

# The Effect of Clickbait\*

Kevin Munger, Mario Luca, Jonathan Nagler, Joshua Tucker<sup>†</sup>

September 26, 2018

## Abstract

“Clickbait” has become a dominant form of online media, with headlines designed to entice people to click becoming the norm. The propensity to consume clickbait is not evenly distributed across relevant political demographics, so the present study presents the results of a pair of experiments: one conducted using Facebook ads that explicitly target people with a high preference for clickbait, the other using a sample recruited from Amazon’s Mechanical Turk. We estimate subjects’ individual-level preference for clickbait, and randomly assign some to read clickbait or traditional headlines. We find that older people and non-Democrats have a higher “preference for clickbait,” but find no evidence that assignment to read clickbait headlines drives affective polarization, information retention or trust in media.

---

\*This research is supported by the John S. and James L. Knight Foundation through a grant to Stanford University’s Project on Democracy and the Internet.

<sup>†</sup>Munger is an Assistant Professor in Political Science at Penn State University; email contact kevinmunger@gmail.com. Luca is a doctoral candidate at Sciences Po. Nagler and Tucker are Professors of Politics at New York University.

# 1 The Rise of Clickbait Media

Trust in the news media has been declining steadily ever since the 1970s (Ladd, 2011), especially among conservatives. This same time period has seen a rise in “affective polarization”—the extent to which Republicans and Democrats dislike and distrust each other (Iyengar, Sood, and Lelkes, 2012). Cultural and technological changes in the media environment have been theorized as causes of the latter trend, with an increasingly fragmented political news industry able to target niche political audiences (Stroud, 2011) and the increasing range of entertainment options and decline of incidental exposure to the nightly news pushing low-political-interest moderates out of the electorate (Arceneaux and Johnson, 2013; Prior, 2007).

Over the past twenty years, the news industry has increasingly become the online news industry. Figure 1 displays the striking decrease in employment by newspapers—from over 450,000 jobs in 1990 to under 200,000 in 2016—and the concomitant increase in jobs in “internet publishing and broadcasting.” Neither trend has significantly changed trajectory in the years since the great recession.

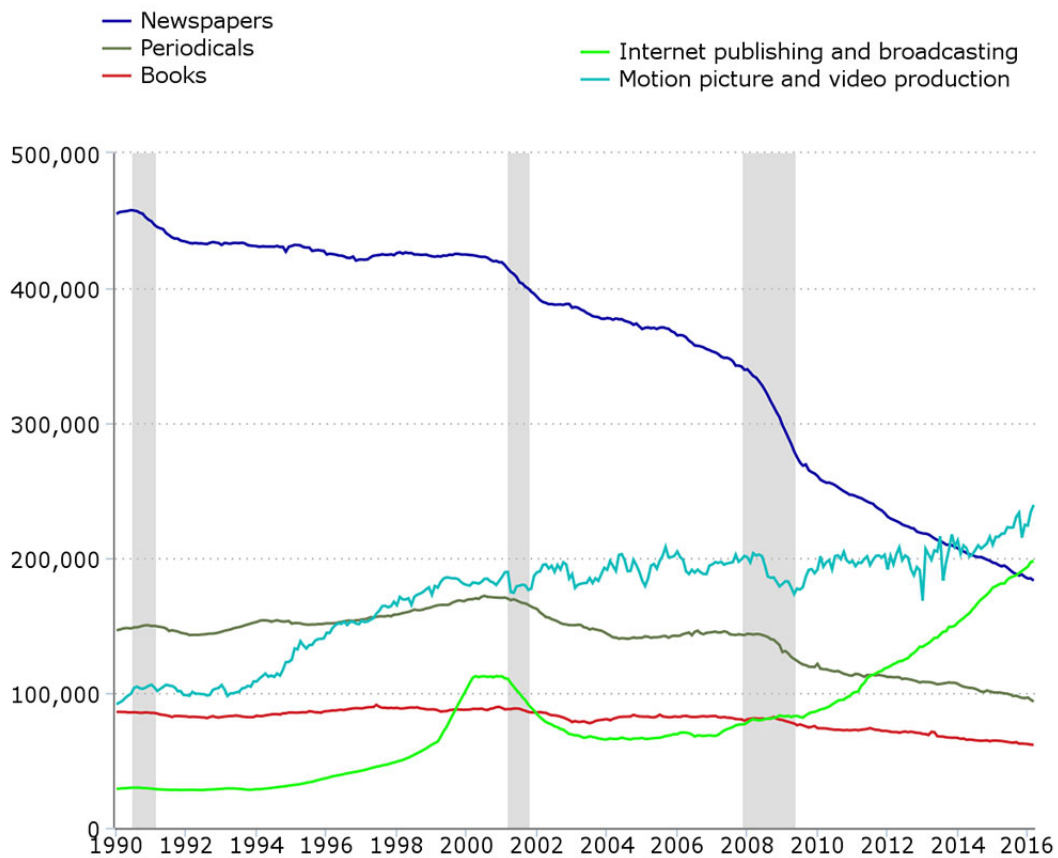
The economic model of the contemporary online news industry is distinct from print journalism. Although a small number of publications are financed by subscription revenue (the *New York Times* gets 60% of its revenue from subscriptions (Ember, 2018)), the primary business model is based on click-based advertisements. Competition comes from trying to attract readers’ eyeballs.

