

The supply of media slant across outlets and demand for slant within outlets: Evidence from US presidential campaign news ^{*}

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

March 18, 2020

Abstract

We conduct across-outlet and within-outlet (and within-topic) analyses of “congenially” slanted news. We study “horse race” news (news on candidates’ chances in an upcoming election) from six major online outlets for the 2012 and 2016 US presidential campaigns. We find robust evidence that horse race headlines were slanted congenially with respect to the preferences of the outlets’ typical readers. However, evidence of congenial slant in the timing and frequency of horse race stories is weaker. We also find very limited evidence of greater within-outlet demand for headlines most congenial to outlets’ typical readers, and somewhat stronger evidence of greater demand for relatively *uncongenial* headlines. We discuss how various aspects of our results are consistent with each of the major mechanisms driving slant studied in the theoretical literature, and may help explain when each mechanism is more likely to come into play. In particular, readers may be more likely to click on uncongenial headlines due to inferring that these stories are particularly informative when they stand in contrast to an outlet’s typically congenial slant.

^{*}We thank Matthew Gentzkow, Jesse Shapiro, Greg Martin, Matt Botsch, Ben Ho, and participants at the 2018 Economics of Media Bias Workshop, 2018 European Public Choice Society, 2018 Behavioral Models of Politics conference, 2018 Maine Economics Conference, the NYC Media Seminar, and a seminar at the City College of New York for very helpful comments, and thank Isaiah West for excellent research assistance. This research was funded by a Bowdoin College Grua/O’Connell Research Award. Declarations of interest: none.

[†]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 can be reached at: jtw95@uw.edu.

1 Introduction

Over the last decade, the economics literature has made substantial contributions to the debate about the causes and consequences of political media slant (Gentzkow et al., 2015; Puglisi and Snyder Jr, 2015a). This literature has focused on studying political slant *across* news outlets, and effects of such slant.¹ However, the ways in which news is slanted, and variation in slant across outlets, issues, and over time, are, of course, still not fully understood. Furthermore, the investigation of *within*-outlet demand for slant has been largely neglected so far. It is unclear whether, for a given trusted outlet, readers prefer more politically “congenial” news—news favoring one’s preferred policies or party—or if consumers are indifferent to the slant of individual stories from a preferred outlet.²

This topic is important because it relates directly to the question of how well the market for news contributes to a well-informed citizenry. Even if consumers primarily get news from ideologically aligned outlets, and even if such outlets distort their news to favor one side of the spectrum, consumers may still be reasonably well informed if they read a variety of stories, including “uncongenial” stories, from such outlets. On the other hand, if consumers restrict attention to congenial news within their favored outlets, then the echo chambers problem, and concerns about voters lacking information, would be more severe.

In this paper, we contribute to filling this gap in the literature. We construct a novel data set and use it to analyze both across-outlet slant and the within-outlet relationship between news slant and demand. Since both slant and consumer demand are correlated with news topics, it is crucial that our analysis also holds the news topic fixed. We do this by identifying a

¹Studies of variation in slant across outlets include Chiang and Knight (2011), Puglisi and Snyder Jr (2015b), and Casas et al. (2016), who use explicit ideological standpoints to measure slant across newspapers, whereas Gentzkow and Shapiro (2010), Greenstein and Zhu (2012), and Garz et al. (2019) measure the slant of media outlets by comparing their language with that of ideological reference texts. A related approach uses citation counts of political think tanks (Groseclose and Milyo, 2005) and core topics (Larcinese et al., 2011; Qin et al., 2018; Gambaro and Puglisi, 2015) to capture outlet-level differences in slant. See, e.g., Dewenter et al. (2019) for a recent example of a paper studying media effects.

²Note that the term congeniality is distinct from the more general term slant in that congeniality refers only to a type of slant that is favorable to a media consumer with a particular political orientation. Congenial news and uncongenial news are standard terms in the political science and communications literatures on media effects.

news issue reported on repeatedly by major outlets across the political spectrum: “horse race” news (news on competing candidates’ chances of winning an upcoming election). We study horse race news for the 2012 and 2016 US presidential elections, constructing an original data set by scraping six news websites for each year. The sites are: the New York Times (NYT), Fox News, Wall Street Journal (WSJ), and Yahoo! News (Yahoo) 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. For each year, we have two sites typically perceived as left-of-center (NYT and HuffPost in 2012, and NYT and WashPost in 2016), two perceived as right-of-center (Fox News, WSJ), and two perceived as relatively neutral, with readerships that are aligned, though not completely, with these reputations.³

We use the term *slant* to refer to favorability of headlines (which we focus on in our analysis, rather than article content, as we explain below) to one party or the other’s chances in the election. This definition of *slant* is a natural one for our context, though somewhat distinct from definitions used in other papers in this literature, discussed below. Our definition does not necessarily imply distortion from truth; however, horse race headlines indeed might distort, mislead, or simply emphasize different factors about a given state of the horse race in various ways via framing and selection. Two headlines about a given poll result can send different messages via different word choices, or by emphasizing different aspects of polling information, such as changes in polls versus levels, particular states versus national results, results for particular demographic groups, or even results for some polls versus others. Outlets can also shape the horse race news they produce by simply deciding when, and to what extent, to report on polls at all. *Slant* can be driven by psychological and information-based demand side factors and/or supply-side factors, which we discuss in more detail below.

While we do not have micro-level click data for our main analysis, we use an innovative

³See, for example, <https://www.allsides.com/media-bias/media-bias-chart> for a reference to the outlets having these reputations. We are able to use four websites (including at least one from each reputational category) for both years: the NYT, Fox, WSJ, and Yahoo. We substitute USAT for Google, and HuffPost for WashPost, in 2012 due to data limitations, as we discuss in Section 3. For evidence of ideological alignment between outlets and readers, see, e.g., [Gentzkow and Shapiro \(2010\)](#), [Gentzkow and Shapiro \(2011\)](#), and [Flaxman et al. \(2016\)](#).

measure of demand: the websites’ publicly reported lists of most popular articles. In addition to public availability, another advantage of the “most viewed” list data over most micro data is that the most viewed lists capture demand for the entire market, and not just a sample of users observed. We interpret our results carefully in light of the aggregate nature of the data and discuss how reader heterogeneity could affect results in various parts of the paper. Furthermore, to better understand individual reader behavior, we conduct a complementary micro-level analysis, using data from incentivized surveys based on real-time, quasi-randomized news for a closely related, but distinct, topic, the presidential debates.⁴

We discuss our web data in Section 3. Since articles attain most viewed status based on clicks, which are primarily based on headlines, we code the partisan congeniality of headlines, and not the entire article.⁵ Headlines are short, so we find it is most effective to use content analysis (human coding) rather than computational analysis to code them. We used a team of highly incentivized and closely supervised master MTurkers for this content analysis. We examine general patterns in slant of the outlets’ reporting, and confirm that slant of all six outlets varies over time within each year, and each outlet’s slant was correlated with poll averages in each year.

In Section 4.1, we present the across-outlet analysis. We assess slant on both the intensive and extensive margins (slant in headlines actually reported, and slant in the decision of whether to report on the horse race at all, respectively). We find robust and substantial evidence of intensive margin congenial slant for the HuffPost in 2012, the NYT in 2016, and Fox in both 2012 and 2016. The evidence for congenial extensive margin slant is weaker and less consistent. Perhaps surprisingly, the strongest evidence of extensive margin slant was for the WSJ, which appeared more likely to report horse race stories when polls were more favorable to Democrats (i.e., uncongenial to Republican readers) in both years.

In Section 4.2 we present the within-outlet analysis. We use simple linear probability

⁴We use debate news, and not (pure) horse race news, because the timing of the debates (and therefore debate news) was known in advance, allowing us to plan when to run the surveys.

