, and all of the point estimates are negative.

A supply-side explanation for the 2016 results would be that Fox and the NYT "overshot" the partisanship of their readers in that year, perhaps unintentionally, or perhaps in an attempt to influence readers. Note that the standard errors for the insignificant estimates for the HuffPost, WashPost, and NYT in 2012 are all less than 0.1 (10 percentage points), implying reasonably good power.

The most viewed data could be misleading if total website traffic changed depending on the congeniality of horse race news at the time. For example, suppose Republicans were less likely to visit foxnews.com on days when HR news for Republicans was less congenial, and Fox HR stories were indeed less congenial those days. Even if these stories were more likely to make the most viewed list than an HR story on a more congenial news day, it is possible these uncongenial stories received less total clicks than the more congenial stories. We do not have daily total website click data to directly address this issue, but can use publicly available data to shed some light on it. Appendix A.1 presents smoothed polynomials of Google Trends data on Google searches for "fox news" and "new york times" for the election seasons of 2012 and 2016. The curves all generally trend up over time, more sharply in the final two weeks before the election. There is no evidence of the trends differing between the two outlets in a given year, or of any correlation between these trends and the trends presented in earlier figures using poll and $Slant_i$ data. Thus, there is no evidence that Google searches are correlated with the congeniality of HR news on a given day. Google searches likely only account for a small fraction of total website traffic. Still, if total traffic was substantially correlated with congeniality we would expect to see some sign of this correlation in the search data as well.

Last, we address the final horse race prediction of Table 1: that outlets might report more or fewer horse race stories on particular days depending on the nature of horse race news that day. This prediction can be tested by using outlet-level time series, which allow us to regress the daily number of horse race stories on a measure of "true slant". In these regressions, we control for the importance of horse race news on a given day using a 4th order polynomial of the total number of horse race stories reported by other outlets (we cannot use day fixed effects because there is only one observation per day). Two measures of "true slant" are used: 1) the average slant of stories reported by other outlets that day; 2) the Pollster average difference in polls (percent planning to vote Republican minus percent planning to vote Democrat) that day. Results are largely similar when we also include recent poll changes in the second specification, but cleaner to report when these are omitted. For the first measure, we also include a 4th order date polynomial to further control for general trends in horse race news interest. We do not do this for the second measure because the date polynomial is much more collinear with the poll average variable. We use Poisson regressions because the left-hand side is a count variable; results are similar when we use OLS with Newey-West standard errors.

Results are reported in Table 7. Of the six estimates for the Fox 2012 sample, two are significant at 10%, and there is one estimate significant at 5% for 2016, all consistent with the prediction. There are also two significant estimates consistent with the prediction for the NYT in 2016 (at 5% and 1%) for the other outlets' slant measure. The strongest results, however, are for the WSJ: The majority of estimates in both years are significant at least at the 5%-level and indicate that the WSJ reported more horse race stories on days when news is *less* favorable to Republicans. This is inconsistent with the idea that the WSJ provides news that feels good to a predominantly conservative readership, and perhaps can be explained by the WSJ's news section being surprisingly liberal, consistent with a finding of Groseclose and Milyo (2005).

In summary, we find substantial evidence that Fox horse race news was slanted right (of the other outlets we study) in both 2012 and 2016, and that HuffPost stories in 2012, and NYT stories in 2016, were slanted left, and some evidence of the other outlets we study being slanted to the left of both Google and Fox in 2016. These effects apply to both the content and quantity of headlines, though the evidence is stronger regarding the content. The content effects can be explained by the psychology or reputation mechanism, and the quantity effects are best explained by psychological factors. We find non-robust evidence that particularly right-slanted stories were more popular for Fox in 2012. We find somewhat more robust evidence that stories with slants that went "against the grain" for their outlets were more popular for Fox and the NYT in 2016. The 2012 effect is best explained by the psychology mechanism, and the 2016 effects by the reputation mechanism. We also find that the WSJ reported more horse race stories on days when news was more favorable to Democrats in both years. This result cannot be explained by our theoretical framework; a factor outside our framework that could explain it could be that the WSJ's news staff is more liberal than the opinion staff, and the conservative reputation of the paper is based more on its opinion content.

5 Concluding remarks

Why do people like congenial news? Is it because this news is just more enjoyable to read? Or because they think it is actually more informative? Or because it is truly more informative? We look for clean evidence to distinguish between these mechanisms. Our results are complex. We find both direct and indirect evidence that consumers prefer news for psychological reasons and because they perceive it as more informative. The evidence for the psychology mechanism is perhaps surprisingly weak and inconsistent. Still, the results imply that theories assuming that congenial news demand is only driven by optimal information-seeking behavior are overly simplistic and likely to be misleading.

