What Drives Demand for Media Slant?*

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

October 9, 2018

Abstract

We conduct within-outlet, within-topic analysis of the relationship between partisan congeniality of news and news demand. We study "horse race" news on the 2012 and 2016 US presidential campaigns from six major online outlets, and data from incentivized surveys on news on the winners of presidential campaign debates. We find some evidence of higher demand for more congenial stories (within-outlet-topic), and some evidence of higher demand for more *uncongenial* stories. We argue the former evidence is most consistent with psychological theories of demand for slant, and the latter is most consistent with the trust theory. We also obtain evidence of systematically congenial outlet-level horse race slant, arguably consistent with both the psychology and trust theories, but inconsistent with the instrumental value theory.

*Our title alludes to that of Gentzkow and Shapiro (2010) ("What Drives Media Slant?"). We thank Matthew Gentzkow, Jesse Shapiro, Greg Martin, Matt Botsch, and participants at the 2018 Economics of Media Bias Workshop, 2018 European Public Choice Society, 2018 Behavioral Models of Politics conference, and 2018 Maine Economics Conference for very helpful comments, and also thank Isaiah West for excellent research assistance. This research was funded by a Bowdoin College Grua/O'Connell Research Award.

†Marcel is Assistant Professor of Economics at Jönköping University, Sweden. He can be reached at: marcel.garz@ju.se.

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

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

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

1 Introduction

The information age has brought to fore an age-old problem: people tend to seek, and be too credulous of, belief-confirming information. The dramatic lowering of search costs and increase in the number of choices means that, now more than ever, people can easily get their political news from ideologically like-minded outlets. This greater ease of "selective exposure" has fomented concern that people will silo themselves in "filter bubbles" and "echo chambers," phenomena widely seen as harmful to democracy.[1]

Are these concerns well-founded? Understanding the effects of partisan selective exposure requires understanding the causes. Gentkzow et al. (2015) propose a taxonomy of three theories. First, consumers may rationally prefer a like-minded news source because it provides optimal instrumental value for decisions (e.g., Burke, 2008; Chan and Suen, 2008; Oliveros and Várdy, 2015; Fang, 2016). For example, in Chan and Suen's model, like-minded news sources provide media consumers with more useful information about whom to vote for.[2]

The second, and perhaps most intuitive, theory is that media consumers may prefer like-minded news because it provides "psychological utility" (Gentkzow et al., 2015). Belief-confirming news may be pleasant, and belief-challenging news unpleasant, due to cognitive dissonance, ego and identity effects, and related factors.[3]

The third theory is that media consumers prefer like-minded sources because they are perceived as objectively more credible (Gentzkow and Shapiro, 2006), though perhaps mistakenly (Vallone et al., 1985; Stone, 2011; Kelly, 2018). Note that in both this case and the instrumental value case, consumers are information-seeking. But these two theories differ in

[1]See, for instance, Pariser (2011), Bakshy et al. (2015), Flaxman et al. (2016), Halberstam and Knight (2016), and Peterson et al. (2018).

[2]Gentkzow et al. (2015) discuss how models making the point that slanted media outlets can provide optimal instrumental value 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 media consumer who uses the news to make a decision may be better off getting this news from a media outlet with a similar political viewpoint. Bruns and Himmler (2016) provide a model showing why rational voters may be willing to pay for instrumental information, even though the probability that an individual vote is decisive is small.

[3]See, for example, Iyengar and Hahn (2009), and from economics, Mullainathan and Shleifer (2005) and Bernhardt et al. (2008). Stroud (2011) discusses three psychological factors that could cause such demand: cognitive dissonance, motivated reasoning, and "cognitive misers" (that processing conflicting information takes more resources than consistent information).

two key ways: 1) the instrumental value theory implies that consumers agree on the accuracy and slants of the different outlets, while the credibility theory implies disagreement; 2) the instrumental value theory only applies to decision-relevant news, while the credibility theory applies to information with either instrumental or intrinsic value.

There are important welfare distinctions between these three theories—which we refer to, for short, as the instrumental, psychology, and trust mechanisms.[4] If consumers choose like-minded news because it has greater instrumental value, then proliferation of diverse media has greater social benefits (Gentkzow et al., 2015), and concerns about selective exposure are more likely to be overblown. If consumers choose like-minded sources because they are more trusted, then consumers may at least be willing to see a variety of stories from these outlets, to the extent that such stories are available. Consumers would then be better informed than if they were only willing to view congenial news from like-minded outlets.[5] If psychology is the primary mechanism, then consumers are likely least well informed, and least capable of performing their jobs as voters.

In this paper, we try to shed some light on the primacy of these three (not mutually exclusive) mechanisms. Our paper is the first, to our knowledge, addressing this subject with real-world, real-time news data. Our empirical strategy is to study news topics reported on repeatedly by various major outlets with clear variation in congeniality of the stories across time within each outlet. We look to see: 1) how congeniality varies on average, for a given topic, across outlets, and 2) how variation in congeniality affects demand within-outlets (and within-topic). Holding fixed the outlet and topic holds fixed two key components of the information value (and perceived information value) of news. Thus, if demand is higher for more congenial articles (within-outlet-topic), this would likely not be due to information, and would instead provide relatively clear evidence of the psychology mechanism. Interpretation of other results is somewhat more complicated, but, we claim, still informative, especially when the two types of analysis are interpreted in conjunction with one another, as we discuss

[4]Gentkzow et al. (2015) refer to the instrumental and trust mechanisms as "delegation" and "reputation", respectively.

[5]We use the terms "congenial" and "like-minded" in distinct ways: "Congenial" refers to a story providing politically favorable news for a given partisan reader, while "like-minded" refers to an outlet being ideologically aligned with a reader.

in Section 3. Regardless, a straightforward description of our contribution is that we present a novel study of within-outlet-topic selective exposure outside the lab.

The first topic that we study, and our main focus, is "horse race" news—news on the chances of candidates winning an upcoming election. We analyze horse race news (headlines specifically) from the 2012 and 2016 US presidential campaigns by six major outlets with varying ideological reputations: the New York Times (NYT), Fox News, Wall Street Journal (WSJ), and Yahoo News in both 2012 and 2016; Google News and the Washington Post (WashPost) in 2016; and USA Today (USAT) and the Huffington Post (HuffPost) in 2012. Horse race news offer the advantage of relatively clear congeniality for one party or the other: if a headline indicates that one party is winning, or gaining ground, then the headline is congenial for that party, and uncongenial for the other party. Furthermore, horse race news has little relevance to most decisions, and so any systematic outlet-level slant is likely evidence of non-instrumentally driven slant. We focus on headlines because they are the main source of information on article content, aside from the outlet itself, observed by readers before deciding whether to click on an article link. Moreover, headline congeniality is relatively easy to code and an accurate predictor of article content congeniality.

We describe the data in Section 4, scraped web headlines and "most viewed" lists to capture link popularity, and present results in Section 5. We find that horse race headlines are indeed systematically skewed in a congenial way for an outlet's typical reader in several cases, consistent with the psychology mechanism and perhaps also the trust mechanism (if consumers have motivated beliefs about electoral outcomes). Fox News horse race headlines generally favored the Republican ticket's chances in both years, as compared to stories from the other outlets. Huffington Post headlines favored the Democrats in 2012, as did the NYT and, to a lesser extent, the Washington Post, in 2016. This favoritism mostly occurred on the intensive margin of reporting (slant for a given horse race headline), rather than the extensive margin (reporting more or fewer horse race stories depending on what the actual news of the day was). We also obtain marginally significant evidence that the WSJ's reporting favored the Democrats' chances in both years. Given the outlet's conservative reputation, this is one case where such news could have plausibly provided more instrumental value to readers. For example, the WSJ's typical reader may have been more likely to donate or support their preferred campaign when horse race news was more uncongenial.

We find only weak evidence of higher demand for more congenial headlines within a given outlet (supporting the psychology mechanism): a non-robust result for just one outlet in one year, Fox in 2012. In 2016, we find evidence of the reverse relationship for Fox News and the NYT, with more *uncongenial* stories being more popular. While this result could have several explanations, one that we find most plausible is that readers were more interested in these less congenial headlines because they were perceived to signal particularly credible stories, consistent with the trust mechanism. And as discussed in Section 3, this result and interpretation are still also consistent with the psychology mechanism also being important, given that average slant from these outlets was congenial. That is, readers may care about both news congeniality for psychological reasons and accuracy, and thus prefer that an outlet's content be slanted congenially in general, but particularly trust (and be more likely to click on) the occasional uncongenial story.

An issue with the web data analysis is that we do not know who exactly is doing the clicking. We address this concern with our analysis of the second news topic we examine, US presidential debate winners, presented in Section 6. For this analysis we use micro-level data from incentivized Amazon Mechanical Turk (Mturk) surveys. We ran surveys after the first three 2016 debates, offering subjects an incentivized choice of real and timely headlines for one story on the debate winner from the NYT, one from Fox News, and two Yahoo headlines on other topics. We use the topic of debate winners, and not the horse race, because the timing of debate stories is known in advance. The two news topics are closely related but distinct, allowing the analyses to complement one another in subtle ways. Debate stories may contain more voting decision-relevant information than horse race stories. However, again, as we discuss, greater demand for more congenial news within one's preferred outlet would most likely be due to the psychology mechanism, and greater demand for congenial news across or within-outlets in general could be due to either psychology or trust but is unlikely to have an instrumental motive.

Survey respondents on both sides of the political spectrum—but especially Democrats—

were more likely to read debate news when it was more congenial for their side. These effects, across the spectrum, were mostly driven by changes in the demand for Fox News, whereas NYT debate news was not significantly sensitive to the congeniality of the headline. We interpret these results as more supportive of the psychology mechanism for respondents who were Trump supporters, and more ambiguous for Democrats (supportive of psychology or trust), and in general failing to support the instrumental theory.

We provide concluding remarks in the final section. We also discuss how the psychology and trust mechanisms are perhaps more closely related than they appear: trust is influenced by psychology, and psychological value depends on trust. (For news to make us "feel good" it must be at least somewhat credible.) The latter point has not been widely recognized in the media economics literature, and implies a natural limit to the degree of demand-driven distortion of facts. Still, both the psychology and trust mechanisms lead to misleading reporting on factual issues, and so our results support the importance of policies designed to mitigate echo chambers and enhance exposure to high quality and/or diverse news sources.

2 Related literature

In addition to the media bias literature discussed above, our study contributes to research on partisan selective exposure to information; see Hart et al. (2009) and Stroud (2011) for reviews in psychology and communications research. The economics literature on the demand for congenial information outside of political media is better developed; see Golman et al. (2017) for a review. For instance, Karlsson et al. (2009) find that investors exhibit the "ostrich effect," checking their portfolios more often when markets are rising rather than falling, consistent with what we call the psychology mechanism. More recent studies examine partisan selective engagement, the social media analogue of partisan selective exposure—i.e., liking, sharing and commenting; see Garz et al. (2018) and Pogorelskiy and Shum (2018).