⁵For evidence that headlines are the main driver of article clicks, see, e.g., [Powell et al. \(2019\)](#). Headline and article congeniality are generally well aligned as shown by, e.g., [Tremayne \(2015\)](#), which we confirm in spot checks.

models, regressing a dummy variable for making an outlet’s most popular list on slant-outlet interactions, with outlet and day fixed effects to account for general variation in article popularity driven by these factors. We find non-robust and only marginally significant evidence that stories more congenial to typical readers were more popular for just one outlet, in one year (Fox in 2012). Point estimates for the NYT and HuffPost that year are near zero, with reasonably high precision. For 2016, we find somewhat more significant and robust evidence of non-null effects—that news that was more *uncongenial* for typical readers was more likely to be popular—for both the NYT and Fox. The survey results also largely fail to yield significant results indicating preferences for congenial news from ideologically aligned outlets.⁶

Thus, our main results are that outlets did often slant stories in a way that was congenial to their typical readers, but only on the intensive margin, and there was still variation in slant across stories within outlets. Moreover, stories that were most congenial to an outlet’s typical readers were not most likely to make it to the outlet’s most viewed list. In Section 5, we discuss interpretation of our results with respect to the broader question of what drives supply and demand for slant, based on the framework provided by [Gentzkow et al. \(2015\)](#). Why does variation in slant across outlets make different outlets appeal to different readers? Do media consumers prefer to get news from ideologically aligned outlets because such news “feels good” due to a psychological mechanism, or because consumers perceive this news to be truly more informative? Within the information category, there are two distinct mechanisms: instrumental value for decisions and trust; can we comment on which, if either, our results support? And to what extent do outlets impose slant on consumers for supply-side reasons, i.e., reasons favoring an outlet’s interests other than attempts to appeal to consumer preferences?

Congenial across-outlet slant could be due to outlets pandering to the psychology mecha-

⁶The point estimates of preferences for congeniality from the NYT from Democrats and Clinton voters are near zero, and insignificant for Republicans for Fox, but we do obtain marginally (10%) significant results for a preference for congenial news from Fox for Trump supporters. This asymmetry in results across parties is directionally consistent with a similar asymmetry in the results for the most viewed list data, as Fox was the only outlet with any evidence of within-outlet preference for congeniality. We also find that Democrats and Clinton voters were more likely to choose Fox news when it was congenial to Democrats, but Trump voters were consistently unlikely to choose NYT news, and were more likely to choose a non-political news option when all political options were uncongenial.

nism. Readers may click on more congenial stories because they “feel good” and the slanted distribution supplied by various outlets may maximize their clicks. But we cannot immediately rule out that across-outlet slant is driven by informational factors either in addition to, or instead of, psychological factors. While it seems implausible that news reported with systematically congenial slant would provide greater instrumental value for decisions, it could be more trusted, as consumers may have perceived congenially slanted horse race stories as more informative (if consumers had priors biased toward their preferred candidate being more likely to win). However, the incentive for outlets to pander to these priors would decline as the election approaches, and we find it was not the case that actual slant declined in this way.

We interpret the lack of within-outlet demand for additional congeniality for typical readers, and possible within-outlet demand for uncongeniality, to be likely due to satiation of the psychology mechanism and perceptions that less congenial stories were relatively credible and informative. Some outlets’ stories may have been slanted, on average, more than the optimal slant with respect to maximizing readers’ preferences due to supply-side factors like an outlet’s desire to promote a bandwagon voting effect, or simple miscalculation. Heterogeneity of reader behavior may also help explain these within-outlet demand results: demand for news that was more uncongenial to typical readers may have been driven partly by relatively moderate or even “counter-partisan” readers (though we argue this latter possibility is unlikely to be the main explanation).

Moreover, we discuss how the extensive margin (WSJ) results perhaps best support the instrumental information theory of slant demand. We also discuss how several subtle contextual factors could influence the relevance of each mechanism. Our results hence fail to support a simple, but *a priori* plausible, story in which the most partisan stories on partisan outlets were most popular for purely psychological reasons. By contrast, our results provide support for each of the major mechanisms potentially driving slant considered in prior literature, arguing against an either/or approach to understanding drivers of slant.

2 Related Literature

For reviews of research on partisan selective exposure in psychology and communication studies see [Hart et al. \(2009\)](#) and [Stroud \(2011\)](#), respectively. Most of this research relies on experimental settings. See [Golman et al. \(2017\)](#) for a review of economics literature on the demand for congenial information outside of political media. 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. [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. In contrast to most of the existing literature, our study offers an investigation of partisan selective exposure to news based on real-world data.

Two recent papers present experiments that study demand for belief-confirming information that by design abstract from psychological factors. [Charness et al. \(2018\)](#) study how variation in the nature of bias, which drives the optimality of seeking information sources biased either towards or against one’s priors, affects choice of information source in a lab experiment. They find that subjects tend to excessively choose belief-confirming sources, despite the absence of any psychological motive. [Montanari and Nunnari \(2019\)](#) study a trade-off between alignment of a source’s bias and source reliability using an online (Prolific) experiment. They find, perhaps surprisingly, excessive willingness to choose more reliable sources even when misaligned with the subject’s prior and, perhaps less surprisingly, excessive responsiveness to confirmatory advice from sources aligned with the subject’s prior.

A few studies explicitly investigate mechanisms underlying news demand, rather than demand for information in general. Using artificial news articles on controversial political topics, [Metzger et al. \(2015\)](#) find some support for the role of psychological utility, but stronger support for the importance of trust in driving news choices. [Simonov and Rao \(2018\)](#) show that exposure to pro-government news in Russia is shaped by the persistent preferences of consumers for certain outlets. [Thaler \(2019\)](#) conducts an MTurk experiment to cleanly distinguish motivated reasoning from Bayesian inference with heterogeneous priors as explanations

for trust (or lack thereof) in “fake news”, and finds that the results largely support motivated reasoning, consistent with what we refer to as the psychology mechanism of demand for slant.

Another recent paper that is especially closely related to ours is [Chopra et al. \(2019\)](#). The authors conduct a survey experiment of demand for NYT news on effects of Trump’s healthcare plan with a large sample size. They find that respondents prefer suppressed information, also consistent with the psychology mechanism. Our papers complement each other in that we examine different news issues, outlets, and use different methods; our paper also concludes that psychology is important, and discusses evidence of other mechanisms.

Finally, our research relates to previous studies of bias in horse race reporting. [Searles et al. \(2016\)](#) and [Tremayne \(2015\)](#) compare poll coverage with underlying poll results during the 2008 and 2012 US presidential campaigns, respectively. We confirm the result obtained by both studies that the coverage is slanted in favor of news outlets’ average readers. However, in addition to using data from different outlets and a more recent year, our study is novel in that we combine data on slant with information on the popularity of stories among readers.

3 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 determined. The websites of Fox, WSJ, NYT, WashPost, Google, and Yahoo were scraped three times daily until the day of the election, November 8. 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.⁷ 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

⁷Given 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.

recent period of time.⁸

News headlines for the 2012 presidential election were collected by scraping snapshots of the outlets' websites stored by web.archive.org, also dating between July 27 and election day of that year.⁹ We were forced to make substitutions for Google and WashPost due to their snapshot data being unavailable and replaced them with USAT and the HuffPost, respectively.¹⁰ 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 is 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 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

⁸Both Yahoo and Google aggregate news from other sources, occasionally including those in our sample, but these comprise a small fraction of Yahoo and Google's stories in total.

⁹We cannot scrape directly from the websites since, even if horse race stories from earlier years are still available on those sites, the sites do not provide access to the actual homepages presented to readers and most viewed lists from prior dates.

¹⁰Since 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`. Scraping historical Google and WashPost data using the websites' sitemaps is not technically feasible, to our knowledge.

of being an opinion piece in the headline, due to our focus on hard news stories.¹¹

Because headlines are short, the degree to which they favored one party or the other is relatively easy to manually code, and difficult to do so computationally. We therefore used manual content analysis for the coding of headline congeniality. 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. 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.¹² These workers each passed initial 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

¹¹We 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}. Excluding opinion stories should reduce the prevalence of within-outlet variation in an article’s credibility and information content. For example, some readers might perceive certain opinion writers to be particularly credible, which could be correlated with reader demand and headline congeniality, as some opinion writers might be more likely to write about the horse race when it is going well (or poorly) for the preferred candidate. Excluding opinion pieces therefore makes the data set more homogeneous and “cleaner”, reducing the chance of omitted variable bias (however, we do have complete data on opinion pieces from 2012 and results are very similar when these pieces are included). With the exceptions of opinion pieces, readers do not see the author (if any) before clicking on the headline, in most cases.

¹²Some headlines did refer to Fox News polls; results are largely similar when these are dropped.

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. Krippendorff’s α , a standard measure of inter-coder reliability, for all the coded headlines is 0.313.¹³ 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 0.80.¹⁴ 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 0.900.¹⁵ 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 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. As noted in the introduction, $Slant_i$ simply measures favorability of a headline to one party or the other and not necessarily distortion or any type of misreporting.