The results indicating that mean slants were skewed for several outlets, both for all stories and for the most viewed stories of those outlets, imply that those outlets' readers were likely not exposed to accurate information about the candidates' chances throughout the campaigns. Although horse race reporting may seem innocuous, this distorted information may still have been harmful to social welfare, whether the distortions were driven directly by cognitive dissonance or motivated reasoning, or by reputational factors.[12] For example, distorted horse race reporting and news consumption could lead to distrust of election results and conspiracy theories.

Future work building on our paper could consider more detailed individual click-level web data, or other types of media data, such as social media data. And the complexity of our results

[12] Reputational factors could also be indirectly related to cognitive dissonance. If a partisan has a skewed prior about her preferred candidate's chance of winning, this is likely due to biased prior exposure and processing of information about the candidates' chance.

demonstrates (or reminds us) that both media and reader behavior can vary substantially over time and across outlets. In 2016, Fox may have moved to the right of its readers, and the NYT may have further left. The WashPost and NYT have similar reputations, but there seemed to be significant differences in their reporting. The WSJ's reporting appears quite distinct from its reputation (as opposed to that of Fox and the NYT). These results confirm the importance of considering context in interpretation, and the limitations for extrapolating results. While this is always true in empirical work, this point may be especially relevant to the analysis of media in recent years due to the fast-changing environment and uniqueness of the various outlets.

Table 5: Estimated mean differences in slant across outlets

Outlet	$Slant_1$	$Slant_2$	$Slant_3$	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012						
Fox	0.405***	0.515***	0.795***	0.327**	0.440***	0.671***
	(0.127)	(0.126)	(0.165)	(0.134)	(0.132)	(0.189)
WSJ	-0.150	-0.020	0.305	-0.392*	-0.276	0.016
	(0.227)	(0.210)	(0.303)	(0.223)	(0.207)	(0.330)
USAToday	-0.151	-0.023	0.062	-0.185	-0.067	0.039
	(0.176)	(0.179)	(0.222)	(0.192)	(0.198)	(0.236)
NYT	-0.187	-0.031	0.043	-0.258*	-0.141	-0.098
	(0.131)	(0.140)	(0.211)	(0.134)	(0.143)	(0.221)
HuffPost	-0.400***	-0.262**	-0.112	-0.444***	-0.262**	-0.146
	(0.103)	(0.112)	(0.167)	(0.110)	(0.120)	(0.182)
Poll controls?				✓	✓	✓
Adj R^2	0.352	0.330	0.431	0.393	0.372	0.472
N	425	382	277	400	363	267
Panel B: 2016						
Fox	0.313*	0.342	0.685***	0.328*	0.389*	0.712***
	(0.172)	(0.222)	(0.245)	(0.173)	(0.219)	(0.248)
WSJ	-0.276	-0.353	0.152	-0.265	-0.379*	0.083
	(0.208)	(0.231)	(0.269)	(0.202)	(0.228)	(0.280)
Google	0.018	0.050	0.398**	0.055	0.119	0.431**
	(0.136)	(0.161)	(0.157)	(0.134)	(0.160)	(0.178)
NYT	-0.329**	-0.156	-0.024	-0.322**	-0.131	0.048
	(0.158)	(0.201)	(0.187)	(0.157)	(0.200)	(0.204)
WashPost	-0.159	-0.080	0.148	-0.109	0.001	0.194
	(0.149)	(0.174)	(0.171)	(0.148)	(0.170)	(0.193)
Poll controls?				✓	✓	✓
Adj R^2	0.214	0.270	0.367	0.236	0.308	0.405
N	696	501	381	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: $Slant_i$. The reference category for the outlet dummies is Yahoo. All models include day fixed effects. The poll controls include the Pollster daily average poll difference in the support for the Republican and Democratic candidate, as well as the average change in this difference over the previous week. The 2016 models also include dummies for Yahoo stories' first date occurring during one of two time-frames in which Yahoo data collection changed (see appendix). Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table 6: Effects of outlet-slant interactions on the probability of making most viewed list