In contrast to most of the existing literature, we do not only investigate partisan selective exposure per se but also the underlying explanations. A few studies investigate these explanations in laboratory experiments. Using artificial news articles, Metzger et al. (2015) find some support for the psychology mechanism, but stronger support for the importance of trust in driving news choices. Masatlioglu et al. (2017) investigate subjects' behavior in lotteries and show that the desire to manage anticipatory emotions dominates the preference for instrumental information in many situations, which emphasizes the psychology mechanism. In contrast to these studies, we investigate the demand for congenial information based on observational data, and surveys that exploit real and timely news articles. In this regard, our study closely relates to Simonov and Rao (2018), who use observational data from Russian news outlets to investigate preferences for pro-government bias.

Finally, our research relates to previous studies that investigate bias in horse race reporting (Tremayne, 2015; Searles et al., 2016). In contrast to these studies, our investigation is not limited to differences in slant across across outlets. We also investigate the within-outlet relationship between congeniality and news demand, which allows us to shed light on the underlying demand mechanisms.

3 Theory and empirical strategy

Our empirical strategy involves two main types of analysis: 1) estimation of average outlet-level headline slant and 2) estimation of the headline slant-news demand relationship within outlets.[6] We present a slightly modified version of the model of Mullainathan and Shleifer (2005) in the appendix to illustrate how these two types of analysis are related, and therefore must be interpreted in conjunction with one another. We summarize the model in the next subsection, then discuss broader interpretation, and then summarize the key claims that we take to the data.

[6]We study slant of just headlines, and not article content, for both types of analysis. We do this to be consistent across analyses and because headlines are likely the main driver of variation in demand (clicks) within outlets, the subject of our second analysis. Furthermore, headline and article congeniality are generally well aligned as shown by, e.g., Tremayne (2015), which we confirm in spot checks. A typical horse race article will summarize new poll results, while the article discusses details of the results for different demographic groups.

3.1 Model

The key model assumptions are as follows. A representative consumer's utility from reading news is increasing in the story's accuracy and/or its congeniality. The mechanism for why congeniality may increase utility is not directly modeled. News is reported repeatedly by a single outlet. The outlet is engaged in monopolistic competition (only a particular segment of the entire news readership is "at play") with other outlets not modeled. The reader sees each story's headline before deciding whether to click on and read the story. The outlet strategically chooses an average level of slant across stories, and the congeniality of each news story's content is driven by that of the true news for the story, average slant, and idiosyncratic noise. The headline is an imperfect signal of each component of the story: if a headline is more congenial, this signals that the story contains both more congenial true news (which is always good for the reader) and more congenial noise (which may be bad for the reader).

The marginal effect of headline congeniality on the reader's expected utility from clicking and reading the story decreases as the outlet's average slant becomes more congenial. This marginal utility may be negative if average congenial slant is large enough. Click probability is increasing in expected utility, so the sign of the marginal effect of headline congeniality on reader utility is the same as that of the marginal effect of headline congeniality on clicks. Consequently, there can be a negative relationship between headline congeniality and news demand, even if the reader has a preference for congenial news. Relatively congenial headlines can signal that stories are too inaccurate for the reader's tastes. If the outlet uses a smaller, though perhaps still non-zero, congenial average slant, then the marginal utility of headline congeniality is more likely positive, and demand is more likely to increase as headline congeniality grows.

3.2 Interpretation

The model implies that either average outlet-level congenial slant or a positive within-outlet congeniality-demand relationship, or both, would indicate a reader preference for congenial news. The natural interpretation of this preference, consistent with Mullainathan and Shelifer's discussion, is what we call the psychology mechanism. Reading more favorable news about one's preferred candidate could be more enjoyable because it increases anticipation utility, such as increased optimism in the desired electoral outcome, or confirmation of the broader validity of one's political views. A positive within-outlet congeniality-demand relation would indicate that average slant is not high enough to make the marginal utility of headline slant negative.

The model also implies that even with a reader psychological preference for congeniality, an outlet's average slant can be too congenial for readers' tastes, which could cause demand to be higher for relatively uncongenial news from the outlet. This would be evidence of the trust mechanism: readers put more trust in stories with relatively uncongenial headlines from a typically congenial outlet, since these headlines signal relatively high accuracy. Note that an outlet's average slant could be higher or lower than the slant that would maximize reader utility, and total clicks, for several reasons beyond our simple model. These include a desire to influence the public's beliefs, to manage the outlet's reputation, to pander to politicians for access or perhaps other reasons, or a simple failure to optimize.[7]

An additional important factor to consider is that readers may have biased beliefs about the state of the horse race (see, e.g., Stiers and Dassonneville, 2018) for both endogenous and exogenous reasons. If so, readers may trust congenially slanted stories more than they should. This could possibly explain average outlet-level congenial slant, even if readers have no psychological preference for congeniality. However, this factor is unlikely to explain higher demand for relatively congenial stories for outlets with an average congenial slant. Since reader bias about the state of the horse race is at least partly endogenous and based on the outlet's past slant, it would not make sense for readers to find the outlet's especially congenially slanted stories especially credible.

As noted above, we think it is implausible that horse race news demand is typically driven by decision-relevance. However, horse race news could provide useful information for some

[7] Gentzkow and Shapiro (2010) is the most well-known work showing that slant is driven more by demand-side than supply-side forces (for US newspapers in 2005), but even their work does not imply that media outlets perfectly cater to reader's tastes, and regardless their results of course do not necessarily apply equally to different contexts. Martin and Yurukoglu (2017) provide evidence that media bias affects consumer ideology (in addition to causality going the other way).

decisions, such as how much money or time to invest in supporting a campaign. Still, it seems implausible that congenial outlet-level slant is driven by these factors—that readers would find positively distorted or selected news more useful for these decisions. This type of instrumental value seems most likely to occur when the news indicates that a campaign is relatively close, or perhaps when the news is negative for the preferred candidate indicating help is more needed. Even then instrumental value would likely only be relevant for a small fraction of consumers, perhaps those who are more affluent and are more likely to make substantial donation decisions.[8]

3.3 Summary

In summary, the main theoretical ideas that we use to motivate and interpret the empirical analysis are:

1) Average congenial slant for partisan outlets is likely due to the psychology mechanism and/or the trust mechanism, and is unlikely to be due to the instrumental mechanism.

2) Higher demand for relatively congenial stories from generally congenial outlets is most likely explained by the psychology mechanism.

3) Higher demand for relatively uncongenial stories from generally congenial outlets is most likely explained by the trust mechanism.

Again, we acknowledge that there are other potentially relevant factors beyond the three demand mechanisms we focus on, the supply-side factors discussed, and the model presented in the appendix. Due to the complexity of the context we cannot discuss all relevant factors here, but two others worth noting are readership heterogeneity and surprise utility. We address heterogeneity with our micro-level survey data. We address surprise utility—that information-seeking consumers might be more likely to click on a headline when it is more surprising or extreme (Ely et al., 2015)—in Section 5.2.

[8] One other type of instrumental value of horse race news is social learning or bandwagon effects: some readers may be more likely to support candidates doing well in recent polls for these reasons. This instrumental value would not be correlated with congeniality of news.

4 Data

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, WSJ, NYT, WashPost, Google, and Yahoo were scraped three times daily until the day of the election, November 8. See, for example, Pew (2014) for evidence of variation in ideological readership among these outlets. We downloaded outlets' landing pages, politics sections, and most viewed lists. We chose these six outlets because of their prominence, ideological diversity, and because each of them publicly reports "top," "trending," "most popular," or "most viewed" stories; we use "most viewed" as shorthand to refer to all of these categories. 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. We were forced to make substitutions for Google and WashPost due to their snapshot data being unavailable and replace them with USAT and the HuffPost, respectively.[9] Snapshots are available, but much more sparse, for 2008 and earlier election years, and so we do not collect snapshots for any years prior to 2012.

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. Next, we identified a set of articles that were likely to be horse race stories. (Ultimately, whether or not a story is about the horse race will be determined by human coding, as we discuss below.) We used a fairly broad set of keywords to make the initial set inclusive and then narrowed

[9]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 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. Scraping historical Google and WashPost data using the websites' sitemaps is not technically feasible, to our knowledge.

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. Through manual checks of random subsamples, these keywords were determined to lead to a very small fraction of false negative classifications (actual horse race 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, as compared to the cost of false negatives, lower sample size. 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. We used the following headline keywords, determined 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}. This should reduce the prevalence of within-outlet variation in an article's credibility and information content. For example, certain readers might perceive certain opinion writers to be particularly credible. With the exceptions of opinion pieces, readers do not see the author (if any) before clicking on the headline, in most cases.

Because headlines are short, the degree to which they favored one party or the other is relatively easy to code accurately manually, and difficult to do so computationally, as we discuss further below. We therefore used manual content analysis for the coding of headline congeniality used for analysis. We assigned three "master" MTurk workers, having them each 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 "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 A1 in the Online Appendix. We did not allow the Mturkers to see the names of the outlets, so that ratings are based on headline content and comparable across outlets.[10] These workers each passed initial

[10]Some headlines did refer to Fox News polls; results are largely similar when these are dropped.

screens of the quality of their work. To incentivize continued effort, while avoiding excessive monitoring and potential demand effects, we kept instructions intentionally vague. We did not specify additional payment for particular results, but simply offered the incentive of generously paid additional work (coders were paid $3 for each batch of 40 headlines) if the work was done "carefully and reasonably." We monitored the coding done by these MTurkers by choosing four headlines with relatively unambiguous favorability to one party or the other, 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 across the three coders); 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 coders for the vast majority of the coding for consistency and because accuracy may have improved with experience.

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.[11] 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, which is perhaps unsurprising (it is easier to code whether a headline simply favors one side or the other than the strength of favoritism toward that side). Consequently, we only use the three-point scale going forward, with very good or 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

[11] 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.

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 $Slant_1$. Most, but not all, $Slant_2$ headlines are also $Slant_1$ headlines. Note that $Slant_i$ simply measures favorability of a headline to one party or the other and not necessarily distortion or any type of misreporting. That is, $Slant_i$ corresponds to h in the appendix model and not s.

Figures 1 and 2 present smoothed plots of daily means of each slant measure versus daily poll averages (i.e., percent planning to vote Republican minus percent planning to vote Democrat; obtained from R's "pollstR" library). The plots are quite similar to one another, and to the polling average, supporting the validity of all three slant measures. However, a number of results presented in Section 5 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 A2 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.

We also evaluate but discard the option of using a coder-independent, pure text-based measure of slant. We coded headlines as favoring a candidate if they contain the candidate's last name and "win" (which could be part of winning or winner) and "lead" (which could be part of leading) and do *not* contain the opponent's name, or "[candidate's name] lead" (e.g., "Trump leads" or "Trump leading"). This measure of slant correlates with our $Slant_i$ measures (average correlation of 0.33), but not with the daily poll average (correlation coefficient = 0.01). The correlation of each of $Slant_i$ with the polls is approximately 0.42, which supports the superiority of the human-coding approach and the use of these measures in the analyses.

Table 1 reports the number of unique most viewed and other articles per outlet, and

Figure 1: Mean $Slant_i$ and poll average versus day relative to election day (day 0) in 2012.