Figure 1 presents smoothed plots of daily means of each slant measure versus daily poll

¹³Krippendorff’s α measures similarity of ratings across coders. This captures both the quality of coding and objectivity of content to be coded; the basic intuition is that if content is objective and coders do good work, we would expect their ratings to be more positively correlated. See, e.g., White and Marsh (2006) for further discussion.

¹⁴For 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.

¹⁵See Table A1 for explanation of the terms irrelevant and ambiguous provided to the coders.

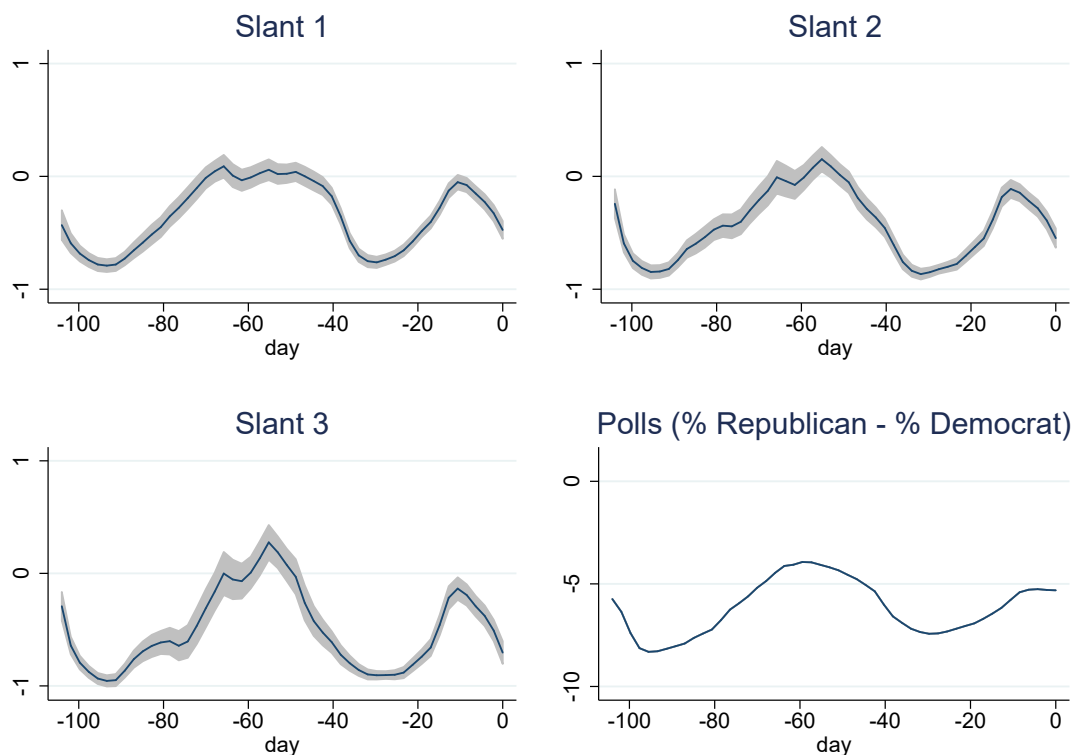
averages (percent planning to vote Republican minus percent planning to vote Democrat; obtained from R’s “pollstR” library) in 2016.¹⁶ The plots are quite similar to one another, and to the polling average, supporting the validity of all three slant measures (this is also true for the 2012 data, and we omit this figure only in the interest of brevity). However, a number of results presented in Section 4 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 provide direct examples of 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 experimented with several formulaic text-based measures of slant, and found that coding 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”) yielded the largest correlations with our manually coded slant measures (average correlation of 0.33 with $Slant_i$). However, the correlation with the daily poll average was negligible (0.01), and correlations were similarly very small for other text-based measures. The correlation of each of the manually coded $Slant_i$ ’s with the polls is approximately 0.42, which supports the superiority of the human-coding approach (as expected, given the small amount of text in each headline) and the use of these measures in the analyses.

Appendix Figures A1 and A2 present smoothed $Slant_1$ plots broken out by outlet and most viewed/other articles. Patterns of slant over time are generally similar across outlets, and in unreported results we find that the slant of each outlet’s articles is significantly correlated with

¹⁶The pollstR library uses the Huffington Post’s Pollster API, which incorporates thousands of polls based on a set of standard, non-partisan criteria (see <https://elections.huffingtonpost.com/pollster/faq>).

Figure 1: 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.

poll results (whether studied separately by outlet or pooling articles for all outlets). These results suggest that for each outlet, horse race stories are driven to a substantial degree by what is truly occurring in the campaign at the time. Table 1 implies a somewhat different story. This table reports the number of unique most viewed and other articles per outlet, and their means for various slant measures.¹⁷ The average slants vary across outlets substantially,

¹⁷The 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

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. We analyze these differences more formally in the next section.

4 Analysis

4.1 Across-outlet slant

There are two basic ways across-outlet congenial slant could occur: on the intensive margin (variation in slant of horse race stories actually reported, given the timing of reporting) and on the extensive margin (variation in the number and timing of horse race stories reported). We state explicit hypotheses regarding these two types of slant as follows:

Hypothesis 1. *Intensive margin congenial slant: horse race headlines reported by left-leaning outlets are relatively favorable to Democrats, and horse race headlines reported by right-leaning outlets are relatively favorable to Republicans, ceteris paribus.*

Hypothesis 2. *Extensive margin congenial slant: left-leaning outlets are relatively likely to report horse race stories favorable to Democrats, and right-leaning outlets are relatively likely to report horse race stories favorable to Republicans, ceteris paribus.*

To address Hypothesis 1, we use a story-level data set to estimate a separate regression for each definition of $Slant_i$ ($i \in 1, 2, 3$) of the following form:

$$Slant_{ijkt} = \alpha + Outlet_j + Day_t + \epsilon_{ijkt}. \quad (1)$$

$Slant_i$ of each story k published by outlet j on day t is regressed on outlet and day fixed effects, yielding the estimates of the $Outlet_j$ ’s and Day_t ’s. The day fixed effects account for mean slant across outlets on the first day that the story was reported driven by new poll results of headlines for a given story as they are on a single version of a headline.

or other factors. We do not include daily polls and poll changes as covariates, 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 to preserve degrees of freedom.

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 $Outlet_j$ estimates thus represent the mean difference in slant for outlet j 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, consistent with Hypothesis 1. 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, also consistent with Hypothesis 1. The point estimates for each non-Yahoo fixed effect 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 on a scale from -1 to 1.

To estimate slant on the extensive margin, we construct outlet-level daily time series data sets. We cannot use story-level observations here as we need to account for days in which horse race stories were not reported. We run separate regressions for each outlet of the following form:

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

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 in addition to, or instead of, the other measures of “true slant.” 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.¹⁸ 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, only one estimate is significant, and at just 10%, but all of the estimates have signs consistent with congenial extensive margin slant providing marginal support Hypothesis 2. There are also two significant estimates consistent with this hypothesis for the NYT in 2016 (i.e., a greater number of stories when the news is more favorable to Democrats). The strongest and more robust results, however, are for the WSJ. Half or more of the 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 news was *less* favorable to Republicans. Moreover, in unreported results, we also find significant *negative* effects of congeniality to Democrats of the previous week’s poll changes on the number of stories reported by the WashPost. We discuss these findings further in Section 5.

4.2 Within-outlet news demand

Given concerns about echo chambers discussed in Section 1 and the focus on psychologically driven demand for news in the literatures outside of economics especially, our hypothesis motivating the within-outlet news demand analysis was as follows:

¹⁸The 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.

Hypothesis 3. *Within-outlet demand for congenial slant: left-leaning outlet readers are more likely to click on horse race headlines favorable to Democrats, and right-leaning outlet readers are more likely to click on horse race stories favorable to Republicans, ceteris paribus.*

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

$$MostViewed_{jk} = \alpha + Outlet_j + \beta_j Slant_{ijk} + \sum_t Day_t^{jk} + \epsilon_{jk}. \quad (3)$$

$MostViewed_{jk}$ is a dummy for whether story k (ever) made its outlet's (j 's) most viewed list, which we regress on outlet fixed effects (to account for different mean probabilities of making the most viewed list by outlet) yielding estimates of $Outlet_j$, a different slant term for each outlet j (i.e, an interaction of a dummy for outlet j and $Slant_{ik}$) yielding the estimates of β_j 's (one for each outlet), 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). These account for unobservables changing horse race news demand over time, such as James Comey's letter to Congress in October, 2016, and interest in the election increasing as it approaches. We also consider a model with a date polynomial. The β_j 's are the coefficients 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. We cluster standard errors by the first date a story was reported.