	$Slant_1$	$Slant_2$	$Slant_3$	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012						
Fox $\times Slant$	0.170**	0.096	0.116	0.208**	0.128	0.166
	(0.083)	(0.103)	(0.135)	(0.093)	(0.121)	(0.166)
WSJ $\times Slant$	-0.010	0.024	0.001	-0.157	-0.117	-0.155
	(0.202)	(0.210)	(0.186)	(0.154)	(0.161)	(0.163)
USAT $\times Slant$	-0.106	-0.087	-0.088	-0.127	-0.124	-0.097
	(0.088)	(0.115)	(0.111)	(0.097)	(0.115)	(0.116)
Yahoo $\times Slant$	-0.043	-0.038	-0.027	-0.046	-0.041	-0.033
	(0.043)	(0.053)	(0.064)	(0.046)	(0.054)	(0.064)
NYT $\times Slant$	-0.013	0.024	0.026	-0.009	0.025	0.024
	(0.072)	(0.076)	(0.087)	(0.073)	(0.074)	(0.092)
HuffPost $\times Slant$	0.010	0.029	0.021	0.004	0.029	0.018
	(0.048)	(0.054)	(0.068)	(0.054)	(0.056)	(0.073)
Competing headline controls?				✓	✓	✓
Adj R^2	0.379	0.408	0.353	0.387	0.403	0.362
N	425	382	277	425	382	277
Panel B: 2016						
Fox$\times Slant$	-0.131*	-0.189**	-0.179*	-0.110	-0.167*	-0.168
	(0.069)	(0.087)	(0.094)	(0.074)	(0.095)	(0.109)
WSJ $\times Slant$	-0.043	-0.187	0.389	-0.064	-0.031	0.275
	(0.244)	(0.523)	(0.293)	(0.246)	(0.418)	(0.377)
Yahoo $\times Slant$	-0.021	-0.009	0.089	-0.024	-0.035	0.040
	(0.087)	(0.116)	(0.141)	(0.095)	(0.155)	(0.258)
Google $\times Slant$	0.016	-0.026	0.009	0.025	-0.011	0.020
	(0.029)	(0.029)	(0.032)	(0.030)	(0.032)	(0.034)
NYT $\times Slant$	0.193*	0.284**	0.511	0.175	0.259*	0.559
	(0.111)	(0.130)	(0.324)	(0.106)	(0.138)	(0.368)
WashPost $\times Slant$	-0.043	-0.013	0.019	-0.042	-0.025	0.019
	(0.056)	(0.072)	(0.082)	(0.058)	(0.077)	(0.097)
Competing headline controls?				✓	✓	✓
Adj R^2	0.286	0.378	0.371	0.292	0.401	0.382
N	696	501	381	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: most viewed (yes/no). The reference category for outlet dummies is Yahoo. All models include day fixed effects. The 2016 models also include dummies for Yahoo stories' first date occurring during one of two time-frames in which Yahoo data collection changed (see appendix). Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table 7: Effects of slant and poll means on the daily number of horse race stories

	Fox	WSJ	NYT	Huff/WashPost
Panel A1: RHS = mean slant of other outlets (2012)				
$Slant_1$	0.310	-0.042	0.312	-0.100
	(0.194)	(0.385)	(0.204)	(0.219)
$Slant_2$	0.277	-0.605	0.188	-0.096
	(0.245)	(0.471)	(0.215)	(0.184)
$Slant_3$	0.114	-0.766**	0.176	-0.369*
	(0.219)	(0.381)	(0.210)	(0.202)
Panel A2: RHS = mean Republican poll advantage (2012)				
$Slant_1$	0.080	-0.604***	-0.061	0.060
	(0.101)	(0.151)	(0.094)	(0.079)
$Slant_2$	0.195*	-0.447***	-0.043	0.040
	(0.104)	(0.163)	(0.088)	(0.081)
$Slant_3$	0.196*	-0.460***	-0.102	0.038
	(0.118)	(0.153)	(0.083)	(0.091)
Panel B1: RHS = mean slant of other outlets (2016)				
$Slant_1$	-0.091	-0.417	-0.367	0.224
	(0.217)	(0.559)	(0.281)	(0.222)
$Slant_2$	-0.198	-1.078**	-0.787***	0.216
	(0.240)	(0.432)	(0.279)	(0.235)
$Slant_3$	-0.146	-1.206**	-0.609**	0.233
	(0.237)	(0.492)	(0.255)	(0.194)
Panel B2: RHS = mean Republican poll advantage (2016)				
$Slant_1$	0.204**	-0.335**	-0.155	0.107*
	(0.081)	(0.135)	(0.100)	(0.062)
$Slant_2$	0.134	-0.363**	-0.156	0.122*
	(0.085)	(0.179)	(0.096)	(0.071)
$Slant_3$	0.091	-0.341**	-0.128	0.076
	(0.102)	(0.162)	(0.136)	(0.090)

Note: Poisson regressions, using daily outlet-level time series. Left-hand side variable: number of horse stories of the outlet (listed in the column header). Mean $Slant_i$ of other outlets equals 0 if no horse race stories of type i are available on a given day. All models include the 4th order polynomial of the number of horse race stories reported by all other outlets; models in Panels A2 and B2 also a 4th order date polynomial. *, **, *** denote 10%, 5%, 1% significance.

References

E. Bakshy, S. Messing, and L. Adamic. Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348:1130–1132, 2015.

P. Barberá and G. Sood. Follow your ideology: A measure of ideological location of media sources. In *MPSA Conference*, 2014.