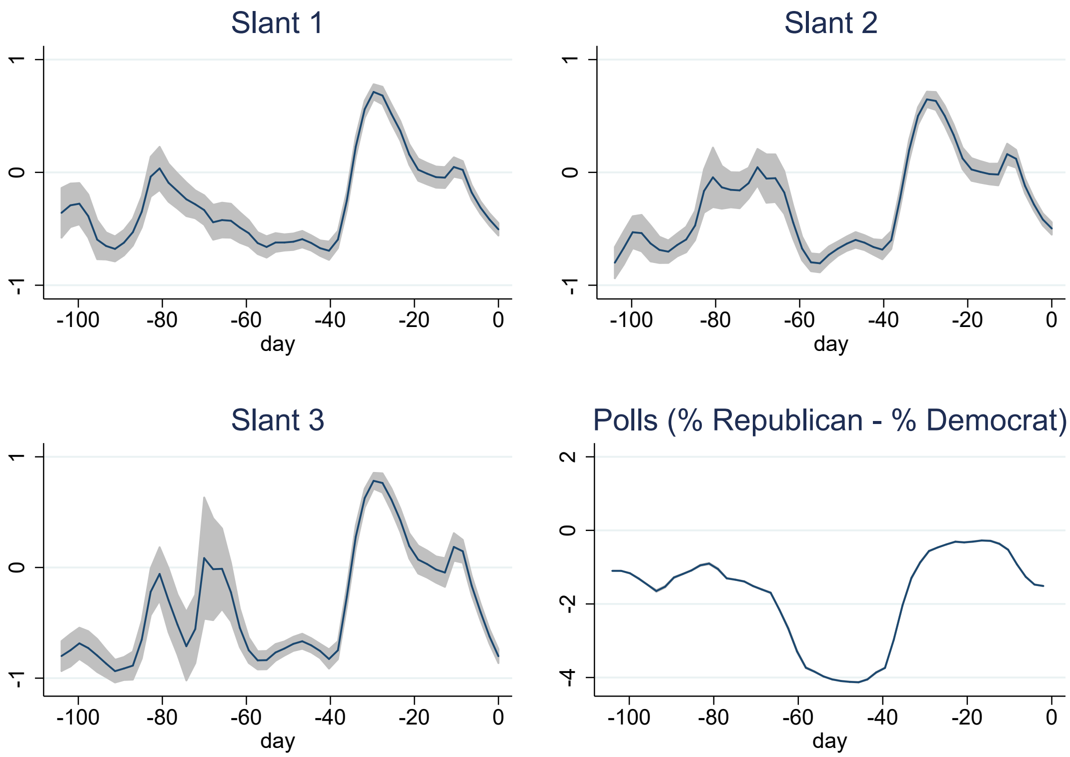

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. Positive values of *Slant* denote better chances of winning for the Republican candidate, whereas negative values indicate better chances for the Democratic candidate.

their means for various slant measures.[12] 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 3 and 4 present smoothed polynomials with 95% confidence bands of the relationship of $Slant_1$ and

[12]The number of headlines reported in this table is less than the corresponding number referred to in the Krippendorff α analysis (2,025) because the sample used in this table, and for most of the subsequent analysis, differs 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 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 2: Mean $Slant_i$ and poll average versus day relative to election day (day 0) in 2016.

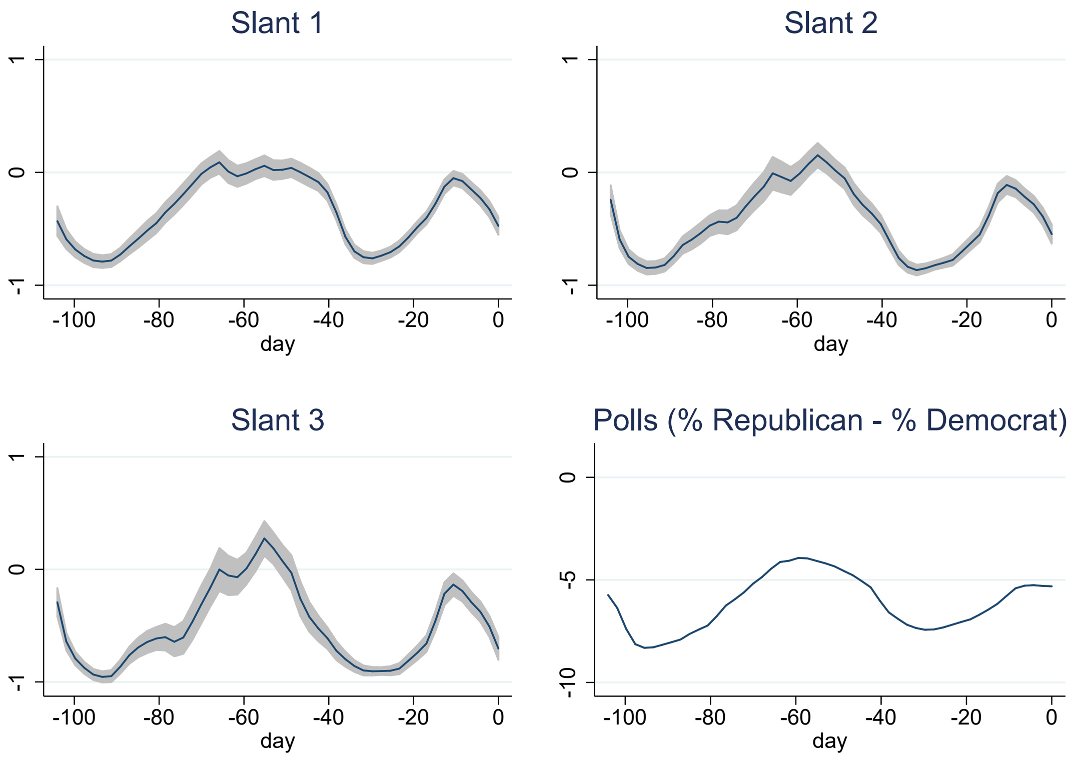

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. Positive values of *Slant* denote better chances of winning for the Republican candidate, whereas negative values indicate better chances for the Democratic candidate.

days to the election for stories that made the most viewed list that day, and for all other stories, for each outlet and year (the graphs are similar when using $Slant_2$ and $Slant_3$). The figures use scraped headline-level data, not story-level data, i.e., 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, both within and across contiguous days, but the bands can be misleading as they can be small (or non-existent) due to limited data. Still, the trends and confidence bands are useful for illustrating broad trends in, and the 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 1. The most viewed stories show the latter bump as well, while only Fox's most viewed stories reflects the earlier bump. The most viewed stories for

Figure 3: Mean $Slant_1$ versus day relative to election day (day 0) by outlet in 2012

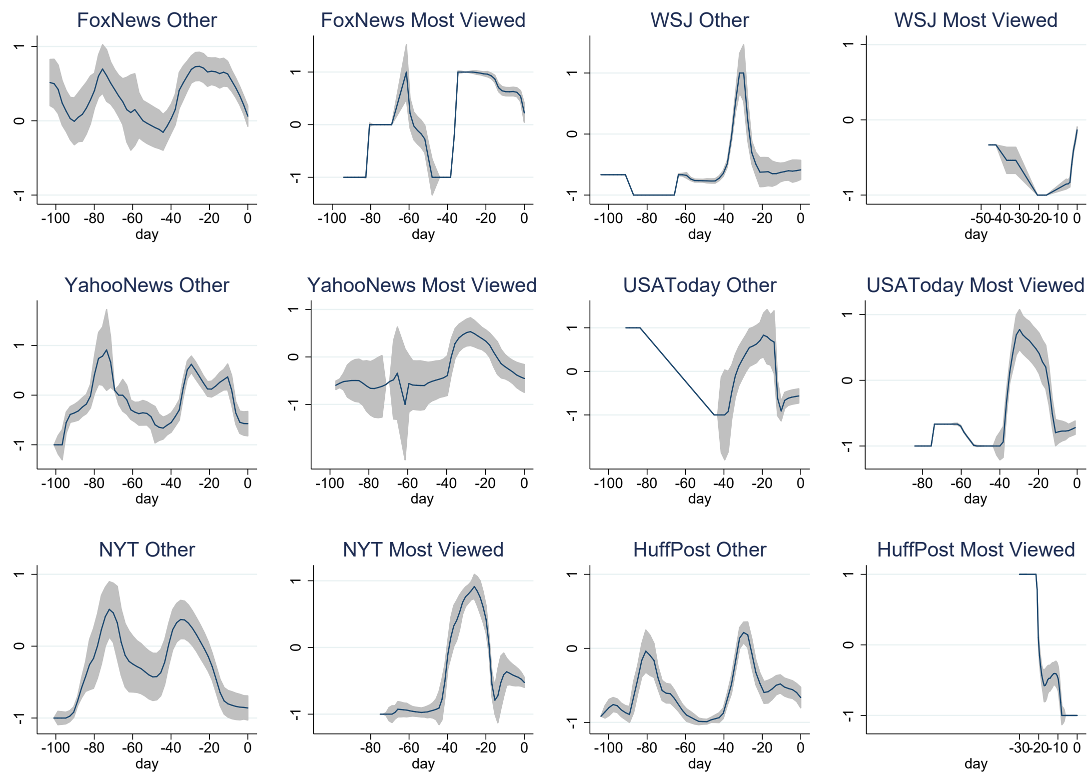

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. Positive values of *Slant* denote better chances of winning for the Republican candidate, whereas negative values indicate better chances for the Democratic candidate.

HuffPost and WSJ all come from the last 30–40 days of the sample time-frame. For 2016, almost all outlets' most viewed and "other" news show pro-Republican bumps around 60 days prior to the election (50 days for WashPost), and also around 10 days prior to the election, with the NYT "other" stories being a notable exception. The pattern of slant in Fox's most viewed stories was particularly volatile. 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 regardless.

Figure 4: Mean $Slant_1$ versus day relative to election day (day 0) by outlet in 2016

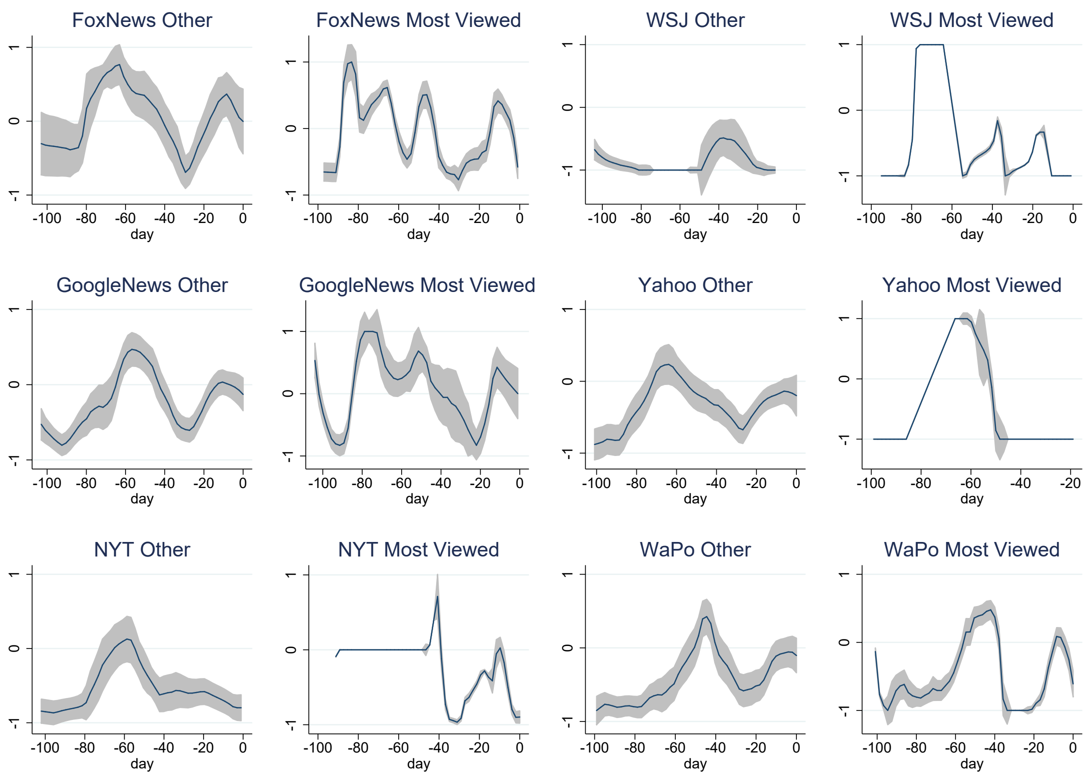

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. Positive values of *Slant* denote better chances of winning for the Republican candidate, whereas negative values indicate better chances for the Democratic candidate.