Results are presented in Table 4. There is some evidence supporting Hypothesis 3 for Fox in 2012, but this is significant for $Slant_1$ only, and only at 5%. The magnitude 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 stories *less* congenial to typical readers were more likely to be most viewed. The significant effects for Fox are 13–19 percentage points per unit decrease in congenial slant, and for the NYT, 19–

31 percentage points.¹⁹ The standard errors for the insignificant estimates for the HuffPost, WashPost, and NYT in 2012 are all less than 0.1 (10 percentage points per unit change in slant), 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 A3). We also consider a specification in which we split out the measure of slant into congenial and uncongenial effects for each outlet. It is possible that avoiding negative news is more important than seeking positive news, or vice versa. We analyze this specification (results unreported but available on request) and find no systematic patterns.

4.3 Micro-level data

While (as noted above) the outlets' reputations are aligned with the ideologies of the majority of the outlets' 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 was 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 more moderate readers or even from those who were typically ideologically misaligned with the outlet. For example, the clicks on NYT stories favorable to Republicans may have mostly come from Republican readers.

To partially address this issue we collect additional micro-level data from incentivized surveys on a related topic, presidential debates. Our survey design allows us to study the effects of congeniality on within-outlet news demand, with real, timely, presidential election-related news. In the morning (between 9:00 AM and 10:00 AM) following each of the first three 2016 US presidential election debates, we conducted a survey on MTurk on interest in debate news versus other news.²⁰ We used debate news, rather than horse race news, for the surveys

¹⁹However, when we use an alternative clustering specification, clustering by outlet with wild bootstrap tests (due to the small number of clusters), nearly all estimates for both years are insignificant at the 10% level.

²⁰The first and third debates were between the presidential nominees, and the second one was between the vice president candidates. See Appendix Table A5 for sample statistics for key variables. Compared to the population, survey respondents recruited on MTurk tend to

because the timing of debates was known well in advance. This allowed us to prepare to post 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.

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 A6. 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, so they should pick the article they are (were) truly most interested in. 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 A6. 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 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.

should reflect the articles that they would be relatively likely to read in similar real-world situations.

We analyze the data to test whether interest in debate news from the ideologically aligned outlet increased when it was more congenial. We also assess whether interest in debate news from the other outlet increased when it was more congenial, but we note below, external validity of this analysis is more questionable.²¹ We relegate details of the survey data analysis to the appendix and discuss the main results here.

We find that Democrats were just as willing to get news from the NYT when it was uncongenial, and Republicans were unwilling to get NYT news even when it was congenial. This supports the interpretation of the NYT most viewed data (i.e., the within-outlet demand results) being primarily driven by regular NYT readers. We also find that Democrats were willing to click on congenial Fox news, and respondents who planned to vote for Trump were somewhat less willing to click uncongenial Fox news versus congenial news. This provides some support for Fox stories congenial to Democrats being clicked on by them, making the interpretation of the Fox most viewed data less clear.²² However, clicking on an “uncongenial outlet” is perhaps much more likely in the survey context than in reality since in the survey users were directly presented with headlines from that outlet (which users might avoid in reality). Moreover, Fox readers in reality are mostly homogeneous; [Gentzkow and Shapiro \(2011\)](#) report that 76% identify as conservative, and only 10% as liberal. Finally, we note that substantial fractions of Trump supporters (24%) and Republicans (28%) were willing to

²¹We also manually collected 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. The data (available on request) shows that the number of stories was fairly constant for both outlets in 2012. However, in 2016, there are indications of a 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 of debate links are correlated with the congeniality of the debate outcome for their readers.

²²The asymmetry in results across parties is consistent with the 2012 aggregate results also indicating right-of-center consumers had a greater within-outlet preference for congenial news.

click on the uncongenial Fox story, supporting the existence of strong within-outlet demand for news regardless of its congeniality. Results are similar when we restrict the sample to respondents who only “lean” toward identifying with a party, providing evidence that the web results were unlikely to be driven by the outlets’ ideologically aligned, but relatively moderate, readers.

4.4 Website Traffic

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 race 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 A3 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.

5 Discussion

As noted in Section 1, there are several mechanisms potentially driving demand for slant that have been studied in the media bias theory literature. See Gentzkow et al. (2015) for simple

models of each of these mechanisms. The first is that slant provides direct psychological value to consumers. News that confirms one’s prior beliefs or outcomes that one hopes to occur may be pleasant, and belief-challenging news unpleasant, due to cognitive dissonance, self-image and social identity support, anticipation utility, and related factors; see, for example, [Iyengar and Hahn \(2009\)](#), and from economics, [Mullainathan and Shleifer \(2005\)](#) and [Bernhardt et al. \(2008\)](#).

It is also possible that demand for slant is due to congenially slanted news being perceived as more informative. [Gentzkow et al. \(2015\)](#) distinguish between two underlying channels, which they refer to as “delegation” and “reputation”. Delegation refers to consumers rationally preferring 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. The reputation channel occurs when media consumers prefer like-minded sources because those sources are more trusted—they are perceived as objectively more credible ([Gentzkow and Shapiro, 2006](#)), though perhaps mistakenly ([Vallone et al., 1985](#); [Stone, 2011](#); [Kelly, 2018](#)).²³ Welfare implications of political slant are generally worst when slant is due to psychology, and best when due to delegation.

Finally, it is possible that outlets slant the news for supply-side reasons. Outlet owners or journalists may have their own ideological agenda or owners may benefit economically when the news supports certain policies. Supply-side driven bias is unsurprisingly harmful to consumer welfare, and is generally, but not necessarily, reduced by competitive forces.

Which of these mechanisms do our results support? Evidence supporting Hypothesis 3 (greater within-outlet demand for slant congenial to an outlet’s typical readers) would have been relatively clean evidence in support of the psychology mechanism. Our lack of evidence supporting this hypothesis may thus appear at odds with this mechanism. However, given the

²³Note that in both this case and the instrumental value case, consumers are information-seeking. But these two mechanisms differ in 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.

across-outlet congenial slant, within-outlet effects must be interpreted accordingly. Across-outlet congenial slant could also be due to outlets pandering to reader psychology, and could even satiate the psychological desire for “good news”, which could explain the lack of additional demand for particularly congenial news from such typically congenial outlets. In Appendix Section A.3 we present an illustrative model to clarify this point. In the model, readers have a preference for both congenial slant and news accuracy. Consequently, their within-outlet demand for congenial slant declines as the mean congenial slant of an outlet’s stories grows.²⁴ The model helps to clarify how outlets may have exceeded demand-maximizing mean congenial slant due to simple miscalculation or due to supply-side motives such as a desire to promote the bandwagon effect (Morton et al., 2015).²⁵

Could across-outlet congenial slant for horse race news also be explained by an information mechanism? Delegation seems to be an unlikely explanation. While horse race news could be relevant to some decisions such as campaign activism, horse race news has been criticized for being relatively non-substantive (Zoizner, 2018). It is especially implausible that a general slant favoring a reader’s ex ante preferred party would be useful for decisions. However, this slant could be at least perceived as yielding more accurate news, if readers had priors biased toward favoring the chances of their preferred candidates. Thus, congenially slanted news by an outlet could increase or confirm a reader’s trust in that outlet. There is evidence that individuals do have such priors and draw inferences like this about the quality of news (Stiers and Dassonneville, 2018; Madson and Hillygus, 2019). However, this type of pandering to readers’ beliefs should decline as the election approached, as the election outcome provides feedback on the accuracy of prior horse race reporting. In unreported results, we find that this is not the case (congenial slant does not decline as the election approaches), supporting psychology being at least part of the explanation for across-outlet slant.

A key question that remains is why would demand ever be higher for news that was uncongenial to typical readers (and why in 2016 only). One possibility is reader heterogeneity;

²⁴We do not model this point, but it is worth noting that variation in demand within-outlets could be part of a profit-maximizing firm strategy.

²⁵By contrast, the “underdog effect” would increase an outlet’s incentive to downplay the chances of the candidate it favors.

however, as noted above, our survey results provide some evidence against this possibility. It seems especially implausible that the greater popularity in 2016 of pro-Republican stories for the NYT, and of pro-Democrat stories for Fox, was driven by ideologically misaligned minorities. This would mean that Democrats would have to have a stronger preference for congeniality than Republicans for Fox stories *and* Republicans would have to have a stronger preference for congeniality than Democrats for NYT stories. These differences in preferences would have to be sufficiently large to outweigh the misaligned readers being minorities (for example, the Democrats' preference for pro-Democrat Fox stories would have to be substantially greater than the Republicans' preference for pro-Republican stories on Fox, since there are many more Republican Fox readers than Democrats).