A. J. Berinsky, G. A. Huber, and G. S. Lenz. Evaluating online labor markets for experimental research: Amazon. com's mechanical turk. *Political Analysis*, 20(3):351–368, 2012.

D. Bernhardt, S. Krasa, and M. Polborn. Political polarization and the electoral effects of media bias. *Journal of Public Economics*, 92(5):1092–1104, 2008.

C. Budak, S. Goel, and J. M. Rao. Fair and balanced? quantifying media bias through crowdsourced content analysis. *Public Opinion Quarterly*, 80(S1):250–271, 2016.

M. Buhrmester, T. Kwang, and S. Gosling. Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011.

J. Burke. Primetime spin: Media bias and belief confirming information. *Journal of Economics & Management Strategy*, 17(3):633–665, 2008.

J. Chan and W. Suen. A spatial theory of news consumption and electoral competition. *Review of Economic Studies*, 75(3):699–728, 2008.

J. Ely, A. Frankel, and E. Kamenica. Suspense and surprise. *Journal of Political Economy*, 123(1):215–260, 2015.

R. Y. Fang. Profit-maximizing media bias. *SSRN Working Paper*, 2016.

S. Flaxman, S. Goel, and J. M. Rao. Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly*, 80(S1):298–320, 2016.

M. Gentkzow, J. Shapiro, and D. Stone. Media bias in the marketplace: Theory. In S. Anderson, J. Waldfogel, and D. Stromberg, editors, *Handbook of media economics*, pages 623–645. Elsevier, 2015.

M. Gentzkow and J. Shapiro. Media bias and reputation. *Journal of Political Economy*, 114 (2):280–316, 2006.

M. Gentzkow and J. Shapiro. What drives media slant? Evidence from US daily newspapers. *Econometrica*, 78(1):35–71, 2010.

M. Gentzkow and J. Shapiro. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839, 2011.

R. Golman, D. Hagmann, and G. Loewenstein. Information avoidance. *Journal of Economic Literature*, 55(1):96–135, 2017.

T. Groseclose and J. Milyo. A measure of media bias. *The Quarterly Journal of Economics*, 120(4):1191–1237, 2005.

A. Guess. Media choice and moderation: Evidence from online tracking data. Working paper, 2016.

Y. Halberstam and B. Knight. Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter. *Journal of Public Economics*, 143:73–88, 2016.

C. Huff and D. Tingley. "Who are these people?" Evaluating the demographic characteristics and political preferences of MTurk survey respondents. *Research & Politics*, 2(3):1–12, 2015.

S. Iyengar and K. S. Hahn. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1):19–39, 2009.

N. Karlsson, G. Loewenstein, and D. Seppi. The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115, 2009.

J. M. Miller and J. A. Krosnick. News media impact on the ingredients of presidential evaluations: Politically knowledgeable citizens are guided by a trusted source. *American Journal of Political Science*, pages 301–315, 2000.

S. Mullainathan and A. Shleifer. The market for news. *The American Economic Review*, 95 (4):1031–1053, 2005.

K. Mullinix, T. Leeper, J. Druckman, and J. Freese. The generalizability of survey experiments. *Journal of Experimental Political Science*, 2(2):109–138, 2015.

S. Oliveros and F. Várdy. Demand for slant: How abstention shapes voters' choice of news media. *The Economic Journal*, 125(587):1327–1368, 2015.

E. Pariser. *The filter bubble: What the Internet is hiding from you.* Penguin, 2011.

Pew. Political polarization & media habits. Pew Research Center, October 2014.

K. Searles, M. H. Ginn, and J. Nickens. For whom the poll airs: Comparing poll results to television poll coverage. *Public Opinion Quarterly*, 80(4):943–963, 2016.

G. Sood and Y. Lelkes. Dont expose yourself: Discretionary exposure to political information. Working paper, 2016.

D. Stone. Ideological media bias. *Journal of Economic Behavior & Organization*, 78(3): 256–271, 2011.

N. J. Stroud. *Niche news: The politics of news choice.* Oxford University Press, 2011.

M. Tremayne. Partisan media and political poll coverage. *Journal of Information Technology & Politics*, 12(3):270–284, 2015.

R. P. Vallone, L. Ross, and M. R. Lepper. The hostile media phenomenon: biased perception and perceptions of media bias in coverage of the beirut massacre. *Journal of personality and social psychology*, 49(3):577, 1985.