5 Analysis

5.1 Slant across outlets

We first address the question if and how mean slant of stories varied across outlets and years. There are two basic ways this could occur: on the intensive margin (how does mean horse race slant vary by outlet, given the "true" horse race news?) and on the extensive margin (how many horse race stories were published, given the "true" horse race news of the day?). To address the former, we use a story-level data set to estimate regressions of the following form:

$$Slant_{ijkt} = \alpha + Outlet_j + Day_t + \epsilon_{ijkt}. \tag{1}$$

Here, $Slant_i$ of each story k is regressed on outlet j and day t fixed effects. The day fixed effects account for mean slant across outlets on the first day that the story was reported driven by new poll results or other factors. We do not include daily polls and poll changes, as they are fully collinear with the day fixed effects. Results are similar, with somewhat smaller standard errors and lower R^2 values, when we replace the day fixed effects with a date polynomial.

We use Yahoo as the reference category in 2012, and Google in 2016, given these outlets' relative neutrality and the large number of horse race stories in their respective years. The outlet fixed effects thus represent the mean difference in slant for that outlet as compared to Yahoo in 2012, and Google in 2016, conditional on mean slant of stories reported on that day across outlets. 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. We cluster standard errors by day.

Table 2 presents results. In both years, there is robust evidence that Fox's stories were slanted to the right of the other outlets. The leftmost stories on average in 2012 were from the HuffPost, and NYT's stories were left of Yahoo's (at 5% significance) in one model. The point estimates for each non-Yahoo FE are all largest for the $Slant_3$ models, implying that Yahoo's slant shifted left for this measure, which is also true for Yahoo in 2016. This could mean that Yahoo used relatively ambiguous wording for headlines more favorable to Republicans. In 2016, the NYT was significantly to the left of Google at the 1% level for two specifications, and the WashPost and WSJ were both left of Google at 5% for two specifications. The magnitudes are on average largest for the WSJ and next largest for the NYT. The magnitudes of the significant effects are typically large, often close to 0.5 (again the total scale ranges from -1 to 1).

As discussed in Section 3, these results could be due to the outlets satisfying consumers' psychological demand for congenial news, or to appear more trustworthy to consumers with congenially biased priors on the electoral outcome. In the latter case, this pandering to readers' beliefs might decline as the election approached. The election outcome provides feedback on the accuracy of prior horse race reporting, which could increase the cost of pandering as the election approached. We examine this possibility by interacting the outlet dummies with the number of days until the election (multiplied by -1, to ease the interpretation of the coefficients). However, we largely fail to find significant interaction effects (Table B1). In fact, the few coefficients that are statistically significant—most of which are for the NYT—indicate the opposite effect; i.e., more congenial slant when the election is approaching. This pattern opposes the trust mechanism and further supports the psychology mechanism.

To estimate slant on the extensive margin, we construct outlet-level daily time series data sets. These data sets account for both days in which horse race stories were and were not reported. We run separate regressions for each outlet of the following form:

$$\#HR\,Stories_{jt} = \alpha + \beta Slant_t^{true} + f(\#HR\,Stories_{-j,t}) + \epsilon_{jt}. \tag{2}$$

The left-hand-side variable is the number of horse race stories reported by outlet j on day t. $Slant_t^{true}$ is a measure of "true" slant on day t, equal to either: 1) the average slant of stories reported by other outlets that day or 2) the pollstR average difference in polls 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. The term $f(\#HR\,Stories_{-j,t})$ is a control for the importance of horse race news on day t, a flexible polynomial of the total number of horse race stories reported by other outlets that day.[13] For models using other outlets' slant as the measure of $Slant_t^{true}$, we also include a 4th order date polynomial, to further control for general trends in horse race news interest. We do not include this polynomial when we use the pollstR measure of $Slant_t^{true}$ because these variables are highly collinear. We use Poisson regressions because the left-hand side is a count variable, with bootstrap standard errors; results are similar when we use OLS with Newey-West standard errors.

Results are presented in Table 3. For the Fox 2012 sample, one estimate is significant at 10%, and all of the estimates have signs consistent with congeniality slant. There are also two significant estimates consistent with congeniality for the NYT in 2016. These results provide some additional support for the psychology mechanism. The strongest and more robust results, however, are for the WSJ. Here, the majority of estimates in both years are significant at least at the 10%-level and indicate that the WSJ reported more horse race stories on days when

[13] The degree of the polynomial is arbitrary and results are similar for other flexible polynomials. We cannot use day fixed effects in these models because there is only one observation per day.

news was *less* favorable to Republicans. Assuming a demand-side explanation for this result and that the WSJ's readers are majority Republican, the result is most plausibly explained by this uncongenial news offering greater instrumental value. The WSJ's relatively affluent readers may have wanted to know when they most needed to write checks for the Republican candidate, which was when he was down in the polls. There is also evidence of a similar but less robust effect for the WashPost. However, these results might also have a supply-side explanation, more plausibly for the WSJ: their news staff could have been acting on their own left-leaning ideology, consistent with a finding of Groseclose and Milyo (2005).

5.2 Within-outlet news demand

We estimate the within-outlet relationship between variation in slant and story popularity with variants of the following linear probability model:

$$MostViewed_{jk} = \alpha + \beta_j Outlet_j + \gamma_j Outlet_j \times Slant_{ik} + \Sigma_t Day_t^{jk} + \epsilon_{jk}. \tag{3}$$

$MostViewed$ is a dummy for whether story k (ever) made its outlet's (j's) most viewed list, which we regress on outlet dummies (to account for different mean probabilities of making the most viewed list by outlet), interactions of a dummy for $Outlet_j$ and $Slant_{ik}$, and a fixed effect for *each* day that a story was available online ($Day_t^{jk} = 1$ if story jk was available on day t). The coefficients on the interactions (the γ_j's) are the estimates of interest, as each represents the marginal effect of slant on the probability of being most viewed for outlet j. 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 4. There is some evidence of a congeniality effect for Fox in 2012, but this is significant for $Slant_1$ only, and only at 5%. The magnitudes of the $Slant_1$ estimates imply that a one unit increase in slant predicts a 17 percentage point increase in being most viewed. There are no other significant results for that year. For 2016, there are several results significant at 10% or 5% for Fox and NYT, but each indicates that *less* congenial stories were the ones 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, 19-31 percentage points. 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 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 (see Table B2). In unreported results, we examine specifications in which we split out slant into congenial and uncongenial effects for each outlet, and find no systematic patterns.

Thus, we only obtain one significant result, for Fox in 2012, that can be uniquely attributed to the psychology mechanism. But even this result is non-robust. The somewhat more robust significant 2016 results for the NYT and Fox are most plausibly explained by the trust mechanism. Given that mean slant for these outlets was congenial, headlines with relatively uncongenial slants may have been more popular because readers expected the stories to be particularly trustworthy. If a like-minded outlet publishes a headline that is uncongenial to the outlet's average reader, this reader will assume that the underlying evidence must be particularly compelling.

As noted in Section 3, another factor that could drive reader interest is surprise, and it is possible that relatively uncongenial stories from these outlets were considered more surprising, and not necessarily more accurate. We investigate this possibility by constructing a variable intended to directly measure the degree to which the slant of a story was surprising for each outlet given trends and polls. We do this by, first, estimating a separate regression model for each outlet with $Slant_i$ on the left-hand side, explained by average pollstR levels and the previous week's changes in relative Republican support, as well as a third-order date polynomial. We use predicted values from these models to extract the surprise component in slant, by calculating the absolute value of the difference of actual slant of any horse race story reported and the predicted value of slant for that story. We then replace $Slant_i$ with the surprise component in the models of Table 4. This captures surprise in headline slant given the outlet's trend in slant and how the outlet typically reports on recent poll results. Results are insignificant, providing additional support for the trust mechanism interpretation (see Table B3).

To further address the role of trust, Tables B4 and B5 provide evidence on the relationship between horse race reporting and true horse race news. The tables show that outlet-slant interactions are highly predictive both of the daily poll average and its weekly change, even conditional on all other outlets' slant. We obtain very similar results when we run the regressions separately for each outlet-year. These results further support the importance of the trust mechanism: Readers can easily verify the accuracy of horse race stories, especially when the election is close. Outlets can build trust by providing unbiased reports, even if it means publishing uncongenial headlines.

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 horse news for Republicans was less congenial, and Fox horse race stories were indeed less congenial those days. Even if these stories were more likely to make the most viewed list than a horse race story on a more congenial news day, it is possible these uncongenial stories received fewer total clicks than the more congenial stories. We do not have daily total website click data to directly address this issue, but can use other publicly available data to shed some light on it. Figure A1 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 horse race 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.

6 Surveys

6.1 Design

When analyzing the web data, we use the outlet's ideological reputations to determine headline congeniality. While reputations are well-aligned with the ideologies of the outlets' typical readers, there is still ideological heterogeneity in the outlets' readerships. Unfortunately, our web data do not allow us to address heterogeneity since the data do not contain information on who is doing the clicking. It is possible that what we refer to as demand for uncongenial news was actually driven by clicks from an outlet's ideological minority. For example, the clicks on NYT stories favorable to Republicans may have mostly come from Republican readers. To address this issue, and complement the web data analysis more generally, we collect additional micro-level data from incentivized surveys.