We think that both of the information mechanisms, especially reputation, more plausibly explain greater within-outlet demand for uncongenial news. If there was variation in credibility of stories within outlets, and headlines provided a signal of credibility of the article's content, then uncongenial headlines might have signaled high credibility for stories from an outlet that is typically congenial. For example: if readers knew the NYT had typically published stories implying that Clinton's chances of winning the upcoming election were high, and saw an NYT headline implying Clinton's chances were now ambiguous or even low, readers might have inferred that the underlying evidence for the story was particularly compelling. The apparent growth in the NYT's mean slant from 2012 to 2016 is consistent with this claim, as it would imply greater demand for uncongenial news in 2016. (The across-outlet estimates for Fox do not imply this type of growth in slant across years, however.)

The negative within-outlet slant-demand relationship could arguably even be compatible with instrumental information seeking, if readers drew inferences on candidate quality from horse race news, as negative horse race news about a preferred candidate is more likely to be decision relevant than positive. For instance, it is perhaps possible that Fox readers could have been more likely to click on pro-Clinton headlines because of the value of this information for changing choices related to campaign involvement (donations, activism), or even whom to vote for or whether to turn out to vote (Bursztyn et al., 2017). Both mechanisms could be stronger in 2016 due to stronger average congenial slant in 2016 or changes in other contextual

factors such as the nature of the candidates or overall media landscape.

Another question that remains is what explains the extensive margin result for the WSJ, implying it provided news that was biased toward being uncongenial to Republican readers. This result is also plausibly explained by the delegation mechanism. If the WSJ’s Republican readers were relatively affluent, then these readers may have been more likely to make campaign contributions to Republican candidates when they were down in the polls. In fact, it is especially plausible that this type of instrumentally valuable is relevant to WSJ readers given their demographics, and especially plausible for extensive margin slant since this is more consistent with selection rather than distortion of information.^{26 27}

We close this section by discussing two factors that have been neglected thus far. The first is another mechanism that could drive horse race news demand: surprise (Ely et al., 2015). If consumers were more likely to click on more surprising news, this could have also caused a negative within-outlet relationship between congeniality of slant and demand for typically congenial outlets. We investigate this possibility by constructing a measure of 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 (available on request) are insignificant, which suggests that surprise did not play a major role in explaining the popularity of uncongenial stories in 2016.

Second, another factor that could affect demand worth noting is article placement on web

²⁶An alternative explanation for this result is that the WSJ is a distinct type of outlet, attracting right-leaning readers via its business coverage even while the news leans left for other reasons (Groseclose and Milyo, 2005).

²⁷The WashPost also exhibited some evidence of uncongenial slant on the extensive margin, but unlike the WSJ, this was offset by congenial slant on the intensive margin.

pages. It is certainly plausible that articles shown at the top of pages, or articles with larger headlines and photos, were more likely to be clicked, independent of their content, and so this factor is relevant to our within-outlet analysis. Unfortunately we did not scrape data on placement, so we cannot control for this variable directly. How might this omission affect our results? If placement was uncorrelated with congeniality, which is plausible, especially given our results implying that extensive margin slant is relatively weak, then not controlling for this factor would not bias results. If more congenial stories were placed more prominently, then the within-outlet congeniality-demand relationship could be biased upward. This would mean that demand for uncongenial stories was even higher than it appears, and our conclusion that there was substantial demand for uncongenial news within-outlets would be “safe.” If uncongenial stories were placed more prominently, the opposite bias could occur, implying that demand for congeniality was higher than that implied by our estimates. This possibility seems very unlikely given the absence of uncongenial bias in both types of across-outlet congenial slant that we estimate for Fox and the NYT. Moreover, even if this were the case, we could still conclude from our analysis that partisan readers were often exposed to uncongenial content within their preferred outlets. Examining the relationship between article placement and demand is likely a worthwhile topic for future research.

6 Additional remarks

Rather than recap our results again, we provide a few remarks on issues that we have not yet noted. Our results clarify the connection between two of the mechanisms potentially driving demand for slant, trust and psychology. Although some media bias theory papers model these mechanisms as distinct factors, they are closely related, and in fact each can influence the other. It is more straightforward to see that trust can be influenced by psychological factors. Cognitive dissonance and confirmation bias might affect reactions to information observed in the past; these forces may even 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—the evidence of demand for “bad news” for consumers

across the political spectrum—help us to see is that the psychological value of news also likely 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 that there can be natural limits to the degree of psychology-driven news distortion: if distortion becomes too extreme, it destroys credibility, defeating the purpose. This trade-off is likely especially relevant to issues reported on repeatedly for a long period of time (such as US presidential campaigns), and especially when true outcomes are eventually revealed, both of which are the case for horse race news.

We may also have found within-outlet demand for bad news to be relatively high due to features of our sample. The outlets we study are relatively mainstream.²⁸ We might expect that within-outlet demand for congeniality is higher within more extreme outlets. Moreover, the US news cycle for presidential horse race news is fairly long, lasting from when the parties’ nominations are effectively determined in spring until the election in early November. This invariably yields ups and downs and the sheer length of the campaign may lead to greater demand for variation in news including relatively uncongenial news; this demand might be lower in settings with shorter campaign periods.

The complexity of our results also demonstrates (or reminds us) that both media and reader behavior can vary substantially over time and across outlets and issues. In 2016, Fox (the NYT) may have positioned itself to the right (left) of its readers, despite this not occurring in 2012. This may have been due to differences in candidates, or any number of other factors. The WashPost and NYT have similar ideological reputations, but there were substantial differences in their reporting and reader behavior. The WSJ’s horse race reporting appears quite distinct from its reputation, or at least the reputation of its editorial section, as opposed to that of Fox and the NYT. These nuances in our findings support the importance of considering context in interpretation, and the limitations one should consider in 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

²⁸See, e.g., https://www.adfontesmedia.com/wp-content/uploads/2018/08/Media-Bias-Chart_4.0_8_28_2018-min.jpg.

the various outlets.

References

- 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.
- J. Burke. Primetime spin: Media bias and belief confirming information. *Journal of Economics & Management Strategy*, 17(3):633–665, 2008.
- L. Bursztyrn, D. Cantoni, P. Funk, and N. Yuchtman. Polls, the press, and political participation: The effects of anticipated election closeness on voter turnout. NBER working paper, 2017.
- A. Casas, Y. Fawaz, and A. Trindade. Surprise me if you can: The influence of newspaper endorsements in U.S. presidential elections. *Economic Inquiry*, 54(3):1484–1498, 2016.
- J. Chan and W. Suen. A spatial theory of news consumption and electoral competition. *Review of Economic Studies*, 75(3):699–728, 2008.
- G. Charness, R. Oprea, and S. Yuksel. How do people choose between biased information sources? evidence from a laboratory experiment. *manuscript, downloaded from: http://econ.ucsb.edu/~sevgi/CharnessOpreaYuksel_Dec2018.pdf*, 2018.
- C.-F. Chiang and B. Knight. Media bias and influence: Evidence from newspaper endorsements. *The Review of Economic Studies*, 78(3):795–820, 2011.
- F. Chopra, I. Haaland, and C. Roth. Do people value more informative news? Working paper, 2019.

- R. Dewenter, M. Linder, and T. Thomas. Can media drive the electorate? the impact of media coverage on voting intentions. *European Journal of Political Economy*, 58:245–261, 2019.
- 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. Gambaro and R. Puglisi. What do ads buy? daily coverage of listed companies on the italian press. *European Journal of Political Economy*, 39:41–57, 2015.
- M. Garz, J. Sørensen, and D. F. Stone. Partisan selective engagement: Evidence from Facebook. Working paper, 2019.
- M. Gentzkow, J. Shapiro, and D. Stone. Media bias in the marketplace: Theory. In S. Anderson, J. Waldfogel, and D. Stromberg, editors, *Handbook of media economics*, pages 623–645. Elsevier, 2015.
- M. Gentzkow and J. Shapiro. Media bias and reputation. *Journal of Political Economy*, 114(2):280–316, 2006.
- M. Gentzkow and J. Shapiro. What drives media slant? Evidence from US daily newspapers. *Econometrica*, 78(1):35–71, 2010.
- M. Gentzkow and J. Shapiro. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126(4):1799–1839, 2011.
- R. Golman, D. Hagmann, and G. Loewenstein. Information avoidance. *Journal of Economic Literature*, 55(1):96–135, 2017.
- S. Greenstein and F. Zhu. Is Wikipedia biased? *American Economic Review: Papers & Proceedings*, 102(3):343–348, 2012.