This business model was first embraced by *Gawker*, a brash new media firm which explicitly gamified the competition for clicks in 2008, paying writers bonuses based on pageviews and installing a “big board” that ranked its highest-performing stories in real time (Murtha, 2015). *Gawker* primarily aimed to attract clicks by writing salacious stories, but the concept of “clickbait” is better associated with another digital media upstart: *Upworthy*.

The “fastest-growing media site of all time,” *Upworthy* implemented a new style of headline designed to entice consumption by strategically withholding information (Sanders, 2017). Less than two years after its founding in March 2012 by Eli Pariser and Peter Koechley, *Upworthy* had over 80 million unique visitors each month—more than the *New York Times* or *Washington Post*. In November 2013, however, Facebook announced that it would penalize deceptive headlines in their ranking algorithm, and within a year, *Upworthy*’s business collapsed. In November 2014, the site had only 20 million unique visitors (Sanders, 2017).

Figure 1: Newspaper Employment Declines, Internet Publishing Employment Soars

**Employment in selected information industries, seasonally adjusted, 1990-2016**



Note: Shaded area represents recession, as determined by the National Bureau of Economic Research.  
Source: U.S. Bureau of Labor Statistics.

The following year, Merriam-Webster added “clickbait” to its dictionary, defining it as “something (such as a headline) designed to make readers want to click on a hyperlink especially when the link leads to content of dubious value or interest.”

This conception of clickbait thus has a negative connotation, characterized by something like regret—if a consumer of clickbait stopped to think about their decision to click on such a story, they would probably not do so.

This definition, however, is out of date; Facebook was able to detect this form of deceptive clickbait (by looking for instances in which users clicked a link and then quickly closed it) and penalize internet media firms that employed it.

A more concerning form of clickbait is one that appeals directly to people’s fears, especially as it relates to a threat to a social group to which they belong. We propose to define this type of clickbait headline as *emotional clickbait*: a headline which is designed to appeal directly and explicitly to the emotions of the reader. This form of clickbait serves the twin purposes of inducing excitement by appealing to group competition (Abramowitz and Saunders, 2006; Mason, 2018), and being easily spread among online social networks, which tend to be homophilous.

Emotional clickbait—when the news is about politics or the relevant social groups are politically relevant—is also more concerning to political scientists than the straightforwardly deceptive information gap clickbait headline because of its capacity to polarize and create separate epistemic communities. In the American context, these concerns manifest themselves as affective polarization between Republicans and Democrats (Iyengar, Sood, and Lelkes, 2012) and the potential for filter bubbles (Flaxman, Goel, and Rao, 2016). There is the additional concern that all forms of clickbait erode public trust in the news media. Since the heyday of broadcast journalism, the news media has consistently been shifting away from hard news and towards soft news or confrontational opinions on hard news in order to meet audience demand; this shift was accompanied by a general decrease in trust in news media (Ladd, 2011). On the other hand, this same shift caused increases in political knowledge: soft news attracted an audience that would have previously ignored news altogether (Baum, 2002), and cable news/opinion programming boosts factual knowledge retention by increasing audience arousal (Mutz, 2015).

Our hypotheses (pre-registered through EGAP, number 3175) were that random assignment to read stories with emotional clickbait headlines would exacerbate affective polarization, decrease trust in online news, and increase information retention; in every case, we observed null effects.