A Supplemental materials

Table A1: Summary statistics for survey full sample

Category	Variable (all 0/1)	Mean
Party	Democratic	0.384
	Lean Dem.	0.124
	Independent (no lean)	0.158
	Lean Rep.	0.180
	Republican	0.154
Preferred candidate	Trump	0.261
	Clinton	0.512
	Not voting/other	0.228
Education	Some HS	0.006
	HS degree	0.121
	Some college	0.227
	2 yr degree	0.113
	4 yr coll. Degree	0.415
	$>$ college	0.118
Gender	Female	0.458
	Male	0.542
Age	18-29	0.326
	30-39	0.382
	40-49	0.160
	50-64	0.113
	$\geq$ 65	0.019

Note: N=638 for all variables (N=226 from survey 1, 216 from survey 2, 196 from survey 3) except preferred candidate (N=637; N=195 from survey 3). Respondents are master MTurkers who answered reading check question correctly.

Table A2: Headline options for each survey

Survey 1:	New York Times: Commentators Give Hillary Clinton Edge in Debate
	Fox News: Hillary won the first debate (it helps to be prepared)
	Yahoo News: Long dog-gone trip: Florida pooch travels to Boston and back
	Yahoo News: Houston gunman had two weapons, thousands of rounds at scene.
Survey 2:	New York Times: Who Won the Debate? Commentators Give Edge to Mike Pence
	Fox News: Pence triumphs in VP debate. And then there was the night's biggest loser...
	Yahoo News: 2 Vermont teachers accused of vandalizing sidewalk
	Yahoo News: Two young girls shot in Cleveland drive-by shooting
Survey 3:	New York Times: Who Won the Debate? Donald Trump Avoids Annihilation
	Fox News: Trump comes out swinging and wins second debate
	Yahoo News: Three police officers shot in Palm Springs, California
	Yahoo News: Record 1,201 couples renew wedding vows in Kalamazoo

Note: The instructions given to survey respondents before being asked to choose a headline were: "Choose one of the following articles to read. You should choose the article that you are more interested in the one you would be more likely to read if you came across these links simply surfing the web, or on a social network, etc. After making your choice, you will have access to the article and a simple question on the articles content. Your payment will be \$0.25 higher if you answer the question correctly. The questions difficulty is the same for each article. Thus, you might as well choose the article you are truly more interested in, as it will be more enjoyable to read, and you will be just as likely (or more likely) to get the extra payment." The median work time was approximately four minutes and so our payment, as an hourly rate, was relatively high for an MTurk task, which typically pay at rates less than \$5 per hour.

Table A3: Multinomial logit survey analysis results

	(1)	(2)
Panel A: Party identity		
Dem. congeniality	0.676**	
	(0.269)	
Rep. congeniality	0.229	
	(0.291)	
NYT × Dem. congeniality		0.388
		(0.351)
Fox × Dem. congeniality		0.983***
		(0.366)
NYT × Rep. congeniality		0.043
		(0.407)
Fox × Rep. congeniality		0.209
		(0.378)
Panel B: Candidate preferences		
Clinton congeniality	0.656**	
	(0.256)	
Trump congeniality	0.579**	
	(0.282)	
NYT × Clinton congeniality		0.306
		(0.339)
Fox × Clinton congeniality		1.035***
		(0.345)
NYT × Trump congeniality		0.323
		(0.404)
Fox × Trump congeniality		0.591*
		(0.356)

Note: All models estimated on respondent-news alternative level data set (three alternatives, Fox, NYT, and other, per respondent, and so 1,911 total observations). "Dem. congeniality" = one for Democrats for both the NYT and Fox alternatives for survey 1, equals -1 for both NYT and Fox stories for survey 2, and equals -1 for the Fox story for survey 3, and 0 otherwise (for NYT survey 3 and for all "other" observations). "Clinton congeniality" takes the same values for those alternatives for Clinton supporters, and "Rep./Trump congeniality" take opposite signs for Republicans and Trump supporters, respectively. Models estimated with alternative-specific conditional logit with left-hand side variable of news choice (Fox, NYT or other), and right-hand side respondent (case)-level variables of party/party-strength, education, age, gender and survey fixed effects. *, **, *** denote 10%, 5%, 1% significance.

Table A4: Instructions for rating slant of horse race headlines

Category	Definition
Very good news for Clinton	"Very good news" about Clintons chances of winning (and very bad news about Trump's chances) that Clinton is substantially more likely to win, or that her chances have greatly improved recently
Good news for Clinton	"Good news" about Clintons chances of winning (and bad news about Trump's chances) that Clinton is more likely to win, or that her chances have improved recently
No change	No change in either candidates chance
Good news for Trump	"Good news" about Trumps chances of winning (and bad news about Clinton's chances) that Trump is more likely to win, or that his chances have improved recently
Very good news for Trump	"Very good news" about Trumps chances of winning (and very bad news about Clinton's chances) that Trump is substantially more likely to win, or that his chances have greatly improved recently
Ambiguous or unclear	Relevant to the chances of one of the candidates winning, but unable to determine which candidate is being favored (if at all)
Not relevant	This headline does not seem to be about the candidates' chances of winning the election

Note: MTurkers were given the following instructions: "The following items are real headlines of reports of major US news outlets from July-November 2016 about the upcoming presidential election that year between Hillary Clinton and Donald Trump. Most, but not all, of these articles are about information about which candidate is more likely to win the election. How do you think a typical reader would perceive these headlines? [using the categories above] We would like for you to evaluate these 40 headlines. We will look over your evaluations and if we believe you have done them carefully and reasonably, we will invite you, by email, to do additional similar work (potentially quite a lot)." For 2012 headlines the references were to Obama and Romney.