- T. Groseclose and J. Milyo. A measure of media bias. *Quarterly Journal of Economics*, 120(4):1191–1237, 2005.
- 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.
- C. Huff and D. Tingley. “Who are these people?” Evaluating the demographic characteristics and political preferences of MTurk survey respondents. *Research & Politics*, 2(3):1–12, 2015.
- S. Iyengar and K. S. Hahn. Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication*, 59(1):19–39, 2009.
- N. Karlsson, G. Loewenstein, and D. Seppi. The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115, 2009.
- D. Kelly. Evaluating the news: (Mis)perceptions of objectivity and credibility. *Political Behavior*, forthcoming, 2018.
- V. Larcinese, R. Puglisi, and J. M. Snyder Jr. Partisan bias in economic news: Evidence on the agenda-setting behavior of U.S. newspapers. *Journal of Public Economics*, 95:1178–1189, 2011.
- G. J. Madson and D. S. Hillygus. All the best polls agree with me: Bias in evaluations of political polling. *Political Behavior*, forthcoming, 2019.
- 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.
- G. Montanari and S. Nunnari. Audi alteram partem: An experiment on selective exposure to information. Technical report, Working Paper, 2019.

- R. B. Morton, D. Muller, L. Page, and B. Torgler. Exit polls, turnout, and bandwagon voting: Evidence from a natural experiment. *European Economic Review*, 77:65–81, 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.
- T. E. Powell, T. G. van der Meer, C. B. Peralta, et al. Picture power? the contribution of visuals and text to partisan selective exposure. *Media and Communication*, 7(3):12–31, 2019.
- R. Puglisi and J. M. Snyder Jr. Empirical studies of media bias. In *Handbook of media economics*, volume 1, pages 647–667. Elsevier, 2015a.
- R. Puglisi and J. M. Snyder Jr. The balanced us press. *Journal of the European Economic Association*, 13(2):240–264, 2015b.
- B. Qin, D. Stromberg, and Y. Wu. Media bias in China. *American Economic Review*, 108(9):2442–2476, 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. Thaler. The “fake news” effect: An experiment on motivated reasoning and trust in news. 2019.
- 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.
- M. D. White and E. E. Marsh. Content analysis: A flexible methodology. *Library trends*, 55 (1):22–45, 2006.
- A. Zoizner. The consequences of strategic news coverage for democracy: A meta-analysis. *Communication Research*, forthcoming, 2018.

Table 1: Average headline slants and article counts by outlet (samples for each $Slant_i$ defined in text)

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 headline slant across outlets

Outlet	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012 (Reference outlet = Yahoo)			
Fox	0.340** (0.130)	0.471*** (0.125)	0.648*** (0.176)
WSJ	-0.306* (0.176)	-0.165 (0.170)	0.059 (0.227)
USAT	-0.117 (0.182)	0.012 (0.175)	0.037 (0.245)
NYT	-0.312** (0.121)	-0.161 (0.131)	-0.142 (0.206)
HuffPost	-0.499*** (0.108)	-0.335*** (0.114)	-0.236 (0.178)
Adj R^2	0.412	0.407	0.494
N	425	382	277
Panel B: 2016 (Reference outlet = Google)			
Fox	0.238** (0.116)	0.308** (0.149)	0.353* (0.191)
WSJ	-0.235 (0.179)	-0.506*** (0.171)	-0.479** (0.203)
Yahoo	-0.082 (0.130)	-0.143 (0.150)	-0.360** (0.149)
NYT	-0.363*** (0.097)	-0.230* (0.120)	-0.411*** (0.120)
WashPost	-0.162** (0.079)	-0.118 (0.092)	-0.260** (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: $Slant_i$, and coefficients presented in table are outlet fixed effects ($Outlet_j$ from equation (1)). 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. 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)				
<i>Slant</i> ₁	0.251 (0.253)	-0.038 (0.504)	0.284 (0.224)	-0.075 (0.265)
<i>Slant</i> ₂	0.309 (0.258)	-0.740 (0.665)	0.182 (0.221)	0.014 (0.210)
<i>Slant</i> ₃	0.202 (0.267)	-0.624 (0.647)	0.224 (0.207)	-0.335 (0.207)
Panel A2: RHS = mean Republican poll advantage (2012)				
<i>Slant</i> ₁	0.077 (0.089)	-0.584** (0.292)	-0.067 (0.091)	0.051 (0.089)
<i>Slant</i> ₂	0.135 (0.098)	-0.437* (0.228)	-0.027 (0.096)	0.037 (0.084)
<i>Slant</i> ₃	0.186* (0.112)	-0.447** (0.193)	-0.040 (0.093)	-0.006 (0.103)
	Fox	WSJ	NYT	WashPost
Panel B1: RHS = mean slant of other outlets (2016)				
<i>Slant</i> ₁	-0.091 (0.269)	-0.417 (0.781)	-0.367 (0.321)	0.224 (0.212)
<i>Slant</i> ₂	-0.198 (0.350)	-1.078* (0.635)	-0.787*** (0.291)	0.216 (0.247)
<i>Slant</i> ₃	-0.146 (0.315)	-1.206 (0.773)	-0.609** (0.273)	0.233 (0.278)
Panel B2: RHS = mean Republican poll advantage (2016)				
<i>Slant</i> ₁	0.204** (0.087)	-0.335* (0.175)	-0.155 (0.105)	0.107** (0.050)
<i>Slant</i> ₂	0.134 (0.085)	-0.363 (0.229)	-0.156 (0.118)	0.122* (0.069)
<i>Slant</i> ₃	0.091 (0.112)	-0.341** (0.165)	-0.128 (0.140)	0.076 (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 a 4th order polynomial for 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