Our confidence in these null results is heightened because they were observed in two separate experiments; the experiments were identical, but conducted on two different online samples. The first sample was conducted on Amazon’s Mechanical Turk, a standard source of research subjects which has been shown to be generally reliable, producing experimental results which closely match results from nationally representative samples (Coppock, 2018; Mullinix et al., 2015; Snowberg and Yariv, 2018).

However, we theorize that age and digital literacy are two crucial moderators of online behavior like the propensity to consume clickbait news—web tracking data has demonstrated that they (in addition to ideology) strongly predicted the propensity to consume Fake News during the 2016 campaign (Guess, Nagler, and Tucker, 2018). These two variables do not sufficiently vary within the Mechanical Turk population (Brewer, Morris, and Piper, 2016; Huff and Tingley, 2015).

As Mullinix et al. (2015) argue in their influential paper on the generalizability of survey experiments, “If one has a well-developed theory about heterogeneous treatment effects, then convenience samples only become problematic when there is a lack of variance on the predicted moderator (p22).” The Mechanical Turk sample contains very few people over 65, and, structurally, it *cannot* contain individuals below a certain threshold of digital literacy (Brewer, Morris, and Piper, 2016).

We first provide evidence that certain types of people are more likely to select an emotional clickbait headline when given the opportunity. Although we did not have a strong theoretical expectation *ex ante*, there is robust evidence that older people and non-Democrats have a higher preference for clickbait. On other characteristics of interest, however, there are significant differences between the two samples, illustrating the importance of thoughtful sample selection when conducting research on digital media effects. The results suggest that the most important pathway by which clickbait could affect political outcomes is by changing *which* or *how many* news stories people consume. However, the magnitude of respondents’ preference for ideologically congruent headlines was much larger than any of these effects, suggesting that the effects of clickbait can be muted in a strongly partisan media environment.

We hypothesized (pre-registered through EGAP, number 3175) that random assignment to read stories with emotional clickbait headlines would exacerbate affective polarization, decrease trust in online news, and increase information retention; in every case, we observed null effects. Below, we discuss possible explanations for these null effects and propose new research designs which could help determine the robustness of these null effects.

## 2 Experiments on the Determinants and Effects of Clickbait News Consumption

We begin by describing our research design. We conducted two related survey experiments, one using Amazon’s Mechanical Turk and the other using subjects recruited using Facebook advertisements.

The survey instrument was designed to take around ten minutes to complete, and contained an attention check and built-in delays to discourage respondents from giving low-quality answers.

The Mechanical Turk sample consisted of 2,803 total respondents across three slightly different experimental setups; in each case, because the pool of MTurk workers contains more Democrats than Republicans, we supplemented the first draw with a sample of self-reported Conservatives.<sup>1</sup> This sample was intended to serve as the baseline, as Mechanical Turk is a standard source of subjects for online survey experiments. Each Mechanical Turk subject was compensated \$1.

The first experimental setup ( $N = 1,140$ ) only allowed us to collect demographic information and non-experimental preference for clickbait questions because the experimental manipulation was unfortunately marred by a design flaw, rendering any inference from the experimental manipulations invalid. The second experimental setup ( $N = 826$ ) fixed the design flaw and allows us to draw the correct inferences. The final experimental setup ( $N = 837$ ) replicates the correct experimental manipulation but drops the pre-treatment preference for clickbait questions; we performed this analysis to check whether this portion of the instrument was dampening treatment effects.

The Facebook sample was recruited through a Facebook advertising campaign run by a Facebook page we created. We paid for an advertisement to appear on the News Feed of our potential subjects. The structure of Facebook’s advertising platform meant that we only pay when a potential subject actually clicks on the ad. The overall cost paid to Facebook for the subject recruitment was \$1,858 for 2,766 subjects who clicked on our ad. We compensated subjects by entering them (the 1,232 who completed the survey) into a lottery to win a \$500 gift Amazon gift card, meaning that the overall cost per subject was \$0.85 for subjects who began the survey and \$1.91 for subjects who completed the survey. The advertisement we used is displayed in Figure 2.

The motivation for the lottery (and the design of the recruitment instrument) was

---

<sup>1</sup>Mechanical Turk allows requesters the ability to specify the demographics of a given sample, including their ideological leaning.