Table A5: Outlet-slant interactions as predictors of 2012 polls and poll changes

	(1)	(2)	(3)	(4)
Fox $\times Slant_1$	0.835***	0.321**	0.664***	0.227*
	(0.206)	(0.152)	(0.141)	(0.128)
WSJ $\times Slant_1$	0.838*	0.744**	0.578*	0.340
	(0.477)	(0.336)	(0.320)	(0.266)
USA $\times Slant_1$	0.446**	0.544***	0.615***	0.050
	(0.208)	(0.180)	(0.178)	(0.117)
Yahoo $\times Slant_1$	0.752***	0.293**	0.543***	0.174*
	(0.176)	(0.133)	(0.112)	(0.096)
NYT $\times Slant_1$	0.402*	0.812***	0.520***	0.595***
	(0.210)	(0.178)	(0.131)	(0.116)
HuffPost $\times Slant_1$	0.334**	0.531***	0.316***	0.268***
	(0.148)	(0.148)	(0.118)	(0.097)
Adj R^2	0.106	0.186	0.582	0.556
N	400	400	400	400
Fox $\times Slant_2$	0.777***	0.301*	0.593***	0.137
	(0.196)	(0.160)	(0.144)	(0.130)
WSJ $\times Slant_2$	0.904*	0.771**	0.547*	0.337
	(0.484)	(0.325)	(0.326)	(0.262)
USA $\times Slant_2$	0.516**	0.798***	0.757***	0.141
	(0.227)	(0.157)	(0.180)	(0.101)
Yahoo $\times Slant_2$	0.775***	0.273**	0.546***	0.123
	(0.206)	(0.126)	(0.132)	(0.095)
NYT $\times Slant_2$	0.494**	0.837***	0.564***	0.581***
	(0.217)	(0.170)	(0.135)	(0.105)
HuffPost $\times Slant_2$	0.451***	0.505***	0.387***	0.279***
	(0.161)	(0.165)	(0.113)	(0.104)
Adj R^2	0.109	0.190	0.607	0.577
N	363	363	363	363
Fox $\times Slant_3$	0.754***	0.233	0.507***	0.074
	(0.196)	(0.164)	(0.160)	(0.141)
WSJ $\times Slant_3$	0.935**	0.703**	0.536*	0.306
	(0.443)	(0.320)	(0.285)	(0.264)
USA $\times Slant_3$	0.516**	0.772***	0.766***	0.203**
	(0.214)	(0.167)	(0.181)	(0.096)
Yahoo $\times Slant_3$	0.853***	0.310**	0.575***	0.151
	(0.209)	(0.127)	(0.140)	(0.098)
NYT $\times Slant_3$	0.489**	0.833***	0.539***	0.600***
	(0.230)	(0.163)	(0.147)	(0.109)
HuffPost $\times Slant_3$	0.466***	0.549***	0.413***	0.309***
	(0.169)	(0.174)	(0.118)	(0.116)
Adj R^2	0.147	0.256	0.632	0.592
N	267	267	267	267

Note: OLS estimates, using story-level data. Dependent variables: Pollster daily average poll difference in the support for the Republican and Democratic candidate (columns 1 and 3); change in this difference over the past week (columns 2 and 4). The models in columns (3) and (4) include third-order date polynomial. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table A6: Outlet-slant interactions as predictors of 2016 polls and poll changes