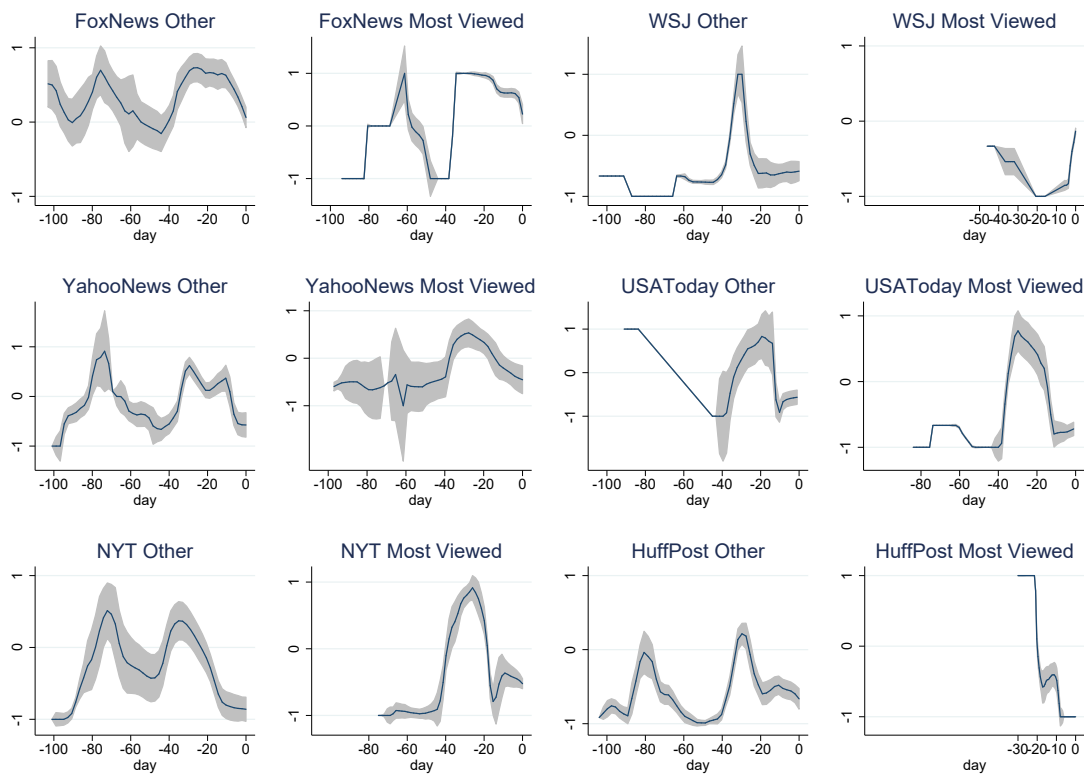
	$Slant_1$	$Slant_2$	$Slant_3$	$Slant_1$	$Slant_2$	$Slant_3$
Panel A: 2012						
Fox $\times Slant$	0.170** (0.083)	0.096 (0.103)	0.116 (0.135)	0.172** (0.085)	0.079 (0.104)	0.100 (0.143)
WSJ $\times Slant$	-0.010 (0.202)	0.024 (0.210)	0.001 (0.186)	-0.042 (0.195)	-0.003 (0.201)	-0.055 (0.180)
USAT $\times Slant$	-0.106 (0.088)	-0.087 (0.115)	-0.088 (0.111)	-0.110 (0.091)	-0.096 (0.120)	-0.118 (0.114)
Yahoo $\times Slant$	-0.043 (0.043)	-0.038 (0.053)	-0.027 (0.064)	-0.027 (0.040)	-0.023 (0.047)	-0.027 (0.056)
NYT $\times Slant$	-0.013 (0.072)	0.024 (0.076)	0.026 (0.087)	0.001 (0.071)	0.037 (0.076)	0.032 (0.085)
HuffPost $\times Slant$	0.010 (0.048)	0.029 (0.054)	0.021 (0.068)	0.020 (0.052)	0.035 (0.058)	0.009 (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* (0.069)	-0.189** (0.087)	-0.179* (0.094)	-0.128* (0.075)	-0.171* (0.099)	-0.176 (0.108)
WSJ $\times Slant$	-0.043 (0.244)	-0.187 (0.523)	0.389 (0.293)	-0.029 (0.247)	-0.091 (0.546)	0.463 (0.285)
Yahoo $\times Slant$	-0.021 (0.087)	-0.009 (0.116)	0.089 (0.141)	-0.034 (0.093)	-0.061 (0.137)	0.022 (0.209)
Google $\times Slant$	0.016 (0.029)	-0.026 (0.029)	0.009 (0.032)	0.023 (0.030)	-0.018 (0.032)	0.018 (0.032)
NYT $\times Slant$	0.193* (0.111)	0.284** (0.130)	0.511 (0.324)	0.194* (0.110)	0.314** (0.129)	0.539 (0.330)
WashPost $\times Slant$	-0.043 (0.056)	-0.013 (0.072)	0.019 (0.082)	-0.048 (0.057)	-0.046 (0.073)	-0.003 (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. 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.

A Appendix

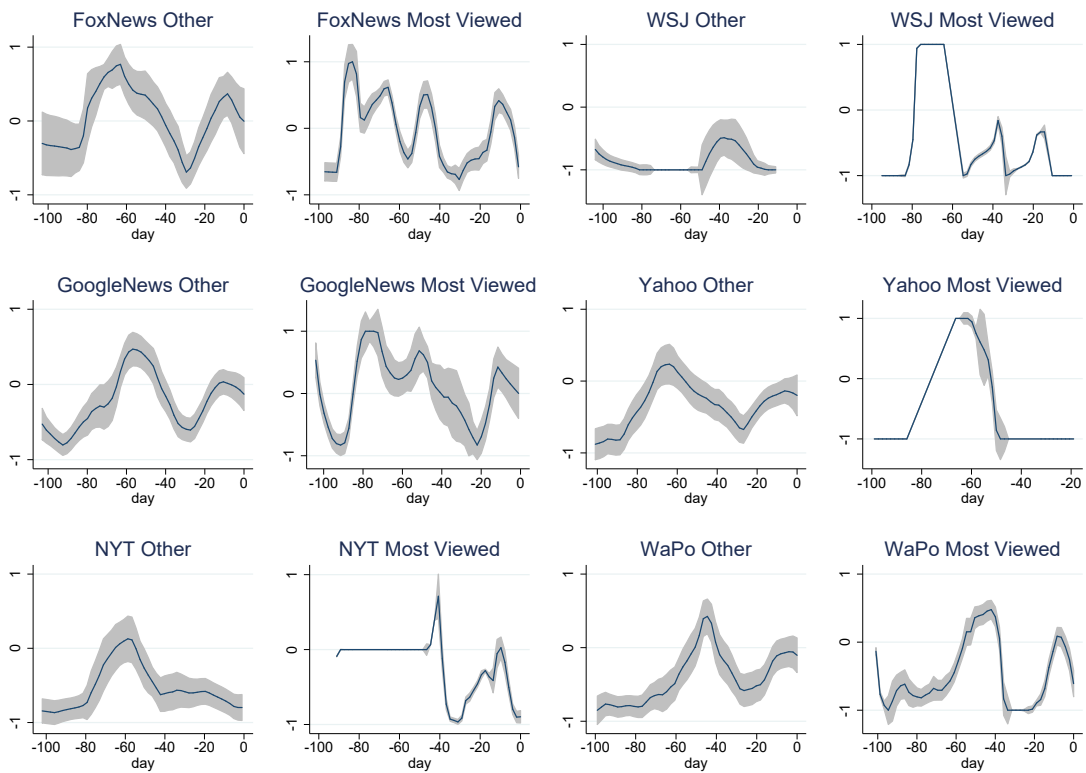
A.1 Additional figures and tables

Figure A1: Mean $Slant_1$ versus day relative to election day (day 0) by outlet in 2012



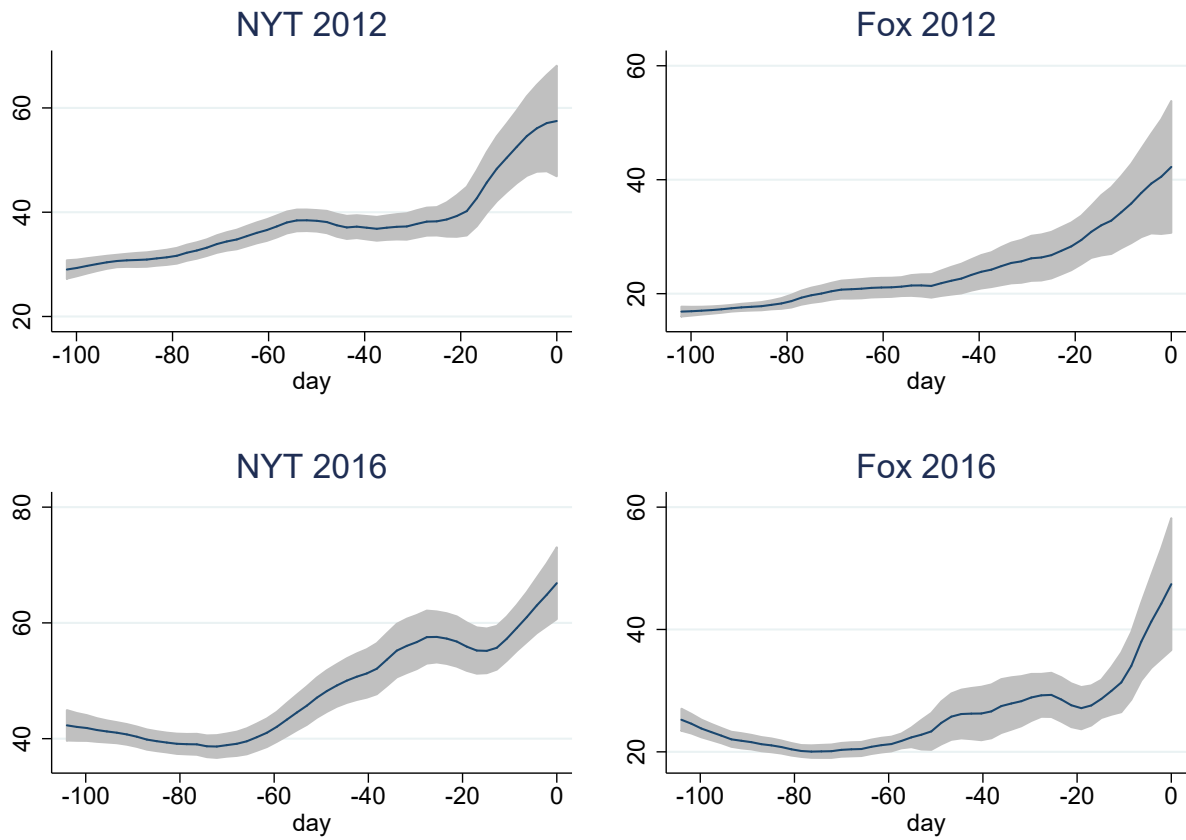
Note: Curves are kernel-weighted smoothed local polynomials with 95% confidence interval bands. Positive values of $Slant$ denote better chances of winning for the Republican candidate, whereas negative values indicate better chances for the Democratic candidate.