Figure 2: Recruitment Instrument for Facebook Sample

The image shows a screenshot of a Facebook page for 'NYU Survey'. The page header includes the Facebook logo, the name 'NYU Survey', and a search icon. Below the header are navigation tabs: 'Page', 'Inbox', 'Notifications', 'Insights', and 'Publishing Tools'. The left sidebar contains the page's profile picture (a purple square with 'NYU SURVEY' in white), the name 'NYU Survey', and options to 'Create Page @Username', 'Home', 'About', 'Events', 'See more', 'Promote', and 'Manage Promotions'. The main content area shows a post from 'NYU Survey' with the text: 'Win up to 500\$ with our 5 minutes survey. Give us your opinion and get a chance to win an Amazon voucher of 500\$! (more info on the website)'. Below the text is a large purple graphic with white text that reads 'WIN UP TO 500\$ WITH 5 MIN SURVEY'. Underneath the graphic is the text 'Online Survey Software | Qualtrics Survey Solutions', a short description of Qualtrics software, and the website 'NYU.QUALTRICS.COM'. At the top of the post area are buttons for 'Like', 'Follow', 'Share', and a menu icon.

Table 1: Summary Statistics of MTurk and Facebook Samples

	MTurk	Facebook
% Female	46%	75%
Mean Age	37	49
75th Percentile Age	43	63
% Finished College	58%	42%
% Republican	33%	21%
% Independent	29%	28%
% Internet > 1/day	96%	93%
% Facebook > 1/day	52%	90%
N	1,903	2,382

twofold. First, one distinct advantage of Mechanical Turk over Facebook for subject recruitment is the former’s built-in system for processing microtransactions. The need to perform an individual \$1 transaction for each subject would have represented a significant additional cost for the experiment.

More substantively, the advertisement was designed to be as eye-catching as possible. Facebook ads can be used with quota sampling to generate valid measures of public opinion (Zhang et al., 2018), but we were particularly interested in a non-representative sample of Facebook users: people who were most likely to click on an eye-catching ad.

The Facebook sample, then, is unbalanced on a number of important dimensions. Table 1 provides the descriptive statistics of the two samples. Some of the distributions are striking; in particular, a full 75% percent of the Facebook sample were women.

However, we cannot be sure whether this 3-1 gender ratio reflects the true rate at which people clicked on ads because there is some uncertainty about the way that the Facebook advertising software operates. As Zhang et al. (2018) points out, Facebook uses a multi-armed bandit algorithm to optimize the efficiency of ad buys throughout the duration of their run. For example, after detecting that women are slightly more likely to click the ad than are men, the algorithm would start displaying the ad to more women.

This does not seem to be driving results in this case, as the proportion of women in the beginning and end of the ad run are identical. Even so, given the opacity of the Facebook algorithm, we should not read the proportions from the Facebook sample as necessarily reflecting the true population of people who might have clicked on the ad.



In each experiment, respondents were directed to a Qualtrics survey in which they first reported demographic information, including partisan affiliation and their frequency of internet/Twitter/Facebook use. We then gave them a series of nine tasks, one of which was an attention check. In each task they were shown four headlines, and asked which they would most like to read. Note that respondents were not actually given links to these stories nor asked to actually read the stories at this point. In each task, there were two political stories (one Democrat-favorable, one Republican-favorable) and two non-political stories (one sports, one entertainment).<sup>2</sup> One of the two political headlines (either the Democrat-favorable or Republican-favorable) in each decision set was turned into a clickbait headline through the addition of an attention-grabbing phrase to the beginning, so that there were four instances in which the Democrat-favorable headline was clickbait and four instances in which the Republican-favorable was clickbait. In the task set up as an attention check, one of the four answers read “Survey taker: always select this option, ignore the other choices.”

The purpose of this part of the survey was to calculate individual-level *preference for clickbait* (PfCB): how often each respondent claimed they would prefer to read the clickbait headline rather than the non-clickbait headline, ignoring the preference for non-political headlines.<sup>3</sup> This process was non-experimental; our goal was to see how PfCB varied across respondent demographics, and to see how the experimental treatment effect (described below) varied with individual-level PfCB.

Respondents were then randomly assigned to one of four treatment conditions plus a fifth “placebo” condition (in which respondents were given a story about sports) through a 2x2 treatment design that varied the partisan leaning of a headline and whether the headline was clickbait. In each case, respondents were presented with a hyperlinked headline; when they clicked the headline, they were directed to a separate tab which displayed the given headline and a news story. The text of the news story was held constant across the conditions.