An ideal experiment for generating micro data might consist of varying each news demand mechanism one at a time, holding fixed the news topic, source, and other relevant factors, and estimating the resulting effects on news demand. Implementing such an experiment with real news would be extremely difficult or perhaps impossible. We approximated such an experiment as follows. 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 on interest in debate news versus other news.[14] We used debate news, rather than horse race news, for the surveys because the timing of debates, and debate news, is known well in advance. The timing allowed us to prepare for posting the surveys shortly after the stories became available. Apart from that, debate news stories share several important similarities with horse race stories: 1) they both provide information on the candidates' chances of winning the election; 2) the partisan congeniality of stories on both topics is relatively clear; and 3) the partisan congeniality of stories tends to vary both across and within outlets. Debate stories differ

[14] See Appendix Table A3 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). A broad variety of experiments done on MTurk have tended to reach similar conclusions as those done on more representative samples (e.g., Mullinix et al., 2015). The two major advantages MTurk offered over a survey firm were: 1) MTurk is much more cost effective, allowing us to obtain a larger sample and 2) MTurk gave us control over the timing of surveys, which, as we explain, is crucial for their validity.

from those on the horse race in that debate stories are more likely to include voting-relevant information, such as the candidates' characteristics and policy position. Thus, debate stories might be have more instrumental value for rational decision-making. However, instrumental value is again more likely to come from *uncongenial* stories since they are more likely to cause the reader to change her decision of whom to vote for. This point is similar to that of Chan and Suen (2008) discussed in Section 1. Thus, we can again infer that a general demand for congenial news likely has non-instrumental motives.

The first and third 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 headline 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. The headlines are provided in Table A4. We use articles from just these three outlets because of their prominence and consistent availability in our web data, and to keep the choice set simple. 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. They disagreed on the third debate. Fox's headline said that Trump won, while the NYT's headline did not declare a winner.

Respondents were 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 which article they picked, and that they would receive an additional payment if they answered the question correctly. Respondents 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 notes to Table A4.

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. Perhaps the most unnatural element of our design was that respondents were presented with news choices from diverse outlets at the same time. This would be unrealistic for news consumers who, for example, go straight to the NYT website to decide which story to read, or have a Facebook feed consisting of ideologically similar sources. Thus, this element could bias observed demand for news from non-like-minded outlets upwards. It is also worth noting that the respondents might have already gotten their debate news prior to the survey. We tried to minimize this possibility by conducting the surveys fairly early in the morning after the debates. 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. Unreported tests indicate that the order did not have significant effects on article choice.

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 637 observations, with 345 identifying as Democrats or leaning toward Democrats, 177 identifying as Republican or leaning toward Republicans, and 115 as independent. Assuming congeniality of news remained neutral for non-partisan respondents, they constitute a quasi-control group that allows us to account for general changes in the importance of debate news over time and/or changes in the appeal of the non-debate news options.

The data and analysis do not map directly to those that we use for horse race news. We certainly cannot estimate average slant over a sample of headlines using the survey data, given the small number of headlines and debates.[15] Instead, we look primarily to see whether: 1)

[15] We did manually collect the number of debate-related links (both articles and videos) from web.archive.org on the Fox and NYT websites in the morning following each of the four debates of 2012 and 2016. Comparing the number of links to the verdicts for each debate allows us to informally examine slant at the extensive margin. These numbers are presented in Table A5. The table shows that the number of stories was fairly constant for both outlets in 2012. However, in 2016, there are indications of a quantitative bias towards congenial information. 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. Since Fox claimed Trump won the third debate, the NYT said the Republican (Pence) won the second, and Fox's headlines favored Clinton after the fourth debate (e.g., "Trump winning on points until terrible mistake"), both Fox's and the NYT's numbers

general interest in debate news increases when it is more congenial; 2) interest in debate news from the like-minded outlet increases when it is more congenial; or 3) interest in debate news from the non-like-minded outlet increases when it is more congenial. As discussed in Section 3, all of these results could be explained by the psychology mechanism. Results #1 and #3 could also be explained by trust, because subjects may perceive more congenial headlines as signals of higher credibility of article content. Result #2 is least plausibly due to trust, and thus finding result #2 and *not* finding result #3 would be the clearest support for the psychology mechanism. The reverse of any of these results—higher interest in uncongenial news—would be best explained by the instrumental mechanism.

6.2 Results

Figure 5 previews the econometric results. Across the three surveys, Democrats were least likely, and Republicans most likely, to pick "Other News" 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 more 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 similar but somewhat sharper when respondents are split out by preferred candidate rather than party (see Figure B1).

A multinomial model would be technically the best choice to formally analyze the survey data, since respondents chose among four unordered alternatives. However, in the interest of simplicity and transparency, we relegate the multinomial analysis to the appendix (see Table B6) and present here results of linear probability models predicting a binary outcome equal

of debate links are correlated with the congeniality of the debate outcome for their readers. The table also indicates that the slant of the outlets may have changed over time.

Figure 5: News choices by debate and party

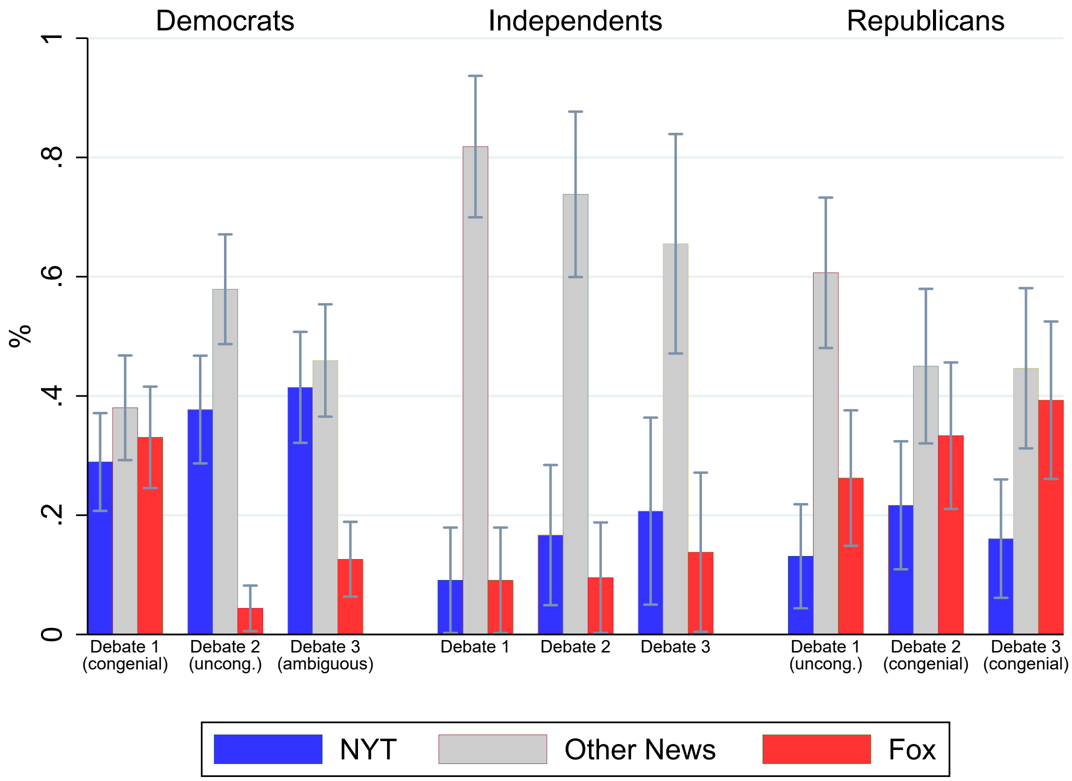

Note: 1) Both NYT and Fox survey 1 headlines said Clinton (Democrat) won first debate, 2) Both NYT and Fox survey 2 headlines said Pence (Republican, Trump's VP) won the second debate, and 3) Fox survey 3 headline said Trump won third debate while NYT survey 3 headline was ambiguous. The error bars denote 95% confidence intervals.

to: 1) Y_i^{NYT} (= 1 if respondent i chose the NYT article); 2) Y_i^{Fox} (defined analogously); 3) whether either type of debate news story is chosen ($Y_i^{debatenews} = Y_i^{NYT} + Y_i^{Fox}$). We run two sets of regressions, one using party identity as a measure of the respondent's politics and one using the respondent's preferred candidate for this, each with the following structure:

$$Y_i = \alpha + \beta_L D_i^L + \beta_R D_i^R + \beta_{S^2} S_i^2 + \beta_{S^3} S_i^3 + \beta_{L,2} D_i^L S_i^2 + \beta_{L,3} D_i^L S_i^3 + \beta_{R,2} D_i^R S_i^2 + \beta_{R,3} D_i^R S_i^3 + \beta_X X_i + \epsilon_i. \quad (4)$$

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) and 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.

Table 5 reports the results. Congeniality drives Democrat and Clinton supporter demand for Fox, but not for NYT news. Democrats were 20-30 percentage points less likely to get Fox news when it was uncongenial than when it was congenial ($\beta_{L,2} = -0.293$ and $\beta_{L,3} = -0.221$ for Y_i^{Fox}). There are no significant effects for Y_i^{NYT} for either Democrats or Clinton supporters, or for any dependent variable for Republicans. For Trump supporters, there is also evidence of a congeniality effect but it is primarily for Fox: Trump supporters were around 20 percentage points more likely to get Fox news in survey 2 as compared to survey 1 ($\beta_{R,2} = 0.194$).

As discussed at the end of Section 6.1, for L types, these results are consistent with either psychology and/or trust effects driving their demand for news from Fox. Democrats may enjoy this news more, or trust it more, when it is more congenial. It is unlikely Democrats found more congenial headlines more worth investigating for decisions. For Trump supporters, the results are most consistent with the psychology mechanism, since it is implausible that more pro-Republican Fox stories were perceived as more trustworthy (or more instrumentally useful).

We note also that while Trump supporters were consistently uninterested in NYT news, Democrats were just as willing to get news from the NYT when it is uncongenial. Thus, the findings support our interpretation of the 2016 web data result that NYT stories uncongenial to Democrats were more popular: the higher probability of these stories making the most viewed list was likely driven by clicks from Democrat readers (cp. Section 5.2, Table 4). In contrast, the survey results do not directly support the interpretation that uncongenial Fox stories were more popular because of clicks by Trump supporters. Still, substantial fractions of Trump supporters (24%) and Republicans (28%) were willing to click on the uncongenial Fox story, indicating their news demand was primarily driven by the desire for information.

7 Concluding remarks

We study slant in headlines of horse race stories and news on the winners of presidential debates. Investigating differences in slant across news outlets—and selective exposure within—allows us to better understand the causes of demand for slant. We find numerous results that we interpret as supporting either the trust or psychology mechanisms, and some that best support just one or the other. The evidence for the psychology mechanism is stronger for right-of-center consumers. This kind of asymmetry is not implausible, as unequal psychological effects across the ideological left-right spectrum are common (e.g., Jost, 2009). We also find some evidence arguably consistent with the instrumental mechanism being important for readers of the WSJ.

While trust and psychology are modeled as distinct factors in some theory papers, they are clearly related. Trust can be influenced by psychological factors. Cognitive dissonance and confirmation bias might affect reactions to information observed in the past. These factors may be the primary reasons that perceptions of news source credibility vary so much across the population (Kelly, 2018). A more subtle point that our results on the reasonably robust demand for "bad news" across the spectrum help us to see is that the psychological value of news also depends on trust. As in the case of self-signaling models that incorporate credibility (Bénabou and Tirole, 2016), "good news" does not make one feel good unless one thinks this good news is believable. This factor implies possibly natural limits to the degree of psychology-driven news distortion: if distortion becomes too extreme, it destroys credibility, defeating this purpose. This trade-off is likely most relevant to issues reported on repeatedly, such as horse race news.