Figure A2: 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.

Figure A3: 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	<i>Slant</i> ₁	<i>Slant</i> ₂	<i>Slant</i> 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: Probability of story making most viewed list (date polynomial instead of day fixed effects)

	<i>Slant</i> ₁	<i>Slant</i> ₂	<i>Slant</i> ₃	<i>Slant</i> ₁	<i>Slant</i> ₂	<i>Slant</i> ₃
Panel A: 2012						
Fox \times <i>Slant</i>	0.029 (0.104)	-0.015 (0.113)	-0.020 (0.106)	0.013 (0.101)	-0.030 (0.112)	-0.030 (0.110)
WSJ \times <i>Slant</i>	-0.117 (0.203)	-0.114 (0.205)	-0.142 (0.174)	-0.125 (0.199)	-0.123 (0.199)	-0.155 (0.171)
USAT \times <i>Slant</i>	-0.080 (0.080)	-0.009 (0.099)	-0.020 (0.094)	-0.076 (0.081)	-0.005 (0.097)	-0.016 (0.092)
Yahoo \times <i>Slant</i>	-0.023 (0.037)	0.007 (0.046)	-0.001 (0.052)	-0.003 (0.035)	0.031 (0.042)	0.022 (0.049)
NYT \times <i>Slant</i>	-0.057 (0.062)	-0.032 (0.070)	-0.039 (0.070)	-0.041 (0.061)	-0.013 (0.069)	-0.018 (0.069)
HuffPost \times <i>Slant</i>	-0.002 (0.035)	0.022 (0.039)	0.015 (0.044)	0.011 (0.035)	0.031 (0.040)	0.023 (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 <i>Slant</i>	-0.124** (0.060)	-0.167*** (0.060)	-0.159** (0.061)	-0.122** (0.060)	-0.151** (0.059)	-0.150** (0.061)
WSJ \times <i>Slant</i>	-0.033 (0.174)	-0.099 (0.609)	0.676*** (0.187)	-0.040 (0.185)	-0.132 (0.598)	0.627*** (0.218)
Yahoo \times <i>Slant</i>	-0.026 (0.079)	0.066 (0.100)	0.111 (0.125)	-0.039 (0.079)	0.037 (0.104)	0.084 (0.131)
Google \times <i>Slant</i>	0.018 (0.028)	-0.005 (0.031)	-0.005 (0.032)	0.037 (0.026)	0.014 (0.030)	0.014 (0.032)
NYT \times <i>Slant</i>	0.212* (0.119)	0.328** (0.146)	0.259 (0.178)	0.196* (0.118)	0.321** (0.145)	0.248 (0.175)
WashPost \times <i>Slant</i>	-0.017 (0.050)	0.027 (0.068)	0.025 (0.072)	-0.027 (0.052)	0.011 (0.069)	0.022 (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.

A.2 Survey

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.

Figure A4 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.

A multinomial model would be technically the best choice to formally analyze the survey data, since respondents chose among four unordered alternatives. Results for such a model are available upon request. In the interest of simplicity and transparency, we limit results here to those of linear probability models predicting a binary outcome equal 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.²⁹

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 A4 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 ($\beta_{R,2} = 0.194$).

²⁹Our main goal was to obtain the desired sample size as quickly as possible after posting the survey so that the news would be as timely as possible and did not encourage respondents to take multiple surveys or restrict them from doing this.

Table A4: Likelihood of picking NYT and Fox debate winner article

	Y^{NYT}	Y^{Fox}	$Y^{NYT} + Y^{Fox}$
Panel A (Party Identity)			
Debate 2 \times Democrat ($\beta_{L,2}$, uncongenial)	0.004 (0.095)	-0.293*** (0.079)	-0.289*** (0.111)
Debate 3 \times Democrat ($\beta_{L,3}$, uncongenial for Fox only)	-0.023 (0.110)	-0.221** (0.100)	-0.243* (0.126)
Debate 2 \times Republican ($\beta_{R,2}$, congenial)	0.009 (0.101)	0.071 (0.104)	0.079 (0.128)
Debate 3 \times Republican ($\beta_{R,3}$, congenial)	-0.095 (0.113)	0.112 (0.121)	0.016 (0.140)
Adj. R^2	0.059	0.085	0.062
N	637	637	637
Panel B (Candidate supported)			
Debate 2 \times Clinton supporter ($\beta_{L,2}$, uncongenial)	0.039 (0.093)	-0.274*** (0.075)	-0.235** (0.106)
Debate 3 \times Clinton supporter ($\beta_{L,3}$, uncongenial Fox only)	0.122 (0.098)	-0.270*** (0.094)	-0.148 (0.116)
Debate 2 \times Trump supporter ($\beta_{R,2}$, congenial)	0.056 (0.100)	0.194* (0.103)	0.250** (0.122)
Debate 3 \times Trump supporter ($\beta_{R,3}$, congenial)	-0.069 (0.099)	0.186 (0.117)	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.3 Model

We present a slightly modified, and simplified, version of the model of [Mullainathan and Shleifer \(2005\)](#), using their notation to the extent possible. The model is intended to just illustrate and clarify a few important aspects of consumer choice in our setting, and does not capture all relevant factors.

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

Table A5: 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
	≥ 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.

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

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

with $\bar{u} > 0$, $\chi \geq 0$, and $\phi \geq 0$. 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). It is natural to assume both strict inequalities hold.

Before reading the story, the consumer must click on it. It is natural to define news demand as probability of a click (equivalent to expected number of clicks for a larger population of consumers). The consumer decides whether to click by first reading the story's headline, h ,

Table A6: 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.

after privately observing her private opportunity cost of the story, $c \sim U[0, \bar{c}]$. She thus clicks on the story to read it if $E(U_r|h) = \bar{u} - \chi E((\hat{s} + \epsilon)^2|h) + \phi(E(d|h) + \hat{s} + E(\epsilon|h)) \geq c$, and thus $G_c(E(U_r|h))$ is news demand conditional on h , with $G_c()$ being the CDF of c , if $E(U_r|h) \geq 0$. We assume that both of the consumer’s conditional expectations $E(d|h)$ and $E(\epsilon|h)$ are strictly increasing in h .

The outlet maximizes demand 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 , though we do not assume the outlet necessarily makes this choice.³⁰

³⁰The model could be extended so that in equilibrium the outlet does not necessarily maximize clicks (e.g., the outlet could balance clicks with a desire to influence consumers, or the outlet could simply have incorrect beliefs about consumer parameters). We avoid such extensions in the interest of simplicity and since they are unnecessary for our purposes.

Figure A4: 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.

The marginal effect of h on news demand is equal to:

$$\frac{\partial}{\partial h} G_c(E(U_r|h)) = g_c(E(U_r|h)) \left(-\chi \frac{\partial}{\partial h} E((\hat{s} + \epsilon)^2|h) + \phi \frac{\partial}{\partial h} (E(d|h) + E(\epsilon|h)) \right).$$

This expression is decreasing in \hat{s} , given $\chi > 0$.³¹ Thus, a larger mean slant, \hat{s} , moderates a positive relationship between congeniality of headline slant and demand, and could even flip the sign (from positive to negative) of this relationship. Intuitively, when \hat{s} is very large, readers are getting too much slant on average, so they are more likely to click headlines for stories that seem relatively unslanted. When \hat{s} is not so large, demand is increasing in h

³¹Note $g_c(\cdot)$ is a constant, and $-\chi \frac{\partial}{\partial h} E((\hat{s} + \epsilon)^2|h) = -\chi \frac{\partial}{\partial h} E(\hat{s}^2 + 2\hat{s}\epsilon + \epsilon^2|h) = -\chi \frac{\partial}{\partial h} E(2\hat{s}\epsilon + \epsilon^2|h)$, so the derivative of this expression with respect to \hat{s} is $-2\chi \frac{\partial}{\partial h} E(\epsilon|h)$, which is strictly negative since $\frac{\partial}{\partial h} E(\epsilon|h) > 0$.

because this is a signal of larger d .

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.