After respondents read the story and closed the tab, they were asked a feeling thermometer question<sup>4</sup> about Republicans, Democrats, online media and traditional

---

<sup>2</sup>The inclusion of non-political stories has been shown by Arceneaux and Johnson (2013) to provide more reliable estimates of media choice.

<sup>3</sup>This process balances the concern expressed in Leeper (2016) for measuring media treatment effects on the relevant population (those who would actually consume the given media) with the fact that we needed to disguise the nature of the manipulation from the respondent.

<sup>4</sup>These questions ask respondents to rate how they feel about the respective groups on a scale from 0 to 100, where 0 is “Very cold or unfavorable feeling” and 100 is “Very warm or favorable feeling.” This question has a long history of use in the ANES, and is the standard measure of affective polarization.

Table 2: Treatment Headlines

R: Baseline	Trump economic policies working
R: Clickbait	Democrats won't like this economic news: Trump policies working!
D: Baseline	Trump economic policies not working
D: Clickbait	Republicans won't like this economic news: Trump policies not working!

media, as well as a multiple-choice question about their trust in online media and traditional media. On the next page, they were asked three multiple-choice questions based on facts presented in story they had been given to read plus an additional, placebo factual question about the sports story.

The text of the story used in this experimental manipulation summarized the findings from the October jobs report and was taken from a politically neutral news source: CNN Money.

The treatment headlines for the experiments are displayed in Table 2.

The intention was to design headlines that would anger the respective partisan groups, then amplify that anger through the emotional clickbait introduction. The literature on affective polarization is still being developed, but a well-established trend is that the gap in partisan affect is driven by decreased evaluations of the out-party. This is what Abramowitz and Webster (2016) call “negative partisanship”—out-partisan animosity is a powerful motivator for a range of political behaviors. Mason (2016) finds experimental support for the presence of anger in response to partisan threats.

The headlines displayed in Table 2 are symmetric, adding only a “not” to switch the partisan leaning. The emotion appealed to in this version of emotional clickbait is—quite explicitly—negative partisan excitement: the idea that your opponents being angry about something implies that you will excited by it (Abramowitz and Saunders, 2006).

### 3 Hypotheses

The first question this study aims to answer is exploratory: what kinds of people are more likely to consume *emotional clickbait*? There is not any strong theory here, so the analysis related to this research question will be descriptive rather than confirmatory.

**Research Question:** *What kinds of people are more likely to consume emotional clickbait?*

The hypotheses about the effects of being randomly assigned to the clickbait treatment conditions below did not vary across platforms. The R file containing all of the code used to analyze the experimental data was included with our EGAP pre-registration (number 3175) as part of the pre-analysis plan. In addition to pre-registering the hypotheses listed below, we specified in advance the precise coding and data manipulation decisions we would make in testing those hypotheses. The specific language of the hypotheses has been changed to match the terminology in the rest of this paper.

**Hypothesis 1** *The Republican-favorable conditions will **increase** reported affect toward Republicans. The Democrat-favorable conditions will **decrease** reported affect toward Republicans.*

In both experiments, the Republican president is the primary political actor mentioned in the headline. Hypothesis 1 thus predicts a change in the way that subjects feel about the Republican party, in the direction of the frame of that headline.

**Hypothesis 2** *The effects predicted in Hypothesis 1 will be larger in the emotional clickbait than the baseline conditions.*

Hypothesis 2 predicts that the addition of the *emotional clickbait* language to the beginning of the partisan headlines will amplify the effects of those frames.

**Hypothesis 3** *Respondents assigned to read emotional clickbait will report lower trust in (online) media.*

Following the findings in Mutz (2015) on the impact of incivil cable news, Hypothesis 3 predicts that respondents who are assigned to click on a story with a clickbait headline will decrease their trust in media (either just online media or both online and offline media), but Hypothesis 4 predicts that this will be accompanied by an increase in respondents' ability to recall specific facts from that news story. Incivility in cable news is theorized to cause decreased trust in media and increased information retention through the mechanism of increased arousal, which also expect to obtain in the context of emotional clickbait.

**Hypothesis 4** *Respondents assigned to read emotional clickbait will retain more information.*

The final hypothesis was not pre-registered; the null results found in the experiment conducted on the MTurk subjects motivated us to think about the limitations of that sample. We hypothesized that the MTurk sample did not sufficiently vary in a potential treatment moderator (digital literacy), and we sought out the Facebook sample in order to test this theory.