Thus, our results imply that psychology forces, broadly defined, were responsible for horse race reporting being objectively skewed for several outlets. Consequently, 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 have been harmful to social welfare. For example, biased horse race reporting and news consumption could lead to distrust of election results and conspiracy theories (Hollander, 2014). Moreover, slant on other topics, with less frequent and clear feedback, is likely more severe.

Future work building on our paper could consider more detailed individual click-level web data or social media data. The complexity of our results 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 moved further left. The WashPost and NYT have similar ideological reputations, but there are differences in their reporting and reader behavior. The WSJ's reporting appears quite distinct from its reputation, as opposed to that of Fox and the NYT. These nuances in our findings emphasize 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.

References

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

R. Bénabou and J. Tirole. Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–64, 2016.

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. Bruns and O. Himmler. Mass media, instrumental information, and electoral accountability. *Journal of Public Economics*, 134:75–84, 2016.

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. 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. Garz, J. Sörensen, and D. F. Stone. Partisan selective engagement: Evidence from Facebook. Working paper, 2018.

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.

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. *Quarterly Journal of Economics*, 120 (4):1191–1237, 2005.

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.

W. Hart, D. Albarracín, A. H. Eagly, I. Brechan, M. J. Lindberg, and L. Merrill. Feeling validated versus being correct: A meta-analysis of selective exposure to information. *Psychological bulletin*, 135(4):555–588, 2009.

B. A. Hollander. The surprised loser: The role of electoral expectations and news media exposure in satisfaction with democracy. *Journalism & Mass Communication Quarterly*, 91 (4):651–668, 2014.

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.

J. T. Jost. "Elective affinities": On the psychological bases of left-right differences. *Psychological Inquiry*, 20(2-3):129–141, 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.

D. Kelly. Evaluating the news: (Mis)perceptions of objectivity and credibility. *Political Behavior*, forthcoming, 2018.

G. J. Martin and A. Yurukoglu. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–99, 2017.

Y. Masatlioglu, A. Y. Orhun, and C. Raymond. Intrinsic information preferences and skewness. Working paper, 2017.

M. J. Metzger, E. H. Hartsell, and A. J. Flanagin. Cognitive dissonance or credibility? A comparison of two theoretical explanations for selective exposure to partisan news. *Communication Research*, forthcoming, 2015.

S. Mullainathan and A. Shleifer. The market for news. *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.

E. Peterson, S. Goel, and S. Iyengar. Echo chambers and partisan polarization: Evidence from the 2016 presidential campaign. Working paper, 2018.

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

K. Pogorelskiy and M. Shum. News we like to share: How news sharing on social networks influences voting outcomes. Working paper, 2018.

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.

A. Simonov and J. Rao. What drives demand for government-controlled news in russia? Working paper, 2018.

D. Stiers and R. Dassonneville. Affect versus cognition: Wishful thinking on election day: An analysis using exit poll data from Belgium. *International Journal of Forecasting*, 34(2): 199–215, 2018.

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):577–585, 1985.

Table 1: Average 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
USAT	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

Table 2: Estimated mean differences in slant across outlets

Outlet	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012 (Reference outlet = Yahoo)			
Fox	0.340**	0.471***	0.648***
	(0.130)	(0.125)	(0.176)
WSJ	-0.306*	-0.165	0.059
	(0.176)	(0.170)	(0.227)
USAT	-0.117	0.012	0.037
	(0.182)	(0.175)	(0.245)
NYT	-0.312**	-0.161	-0.142
	(0.121)	(0.131)	(0.206)
HuffPost	-0.499***	-0.335***	-0.236
	(0.108)	(0.114)	(0.178)
Adj R^2	0.412	0.407	0.494
N	425	382	277
Panel A: 2012 (Reference outlet = Google)			
Fox	0.238**	0.308**	0.353*
	(0.116)	(0.149)	(0.191)
WSJ	-0.235	-0.506***	-0.479**
	(0.179)	(0.171)	(0.203)
Yahoo	-0.082	-0.143	-0.360**
	(0.130)	(0.150)	(0.149)
NYT	-0.363***	-0.230*	-0.411***
	(0.097)	(0.120)	(0.120)
WashPost	-0.162**	-0.118	-0.260**
	(0.079)	(0.092)	(0.120)
Adj R^2	0.250	0.332	0.434
N	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: $Slant_i$. All models include fixed effects for the first date a story was reported. 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 3: Effects of slant and poll means on the daily number of horse race stories

	Fox	WSJ	NYT	HuffPost
Panel A1: RHS = mean slant of other outlets (2012)				
$Slant_1$	0.251	-0.038	0.284	-0.075
	(0.253)	(0.504)	(0.224)	(0.265)
$Slant_2$	0.309	-0.740	0.182	0.014
	(0.258)	(0.665)	(0.221)	(0.210)
$Slant_3$	0.202	-0.624	0.224	-0.335
	(0.267)	(0.647)	(0.207)	(0.207)
Panel A2: RHS = mean Republican poll advantage (2012)				
$Slant_1$	0.077	-0.584**	-0.067	0.051
	(0.089)	(0.292)	(0.091)	(0.089)
$Slant_2$	0.135	-0.437*	-0.027	0.037
	(0.098)	(0.228)	(0.096)	(0.084)
$Slant_3$	0.186*	-0.447**	-0.040	-0.006
	(0.112)	(0.193)	(0.093)	(0.103)

	Fox	WSJ	NYT	WashPost
Panel B1: RHS = mean slant of other outlets (2016)				
$Slant_1$	-0.091	-0.417	-0.367	0.224
	(0.269)	(0.781)	(0.321)	(0.212)
$Slant_2$	-0.198	-1.078*	-0.787***	0.216
	(0.350)	(0.635)	(0.291)	(0.247)
$Slant_3$	-0.146	-1.206	-0.609**	0.233
	(0.315)	(0.773)	(0.273)	(0.278)
Panel B2: RHS = mean Republican poll advantage (2016)				
$Slant_1$	0.204**	-0.335*	-0.155	0.107**
	(0.087)	(0.175)	(0.105)	(0.050)
$Slant_2$	0.134	-0.363	-0.156	0.122*
	(0.085)	(0.229)	(0.118)	(0.069)
$Slant_3$	0.091	-0.341**	-0.128	0.076
	(0.112)	(0.165)	(0.140)	(0.087)

Note: Poisson regressions, using daily outlet-level time series, with bootstrap standard errors. 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 include a 4th order date polynomial. N=103 and 101 for Panels A1 and A2, and N=105 for Panels B1 and B2. *, **, *** denote 10%, 5%, 1% significance.

Table 4: Probability of story 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	0.172**	0.079	0.100
	(0.083)	(0.103)	(0.135)	(0.085)	(0.104)	(0.143)
WSJ $\times Slant$	-0.010	0.024	0.001	-0.042	-0.003	-0.055
	(0.202)	(0.210)	(0.186)	(0.195)	(0.201)	(0.180)
USAT $\times Slant$	-0.106	-0.087	-0.088	-0.110	-0.096	-0.118
	(0.088)	(0.115)	(0.111)	(0.091)	(0.120)	(0.114)
Yahoo $\times Slant$	-0.043	-0.038	-0.027	-0.027	-0.023	-0.027
	(0.043)	(0.053)	(0.064)	(0.040)	(0.047)	(0.056)
NYT $\times Slant$	-0.013	0.024	0.026	0.001	0.037	0.032
	(0.072)	(0.076)	(0.087)	(0.071)	(0.076)	(0.085)
HuffPost $\times Slant$	0.010	0.029	0.021	0.020	0.035	0.009
	(0.048)	(0.054)	(0.068)	(0.052)	(0.058)	(0.077)
Competing headline controls?				✓	✓	✓
Adj. R^2	0.379	0.408	0.353	0.380	0.410	0.340
N	425	382	277	425	382	277
Panel B: 2016						
Fox $\times Slant$	-0.131*	-0.189**	-0.179*	-0.128*	-0.171*	-0.176
	(0.069)	(0.087)	(0.094)	(0.075)	(0.099)	(0.108)
WSJ $\times Slant$	-0.043	-0.187	0.389	-0.029	-0.091	0.463
	(0.244)	(0.523)	(0.293)	(0.247)	(0.546)	(0.285)
Yahoo $\times Slant$	-0.021	-0.009	0.089	-0.034	-0.061	0.022
	(0.087)	(0.116)	(0.141)	(0.093)	(0.137)	(0.209)
Google $\times Slant$	0.016	-0.026	0.009	0.023	-0.018	0.018
	(0.029)	(0.029)	(0.032)	(0.030)	(0.032)	(0.032)
NYT $\times Slant$	0.193*	0.284**	0.511	0.194*	0.314**	0.539
	(0.111)	(0.130)	(0.324)	(0.110)	(0.129)	(0.330)
WashPost $\times Slant$	-0.043	-0.013	0.019	-0.048	-0.046	-0.003
	(0.056)	(0.072)	(0.082)	(0.057)	(0.073)	(0.093)
Competing headline controls?				✓	✓	✓
Adj. R^2	0.286	0.378	0.371	0.286	0.392	0.371
N	696	501	381	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: most viewed (yes/no). All models include the constituent terms of the interactions and 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). Competing headlines controls are fixed effects for number of headlines with $Slant_1$ value for the outlet on first day that story was reported. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table 5: Likelihood of picking NYT and Fox debate winner article

	Y^{NYT}	Y^{Fox}	$Y^{NYT} + Y^{Fox}$
Panel A (Party Identity)			
Debate 2 × Democrat ($\beta_{L,2}$, uncongenial)	0.004	-0.293***	-0.289***
	(0.095)	(0.079)	(0.111)
Debate 3 × Democrat ($\beta_{L,3}$, uncongenial for Fox only)	-0.023	-0.221**	-0.243*
	(0.110)	(0.100)	(0.126)
Debate 2 × Republican ($\beta_{R,2}$, congenial)	0.009	0.071	0.079
	(0.101)	(0.104)	(0.128)
Debate 3 × Republican ($\beta_{R,3}$, congenial)	-0.095	0.112	0.016
	(0.113)	(0.121)	(0.140)
Adj. R^2	0.059	0.085	0.062
N	637	637	637
Panel B (Candidate supported)			
Debate 2 × Clinton supporter ($\beta_{L,2}$, uncongenial)	0.039	-0.274***	-0.235**
	(0.093)	(0.075)	(0.106)
Debate 3 × Clinton supporter ($\beta_{L,3}$, uncongenial Fox only)	0.122	-0.270***	-0.148
	(0.098)	(0.094)	(0.116)
Debate 2 × Trump supporter ($\beta_{R,2}$, congenial)	0.056	0.194*	0.250**
	(0.100)	(0.103)	(0.122)
Debate 3 × Trump supporter ($\beta_{R,3}$, congenial)	-0.069	0.186	0.117
	(0.099)	(0.117)	(0.132)
Adj. R^2	0.078	0.123	0.091
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. The reference category is debate 1. The reports on this debate were congenial for Democrats/Clinton supporters and uncongenial for Republicans/Trump supporters. *, **, *** denote 10%, 5%, 1% significance.