	(1)	(2)	(3)	(4)
Fox $\times Slant_1$	0.467**	0.469***	0.183	0.172
	(0.192)	(0.149)	(0.137)	(0.117)
WSJ $\times Slant_1$	0.613*	0.350	0.333	0.017
	(0.327)	(0.241)	(0.273)	(0.183)
Yahoo $\times Slant_1$	0.735***	0.419**	0.355**	0.013
	(0.159)	(0.183)	(0.146)	(0.106)
Google $\times Slant_1$	0.776***	0.361**	0.451***	-0.002
	(0.145)	(0.144)	(0.101)	(0.068)
NYT $\times Slant_1$	0.638**	0.362	0.434***	0.214
	(0.276)	(0.261)	(0.139)	(0.173)
WashPost $\times Slant_1$	0.531***	0.023	0.436***	-0.086
	(0.128)	(0.133)	(0.106)	(0.077)
Adj R^2	0.160	0.045	0.499	0.591
N	696	696	696	696
Fox $\times Slant_2$	0.641***	0.343**	0.376**	0.106
	(0.231)	(0.172)	(0.152)	(0.138)
WSJ $\times Slant_2$	0.732**	0.336*	0.563**	0.099
	(0.364)	(0.198)	(0.263)	(0.183)
Yahoo $\times Slant_2$	0.888***	0.443**	0.399**	-0.024
	(0.175)	(0.219)	(0.160)	(0.117)
Google $\times Slant_2$	0.866***	0.360**	0.453***	-0.053
	(0.156)	(0.161)	(0.114)	(0.073)
NYT $\times Slant_2$	0.717**	0.388	0.464***	0.211
	(0.286)	(0.311)	(0.142)	(0.178)
WP $\times Slant_2$	0.679***	0.064	0.494***	-0.103
	(0.173)	(0.129)	(0.123)	(0.107)
Adj R^2	0.210	0.044	0.580	0.574
N	501	501	501	501
Fox $\times Slant_3$	0.694***	0.313*	0.453***	0.056
	(0.232)	(0.170)	(0.160)	(0.139)
Fox $\times Slant_3$	0.633*	0.283	0.501*	0.094
	(0.356)	(0.196)	(0.265)	(0.186)
Fox $\times Slant_3$	0.834***	0.355	0.429**	-0.081
	(0.180)	(0.235)	(0.169)	(0.128)
Fox $\times Slant_3$	0.807***	0.326**	0.460***	-0.061
	(0.152)	(0.153)	(0.109)	(0.073)
Fox $\times Slant_3$	0.673**	0.365	0.492***	0.210
	(0.288)	(0.312)	(0.138)	(0.180)
Fox $\times Slant_3$	0.641***	0.025	0.498***	-0.120
	(0.176)	(0.130)	(0.127)	(0.112)
Adj R^2	0.229	0.036	0.568	0.591
N	381	381	381	381

Note: OLS estimates, using story-level data. Dependent variables: Pollster daily average poll difference in the support for the Republican and Democratic candidate (columns 1 and 3); change in this difference over the past week (columns 2 and 4). The models in columns (3) and (4) include third-order date polynomial. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table A7: *Slant* values for headlines from November 7, 2016

Headline	$Slant_1$	$Slant_2$	$Slant_3$	Outlet	MV
polls trump and clinton virtually tied in key swing states	0.00	0.00		Fox	1
momentum buster? fbis comey tells congress	-0.50			Fox	0
email review completed decision not to prosecute clinton stands					
trump supporters say they feel michigan momentum	0.67	0.67		Google	0
president obama makes closing argument if we win florida its a wrap	-0.50			Google	0
poll on eve of election day clinton maintains her edge over trump	-1.00	-1.00	-1.00	Google	0
iowa poll trump opens 7point lead over clinton	1.00	1.00	1.00	Google	0
us presidential election live countdown to the polls	0.00			Google	1
live blog last updated 7.30am aest us election live trump clinton	0.00			Google	0
in final pitch to voters latest polls					
obama if clinton wins florida she will win the election	-0.50			Google	0
our final map has clinton winning with 352 electoral votes. compare your picks with ours.	-1.00	-1.00	-1.00	Google	0
clinton has solid lead in electoral college trumps winning map is unclear	-1.00	-1.00	-1.00	NYT	1
trump and clinton tied in final upshot poll of north carolina	0.00	0.00		NYT	1
clinton cleared on new emails keeps small lead in polls	-1.00	-1.00	-1.00	NYT	0
clinton leads trump by 4 points in lates poll	-1.00	-1.00	-1.00	WSJ	1
1 hillary clinton has enough electoral votes to win the white house in final fix map	-1.00			WashPost	1
postabc tracking poll clinton 47 trump 43 on election eve	-1.00	-1.00	-1.00	WashPost	1
amid lastminute push in va. clinton holds 6point lead in latest poll	-1.00	-1.00	-1.00	WashPost	0
trump urges voters to deliver justice at polls	0.50			WashPost	0

Note: MV = most viewed.