**Hypothesis 5** *The treatment effects in Hypotheses 1-4 will be larger among subjects with lower levels of digital literacy and thus larger among the Facebook sample than among the MTurk sample.*

## 4 Results

### 4.1 Preference for Clickbait

To analyze the individual-level preference for clickbait (PfCB), we compare the results from the MTurk and Facebook experiments. PfCB is calculated by estimating what percentage of the political stories the subjects selected to read were clickbait. We also estimate the individual-level preference for Republican (PfR) news. Note that these two quantities are structurally (negatively) correlated: an individual who selected 8 out of 8 clickbait stories would necessarily have selected 4 out of 8 Republican stories.

Partisans made the expected choices: the mean PfR was .61 for Republicans and .36 for Democrats, including leaners, in both samples. For Republicans (including leaners), the mean PfR was .60 in the Facebook sample and .64 in the MTurk sample. Restricting to strong partisans heightens these trends only slightly. This preference for co-attitudinal news restricts the range of possible values for PfCB.

Still, the overall results were surprising: the overall PfCB was negligible. The rate of selecting the clickbait political stories was in fact slightly lower than the non-clickbait political stories (median PfCB = .50, mean PfCB = .47); this rate did not vary across the samples.

In addition to the restricted range discussed above, this result illuminates a limitation of the current research design: respondents were acutely aware that they were taking part in a study, and they may have made less impulsive choice than they would have in a more naturalistic setting.<sup>5</sup>

---

<sup>5</sup>Furthermore, we do not take these results as evidence that clickbait “doesn’t work.” Dozens of competing media firms have in effect demonstrated that clickbait does work by adopting it as a

With this caveat, though, we can still estimate the subjects’ *relative* PfCB. Table 3, columns 1 (Mturk) and 2 (Facebook), displays the results of an OLS regression taking PfCB as the dependent variable and all of the demographic information collected from users as independent variables. Questions about frequency of Facebook use, Twitter use, Internet use, reading online news stories and reading offline news stories are on an 8-point increasing categorical scale.

Across both samples, the only consistent results predicting PfCB are that older individuals have a higher PfCB and Democrats have a lower PfCB.

Many of the other coefficients are estimated to have significant effects in one sample and negligible effects in the other. More frequent Facebook users have a significantly higher PfCB in the FB sample, while frequent Twitter users have a significantly higher PfCB in the MTurk sample. More frequent consumers of offline news have a higher PfCB in the MTurk sample, but more frequent consumers of online news have a higher PfCB in the Facebook sample.

The most striking result that appears only in the Mechanical Turk sample is the strong but non-monotonic effect of Republican identification on PfCB: all Republicans have a higher PfCB than do Independents, but this trend is especially pronounced among Republican leaners, whose additional PfCB (relative to Independents) is more than double that of solid Republicans.

The explanation for this non-monotonicity can be found in column 3. Partisan ID has the expected results on preference for Republican (PfR) stories: as reported PID moves from strong Democrat to strong Republican, the PfR increases monotonically. Because the PfR is so strong among non-leaner Democrats and Republicans, these respondents have no “degrees of freedom” left to choose between clickbait and non-clickbait stories. Notice that none of the media use variables have an effect on PfR—except for a marginally significant reduction in PfR among frequent offline news consumers from the Facebook sample.

Subjects recruited via Facebook have higher values for PfCB when all of the covariates take the value of 0. However, it is possible that these factors behave differently in subjects in the different samples, so Table 5 in Appendix A combines these two samples analyzed separately in Table 3 with models that fully interacted with a dummy for which sample a subject was drawn from.

---

prominent format for news headlines, sometimes using explicit A/B testing. There does not yet exist reliable descriptive estimates of the prominence of clickbait relative to traditional headlines, however.