A Appendix

A.1 Model

We present a slightly modified, and simplified, version of the model of Mullainathan and Shleifer (2005), using their notation to the extent possible. There is a single news outlet that repeatedly reports stories on the horse race. Individual news story i, n_i, is the sum of three components: $n_i = d_i + s + \epsilon_i$. d_i denotes "true" news (data) on the horse race, s is a constant that is strategically chosen by the outlet (slant) and applies to all of the outlet's stories, and ϵ_i is exogenous mean-zero story-specific noise. Each component has support $[-b, b]$ for some $b > 0$. We henceforth suppress the i subscript.

There is one representative consumer whose utility from reading n is

$$U_r = \bar{u} - \chi(s + \epsilon)^2 + \phi n - c, \tag{5}$$

with $\bar{u} > 0$, $\chi \geq 0$, $\phi \geq 0$, and c a (story-level) random variable with support $[0, \infty)$. The consumer cares about truth and loses utility from the story's slant and noise $(s + \epsilon)$ if $\chi > 0$. If $\phi > 0$, then the consumer gains utility from n being larger because this is more congenial (n is news about the preferred candidate's chances in the election).

Before reading the story, the consumer must click on it. She decides whether to click by first reading the story's headline, h. She then clicks on the story to read it if $E(U_r|h) \geq 0$ (her reservation utility is normalized to zero). She knows s and c, and uses the headline to draw inferences on d and ϵ. We assume that $h \in \mathbb{R}$ is an unbiased signal of n, and that the consumer's updated expectations, $E(d|h)$ and $E(\epsilon|h)$, are unbiased and both $E(d|h)$ and $E(\epsilon|h)$ are strictly increasing in h. All expectations are implicitly conditioned on other available information (the date, recent stories by the outlet, etc).

The outlet maximizes clicks by choosing s to maximize the consumer's expectation of $-\chi(s+\epsilon)^2 + \phi n = -\chi(s+\epsilon)^2 + \phi(d+s+\epsilon)$ across stories, which is done with $s^* = \phi/2\chi$ for an interior solution. The ratio ϕ/χ measures the consumer's preference for congeniality relative to truthful reporting. If $\phi = 0$ and $\chi > 0$, then the consumer only cares about minimizing slant and does not care about congeniality; if $\chi = 0$ and $\phi > 0$, then the consumer only cares about congeniality, and $s^* = b$. Let $\hat{s}$ denote the outlet's choice of s.

Consequently, $E(U_r|h) = -\chi E\left((\hat{s}+\epsilon)^2|h\right) + \phi\left(E(d|h) + \hat{s} + E(\epsilon|h)\right) - c$. The marginal effect of h is $-\chi \frac{\partial}{\partial h} E\left((\hat{s}+\epsilon)^2|h\right) + \phi \frac{\partial}{\partial h}\left(E(d|h) + E(\epsilon|h)\right)$. This expression is decreasing in $\hat{s}$ and thus the overall marginal effect is negative for a larger range of h for larger $\hat{s}$ if $\chi > 0$.[16] When the marginal effect of h on $E(U_r|h)$ is negative, then news demand (click probability) is decreasing as h grows; when the marginal effect is positive, demand increases in h. In other words, whether demand is increasing or decreasing in the congeniality of the headline (h) depends partly on average slant ($\hat{s}$). Intuitively, when $\hat{s}$ is very large, readers are getting too much slant on average, so they are more likely to click headlines that seem relatively unslanted. When $\hat{s}$ is not so large, demand is increasing in h because this is a signal of larger d and also perhaps because the trade-off between costs and benefits of ϵ was not optimized. The marginal effect cannot be evaluated at $\hat{s} = s^*$ without further assumptions, but we conjecture that this would be positive but relatively small for reasonable assumptions and parameter values, since the outlet is on average balancing the trade-off between slant and truth optimally.

Note that the changes to Mullainathan and Shleifer (2005)'s model are as follows. One is very minor, setting price to zero. Another is also minor: rather than the reader preferring news to be as close as possible to a particular scalar (b in their notation), we assume that "more is better" for the congeniality of content. This assumption has no effect on results and simply feels more natural in our context. The addition of an unobserved noise term, ϵ, to news content, and the addition of the headline being seen before a story is clicked on, are more substantive, but do not affect the key qualitative result about there being a trade-off between congenial slant and accuracy.

[16]Note $-\chi \frac{\partial}{\partial h} E\left((\hat{s}+\epsilon)^2|h\right) = -\chi \frac{\partial}{\partial h} E\left(\hat{s}^2 + 2\hat{s}\epsilon + \epsilon^2|h\right) = -\chi \frac{\partial}{\partial h} E\left(2\hat{s}\epsilon + \epsilon^2)|h\right)$, so the derivative of this expression with respect to $\hat{s}$ is $-2\chi \frac{\partial}{\partial h} E\left(\epsilon|h\right)$, which is strictly negative since $\frac{\partial}{\partial h} E(\epsilon|h) > 0$.

A.2 Additional figures and tables

Figure A1: Google Trends data (day 0 = election day)

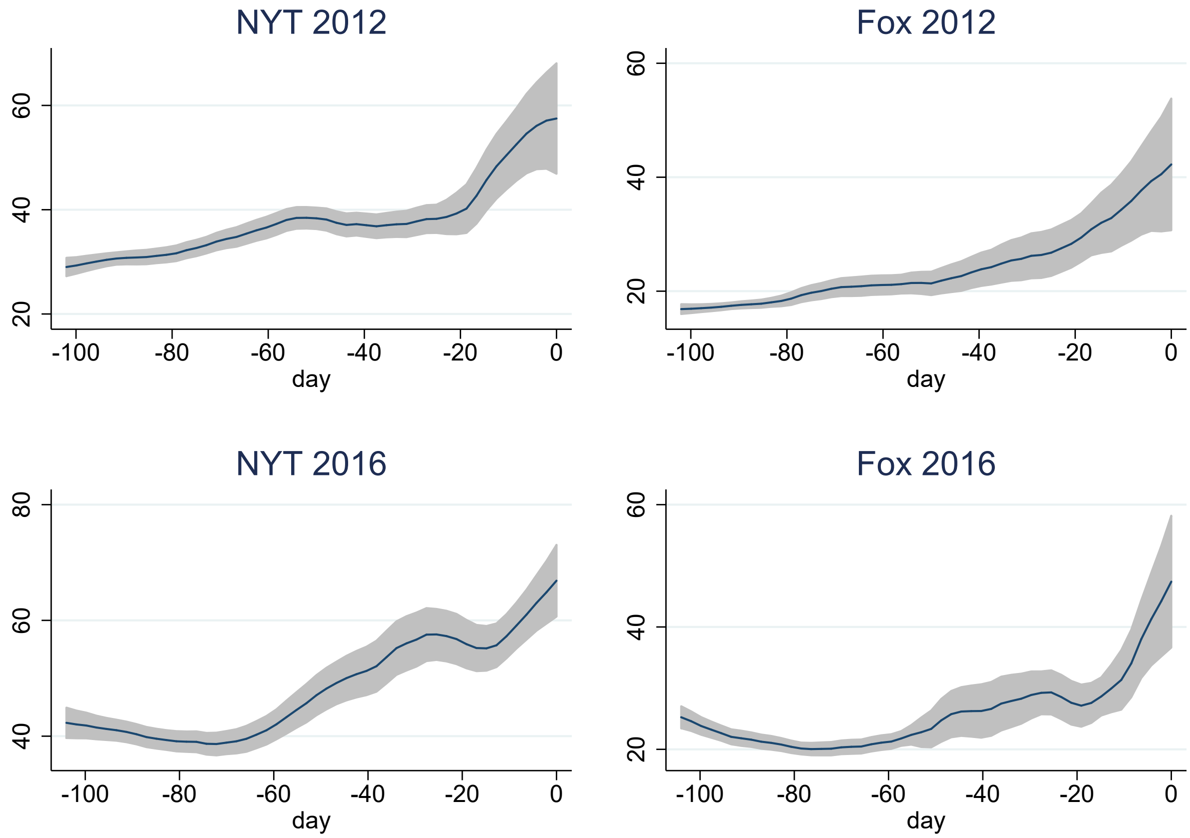

Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. “NYT” = Google searches for “new york times”; “Fox” = Google searches for “fox news”.

Table A1: Instructions for rating slant of horse race headlines

Category	Definition
Very good news for Clinton	"Very good news" about Clinton's 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 Clinton's 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 candidate's chance
Good news for Trump	"Good news" about Trump's 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 Trump's 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 A2: *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? fbi's 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 A3: 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 and gender (N=637). Respondents are master MTurkers who answered reading check question correctly.

Table A4: Headline options for each survey

Survey 1:	NYT: Commentators Give Hillary Clinton Edge in Debate
	Fox: Hillary won the first debate (it helps to be prepared)
	Yahoo: Long dog-gone trip: Florida pooch travels to Boston and back
	Yahoo: Houston gunman had two weapons, thousands of rounds at scene.
Survey 2:	NYT: Who Won the Debate? Commentators Give Edge to Mike Pence
	Fox: Pence triumphs in VP debate. And then there was the night's biggest loser...
	Yahoo: 2 Vermont teachers accused of vandalizing sidewalk
	Yahoo: Two young girls shot in Cleveland drive-by shooting
Survey 3:	NYT: Who Won the Debate? Donald Trump Avoids Annihilation
	Fox: Trump comes out swinging and wins second debate
	Yahoo: Three police officers shot in Palm Springs, California
	Yahoo: 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 article's content. Your payment will be $0.25 higher if you answer the question correctly. The question's 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 A5: 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.

B Supplementary appendix

Figure B1: News choices by debate and preferred candidate

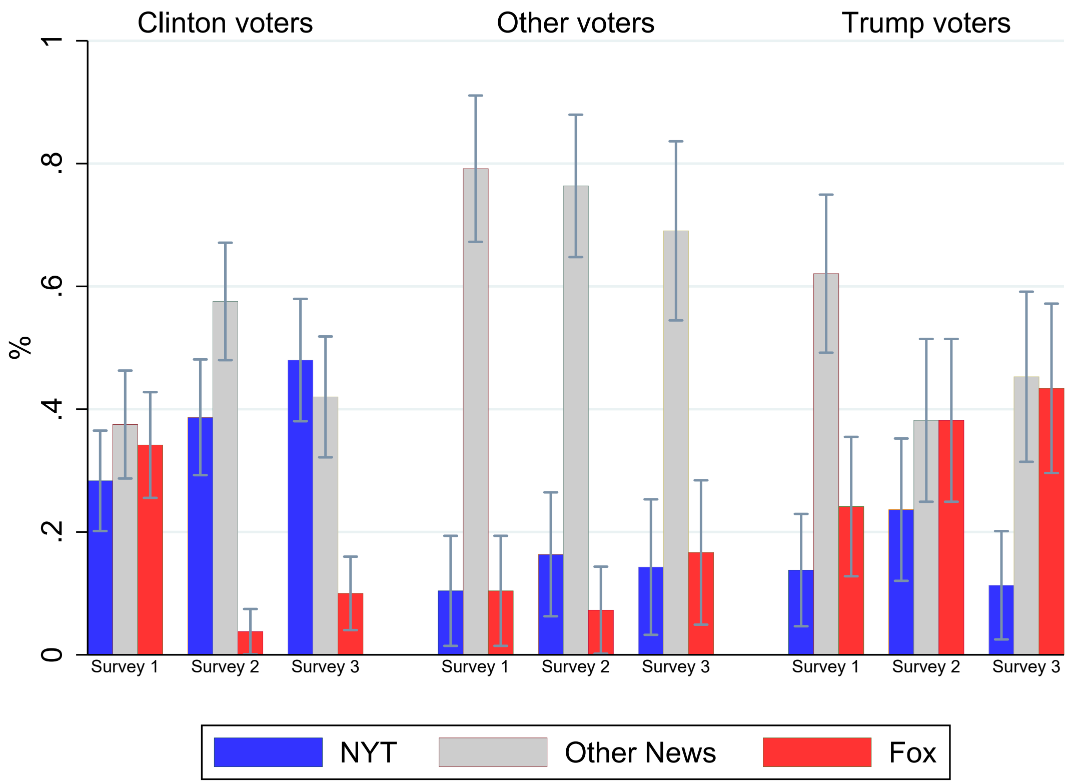

Note: 1) Both NYT and Fox survey 1 headlines said Clinton (Democrat) won first debate, 2) Both NYT and Fox survey 2 headlines said Pence (Republican, Trump's VP) won the second debate, and 3) Fox survey 3 headline said Trump won third debate while NYT survey 3 headline was ambiguous. The error bars denote 95% confidence intervals.