Table A8: Estimated mean differences in slant across outlets in 2016 with Google as reference outlet

Outlet	$Slant_1$	$Slant_2$	$Slant_3$	$Slant_1$	$Slant_2$	$Slant_3$
Fox	0.295**	0.292*	0.288	0.274**	0.270*	0.280
	(0.117)	(0.157)	(0.183)	(0.116)	(0.154)	(0.174)
WSJ	-0.293	-0.402**	-0.245	-0.320*	-0.498***	-0.348
	(0.190)	(0.191)	(0.240)	(0.187)	(0.184)	(0.233)
Yahoo	-0.018	-0.050	-0.398**	-0.055	-0.119	-0.431**
	(0.136)	(0.161)	(0.157)	(0.134)	(0.160)	(0.178)
NYT	-0.347***	-0.206	-0.422***	-0.377***	-0.250**	-0.383***
	(0.108)	(0.134)	(0.136)	(0.100)	(0.122)	(0.134)
WashPost	-0.177**	-0.129	-0.249**	-0.164*	-0.118	-0.237*
	(0.084)	(0.094)	(0.125)	(0.083)	(0.094)	(0.123)
Poll controls?				✓	✓	✓
Adj R^2	0.214	0.270	0.367	0.236	0.308	0.405
N	696	501	381	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: $Slant_i$. The reference category for the outlet dummies is Google. All models include day fixed effects. The poll controls include the Pollster daily average poll difference in the support for the Republican and Democratic candidate, as well as the average change in this difference over the previous week. The 2016 models also include dummies for Yahoo stories' first date occurring during one of two time-frames in which Yahoo data collection changed (see appendix). Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Figure A.1: Google Trends data (day 0 = election day)

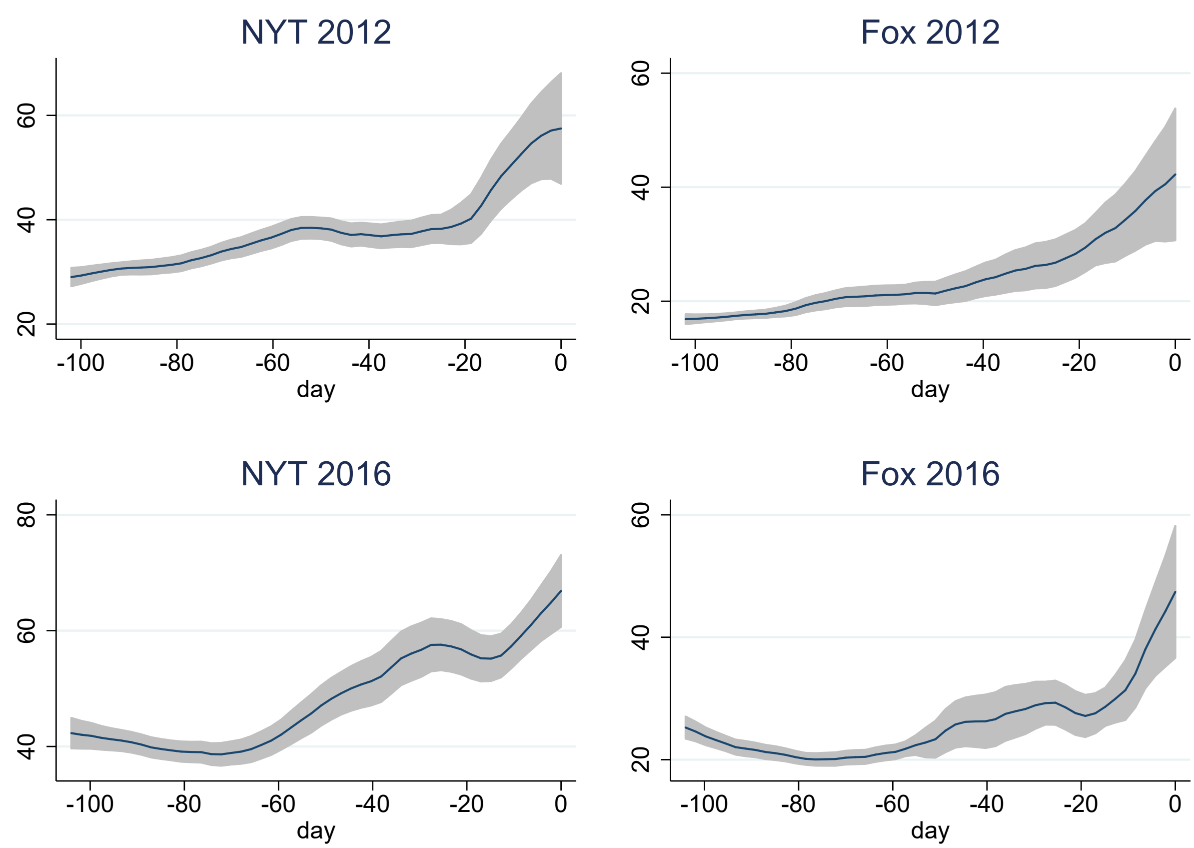

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. "NYT" = Google searches for "new york times"; "Fox" = Google searches for "fox news".