Table 3: Preference for Clickbait

	<i>Dependent variable:</i>			
	Preference for Clickbait		Preference for Republican	
	MTurk	FB	MTurk	FB
Facebook use	0.003 (0.003)	0.015*** (0.006)	0.004 (0.003)	0.007 (0.007)
Twitter use	0.003 (0.003)	0.002 (0.002)	-0.001 (0.003)	-0.001 (0.002)
Internet use	0.009 (0.008)	-0.009* (0.005)	0.002 (0.009)	0.009 (0.006)
Age	0.001* (0.0005)	0.001*** (0.0003)	0.001 (0.001)	-0.0001 (0.0003)
Education	-0.005 (0.007)	-0.014*** (0.004)	-0.001 (0.008)	-0.007 (0.005)
Offline news consumption	0.009** (0.004)	0.001 (0.002)	-0.0001 (0.004)	-0.001 (0.002)
Offline news consumption	0.006 (0.005)	0.005* (0.003)	-0.002 (0.005)	-0.012*** (0.003)
Democrat	-0.052*** (0.016)	-0.020* (0.010)	-0.129*** (0.018)	-0.096*** (0.012)
Lean Democrat	-0.032* (0.019)	-0.032** (0.015)	-0.082*** (0.021)	-0.090*** (0.018)
Lean Republican	0.067*** (0.018)	0.012 (0.017)	0.112*** (0.020)	0.103*** (0.020)
Republican	0.036** (0.018)	0.021 (0.014)	0.135*** (0.020)	0.165*** (0.016)
Constant	0.308*** (0.063)	0.364*** (0.046)	0.450*** (0.071)	0.448*** (0.053)
Observations	1,889	2,256	1,889	2,256
Adjusted R <sup>2</sup>	0.033	0.029	0.117	0.153

*Note:*

\*p&lt;0.1; \*\*p&lt;0.05; \*\*\*p&lt;0.01

## 4.2 Pre-Registered Experimental Results

Turning to the results from the experimental condition, we want to be explicit about our central result: following the analysis code that we pre-registered in our pre-analysis plan, we found null results. After performing the appropriate Bonferroni correction to account for multiple comparisons (and, in fact, in almost every case without this correction), we estimate that each of the hypothesized treatment effects on the relevant outcome variable was significantly indistinguishable from zero.

There is evidence that the reason the treatments did not have the hypothesized effects is not due to a lack of uptake from the factual knowledge questions. Figure 3 displays these results. There were three information retention questions in addition to a sports-related placebo information retention question.<sup>6</sup>

In both the Mechanical Turk and Facebook samples, subjects in the placebo condition answered far fewer questions correctly. This is evidence that subjects were in fact reading the stories carefully and retaining the information, rather than relying on their *ex ante* knowledge.

## 4.3 Post Hoc Analysis of Experimental Results

To further validate that the experimental manipulation was not entirely ignored, and to learn as much as possible from the data, we present the results from post hoc modifications to our pre-registered analysis plan. The primary modification is to interact the treatments with the party identification (on a five-point scale) of the respondents.

Table 4 presents two models that estimate the effects of our four treatment conditions interacted with the party identification of the respondents on affect towards the Republican party, as measured with the 0-100 feeling thermometer.<sup>7</sup> The straightforward effect of party ID dominates, as expected, but in the model using the Mechanical Turk sample, we find the two Republican-leaning headlines cause a significant reduction in warm feelings toward the Republican party; the two Democrat-leaning headlines cause a non-significant reduction.


---

<sup>6</sup>All of the information retention questions were based on information provided in the body of the treatment news story. Because these were taken from existing news stories based on recent political news stories, it is possible that respondents could have known the correct answers *ex ante*. To minimize this problem, the questions concerned specific details from the stories that were not particularly salient to the overall political discussion.


<sup>7</sup>We do not present a similar breakdown with models of the other recorded dependent variables of interest (trust in media and factual information retention) because there are no significant results.

Figure 3: Effects of Clickbait on Information Retention

### MTurk



### Facebook



Bars represent 95% confidence intervals.