Table B1: Mean slant and outlet date interactions

Outlet	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012			
Fox × time to election	0.004	0.005	-0.010
	(0.006)	(0.006)	(0.008)
WSJ × time to election	0.008	0.002	-0.003
	(0.006)	(0.007)	(0.008)
USAT × time to election	0.003	-0.007	-0.024*
	(0.011)	(0.009)	(0.012)
NYT × time to election	-0.008*	-0.007	-0.017**
	(0.005)	(0.005)	(0.008)
HuffPost × time to election	-0.002	-0.004	-0.013
	(0.004)	(0.005)	(0.008)
Adj. R^2	0.415	0.410	0.509
N	425	382	277
Panel B: 2016			
Fox × time to election	-0.005	-0.009**	-0.008
	(0.003)	(0.004)	(0.006)
WSJ × time to election	-0.007	-0.003	-0.005
	(0.005)	(0.005)	(0.006)
Yahoo × time to election	-0.008	-0.007	-0.008
	(0.007)	(0.008)	(0.009)
NYT × time to election	-0.005*	-0.006**	-0.007***
	(0.003)	(0.003)	(0.002)
WashPost × time to election	-0.002	-0.004*	-0.005*
	(0.002)	(0.003)	(0.003)
Adj. R^2	0.251	0.338	0.436
N	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: $Slant_i$. The reference category is Yahoo in 2012 and Google in 2016. All models include fixed effects for the first date a story was reported and the constituent terms of the interacted variables. 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 B2: Probability of story making most viewed list (date polynomial instead of day fixed effects)

	$Slant_1$	$Slant_2$	$Slant_3$	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012						
Fox $\times Slant$	0.029	-0.015	-0.020	0.013	-0.030	-0.030
	(0.104)	(0.113)	(0.106)	(0.101)	(0.112)	(0.110)
WSJ $\times Slant$	-0.117	-0.114	-0.142	-0.125	-0.123	-0.155
	(0.203)	(0.205)	(0.174)	(0.199)	(0.199)	(0.171)
USAT $\times Slant$	-0.080	-0.009	-0.020	-0.076	-0.005	-0.016
	(0.080)	(0.099)	(0.094)	(0.081)	(0.097)	(0.092)
Yahoo $\times Slant$	-0.023	0.007	-0.001	-0.003	0.031	0.022
	(0.037)	(0.046)	(0.052)	(0.035)	(0.042)	(0.049)
NYT $\times Slant$	-0.057	-0.032	-0.039	-0.041	-0.013	-0.018
	(0.062)	(0.070)	(0.070)	(0.061)	(0.069)	(0.069)
HuffPost $\times Slant$	-0.002	0.022	0.015	0.011	0.031	0.023
	(0.035)	(0.039)	(0.044)	(0.035)	(0.040)	(0.045)
Competing headline controls?				✓	✓	✓
Adj. R^2	0.259	0.271	0.231	0.262	0.274	0.231
N	425	382	277	425	382	277
Panel B: 2016						
Fox $\times Slant$	-0.124**	-0.167***	-0.159**	-0.122**	-0.151**	-0.150**
	(0.060)	(0.060)	(0.061)	(0.060)	(0.059)	(0.061)
WSJ $\times Slant$	-0.033	-0.099	0.676***	-0.040	-0.132	0.627***
	(0.174)	(0.609)	(0.187)	(0.185)	(0.598)	(0.218)
Yahoo $\times Slant$	-0.026	0.066	0.111	-0.039	0.037	0.084
	(0.079)	(0.100)	(0.125)	(0.079)	(0.104)	(0.131)
Google $\times Slant$	0.018	-0.005	-0.005	0.037	0.014	0.014
	(0.028)	(0.031)	(0.032)	(0.026)	(0.030)	(0.032)
NYT $\times Slant$	0.212*	0.328**	0.259	0.196*	0.321**	0.248
	(0.119)	(0.146)	(0.178)	(0.118)	(0.145)	(0.175)
WashPost $\times Slant$	-0.017	0.027	0.025	-0.027	0.011	0.022
	(0.050)	(0.068)	(0.072)	(0.052)	(0.069)	(0.076)
Competing headline controls?				✓	✓	✓
Adj. R^2	0.194	0.248	0.224	0.204	0.267	0.233
N	696	501	381	696	501	381

Note: OLS estimates, using story-level data. Left-hand side variable: most viewed (yes/no). All models include the constituent terms of the interactions and a 4th order date polynomial. The 2016 models also include dummies for Yahoo stories' first date occurring during one of two time-frames in which Yahoo data collection changed. Competing headlines controls are fixed effects for number of headlines with $Slant_1$ value for the outlet on first day that story was reported. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table B3: Probability of story making most viewed list and surprise slant

	$Slant_1$	$Slant_2$	$Slant_3$	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012						
Fox $\times Slant_i^{surprise}$	-0.1395	-0.0788	-0.2320	-0.1184	0.0073	-0.0934
	(0.1510)	(0.1678)	(0.2178)	(0.1498)	(0.1816)	(0.2719)
WSJ $\times Slant_i^{surprise}$	0.1823	0.2658	-0.0544	0.0922	0.2048	-0.1173
	(0.8229)	(0.8772)	(0.6059)	(0.9030)	(0.9702)	(0.6861)
USAT $\times Slant_i^{surprise}$	-0.1052	-0.1456	-0.2634	-0.0990	-0.2005	-0.3141
	(0.3127)	(0.3321)	(0.3486)	(0.3159)	(0.3325)	(0.3137)
Yahoo $\times Slant_i^{surprise}$	-0.0643	-0.0313	-0.0934	-0.0477	0.0124	-0.0594
	(0.1099)	(0.1388)	(0.1354)	(0.1043)	(0.1271)	(0.1326)
NYT $\times Slant_i^{surprise}$	-0.1244	-0.2302*	-0.2748	-0.0927	-0.1898	-0.2474
	(0.1332)	(0.1199)	(0.1831)	(0.1322)	(0.1144)	(0.1843)
HuffPost $\times Slant_i^{surprise}$	0.0285	0.0918	0.0815	0.0713	0.1281	0.1058
	(0.0965)	(0.1029)	(0.1623)	(0.0979)	(0.1068)	(0.1719)
Competing headline controls?				✓	✓	✓
Adj. R^2	0.382	0.420	0.366	0.393	0.430	0.356
N	400	363	267	400	363	267
Panel B: 2016						
Fox $\times Slant_i^{surprise}$	0.1308	0.2368	0.2567	0.1090	0.1945	0.1820
	(0.1331)	(0.1877)	(0.1945)	(0.1424)	(0.2065)	(0.2016)
WSJ $\times Slant_i^{surprise}$	0.0549	0.3590	0.4298	0.0856	0.4533	0.5158
	(0.4123)	(0.8140)	(0.4650)	(0.4210)	(0.8240)	(0.4630)
Yahoo $\times Slant_i^{surprise}$	0.1131	0.3316*	0.3998	0.0820	0.3101	0.4132
	(0.1225)	(0.1748)	(0.2897)	(0.1288)	(0.1939)	(0.3415)
Google $\times Slant_i^{surprise}$	0.0117	-0.0100	-0.0282	0.0098	-0.0028	-0.0164
	(0.0656)	(0.0521)	(0.0595)	(0.0711)	(0.0525)	(0.0628)
NYT $\times Slant_i^{surprise}$	-0.0596	0.0610	0.2954	-0.0996	0.0679	0.3409
	(0.2426)	(0.3186)	(0.4757)	(0.2490)	(0.3154)	(0.5073)
WashPost $\times Slant_i^{surprise}$	-0.0781	-0.1304	-0.1324	-0.0927	-0.1488	-0.1958
	(0.1036)	(0.1300)	(0.1497)	(0.1090)	(0.1549)	(0.1733)
Competing headline controls?				✓	✓	✓
Adj. R^2	0.277	0.384	0.360	0.277	0.388	0.360
N	696	471	381	696	471	381

Note: OLS estimates, using story-level data. Left-hand side variable: most viewed (yes/no). The surpise measures are defined as $Slant_i^{surprise} = |Slant_i - \widehat{Slant_i}|$, where $\widehat{Slant_i}$ are predicted values from regressing slant on the levels and changes in the pollstR average difference in polls, and a third-order date polynomial. All models include the constituent terms of the interactions and 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. Competing headlines controls are fixed effects for number of headlines with $Slant_1$ value for the outlet on first day that story was reported. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table B4: Outlet-slant interactions and 2012 polls

	Dependent variable: pollstR poll average			
	level	weekly change	level	weekly change
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)
USAT $\times Slant_1$	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)
USAT $\times Slant_2$	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)
USAT $\times Slant_3$	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
3rd order date polynomial			✓	✓

Note: OLS estimates, using story-level data. The pollstR data refer to the poll average (i.e., percent planning to vote Republican minus percent planning to vote Democrat) on the day the story was published. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table B5: Outlet-slant interactions and 2016 polls

	Dependent variable: pollstR poll average			
	level	weekly change	level	weekly change
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)
WashPost $\times Slant_2$	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)
WSJ $\times Slant_3$	0.633*	0.283	0.501*	0.094
	(0.356)	(0.196)	(0.265)	(0.186)
Yahoo $\times Slant_3$	0.834***	0.355	0.429**	-0.081
	(0.180)	(0.235)	(0.169)	(0.128)
Google $\times Slant_3$	0.807***	0.326**	0.460***	-0.061
	(0.152)	(0.153)	(0.109)	(0.073)
NYT $\times Slant_3$	0.673**	0.365	0.492***	0.210
	(0.288)	(0.312)	(0.138)	(0.180)
WashPost $\times Slant_3$	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
3rd order date polynomial			✓	✓

Note: OLS estimates, using story-level data. The pollstR data refer to the poll average (i.e., percent planning to vote Republican minus percent planning to vote Democrat) on the day the story was published. Standard errors are clustered by the first date the story was available. *, **, *** denote 10%, 5%, 1% significance.

Table B6: 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.