Table 4: Treatments Interacted with Party ID

	<i>Republican Feeling Thermometer</i>	
	(Mturk)	(Facebook)
Shortened Survey	4.255*** (1.258)	
Dem CB	-3.692 (3.788)	-4.423 (3.646)
Dem non CB	-5.107 (3.743)	-1.939 (3.734)
Rep CB	-7.711** (3.664)	2.345 (3.487)
Rep non CB	-7.820** (3.593)	-4.743 (3.601)
Democrat	-22.476*** (3.812)	-18.664*** (3.334)
Lean Democrat	-15.184*** (4.497)	-10.067** (5.084)
Lean Republican	11.564*** (4.379)	18.626*** (5.706)
Republican	24.272*** (4.530)	41.889*** (4.728)
Dem CB X Democrat	2.814 (5.444)	6.898 (4.801)
Dem non CB X Democrat	4.393 (5.335)	1.292 (4.904)
Rep CB X Democrat	3.500 (5.332)	-0.124 (4.693)
Rep non CB X Democrat	9.971* (5.356)	8.403* (4.745)
Dem CB X Lean Democrat	0.439 (6.923)	3.016 (7.295)
Dem non CB X Lean Democrat	0.502 (6.423)	-4.913 (7.235)
Rep CB X Lean Democrat	7.784 (6.249)	-6.994 (7.016)
Rep non CB X Lean Democrat	4.766 (6.684)	2.290 (7.465)
Dem CB X Lean Republican	1.579 (6.345)	16.232* (9.291)
Dem non CB X Lean Republican	13.020** (6.346)	9.794 (7.964)
Rep CB X Lean Republican	12.840** (6.045)	-2.184 (7.785)
Rep non CB X Lean Republican	11.056* (6.210)	7.024 (7.973)
Dem CB X Republican	9.881 (6.346)	-2.072 (6.530)
Dem non CB X Republican	12.913** (6.232)	-2.817 (6.940)
Rep CB X Republican	15.002** (6.082)	-9.679 (6.566)
Rep non CB X Republican	13.425** (6.145)	-2.510 (6.563)
Constant	46.693*** (2.690)	39.474*** (2.578)
Observations	1,608	1,303
R <sup>2</sup>	0.377	0.404

These reductions are more than counterbalanced by the positive and significant interaction terms when the treatment effects are estimated on Republicans and Republican leaners. The ignored category in the party ID variable is Independents, so in the aggregate we can summarize the results of Table 4: the Republican-leaning treatment conditions caused a reduction in warm feelings toward the Republican party among Independents, had no effect on Democrats, and increased warm feelings among Republicans.

We do thus find support for Hypothesis 1, although this is a straightforward replication of the large literature on partisan priming. The fact that this result replicates in our dataset should increase our confidence that the null results in comparing clickbait and non-clickbait headlines are legitimate estimates of the true effect of clickbait.

However, we only find the significant results of partisan cues in the Mechanical Turk sample. One plausible explanation for our failure to replicate these results on the Facebook sample is that there was massive attrition from this sample at the stage of assigning treatment, leaving us with a non-random half of the original sample.

The experimental treatment involved clicking on a hyperlinked headline that opened up the news story in a separate tab. At this point in the survey, 6% of the Mechanical Turk sample dropped out, compared to 31% of the Facebook sample. This dramatically reduces the statistical power to detect any treatment effects in the Facebook case, but because this attrition is non-random, it could also produce biased estimates of treatment effects.

In particular, much of our motivation for using the Facebook sample to confirm our null results from the Mechanical Turk sample was that we were concerned that the latter did not contain enough people with low digital literacy, and that these are precisely the people who should be most likely to consume and be affected by emotional clickbait.

As a result, we cannot be sure that the treatment effects estimated on the sample of Facebook users who completed the survey generalize to *any* relevant population: to the population of “people who click on Facebook ads” from which the entire sample was drawn, or to the population of Facebook users (Coppock, 2018).

## 5 Conclusion

Clickbait news media is here to stay. Although Facebook and other online platforms try to combat deceptive or regret-causing headlines to improve the experience of their users, the fundamental economic dynamic of the social media feed is that it incentivizes media companies to compete on the level of individual stories. Crafting attention-grabbing headlines is essential: with a near-infinite amount of news content available in every users' feed, media companies need to make readers choose their stories. With the proliferation of online media outlets enabled by the reduced cost of producing news content, one strategy has been to create headlines that appeal directly to readers' identities via the mechanism of emotional arousal.

In the context of politics, these identities tend to be partisan. “Partisan emotional clickbait” headlines include explicit cues about how partisans should feel in response to a given piece of news. This phenomenon has become more relevant with the increasing alignment of partisan identities with other social identities; partisans tend to experience political news in terms of it being good or bad for their party, and partisan media reinforces this tendency (Mason, 2018).

We hypothesized that random assignment to read an emotional clickbait headline would provide evidence that this media trend might be able to explain some of the worrying trends among American partisans. The concerns about sample attrition discussed above notwithstanding, the experimental results presented in this paper failed to provide evidence for the hypotheses that emotional clickbait has direct effects on affective polarization, information retention or trust in media. We have complete data from 1,608 Mechanical Turk respondents (half of whom participated in a shortened version of the survey out of concern for dampened treatment effects) and 1,303 Facebook respondents. Using our pre-registered code to analyze this data, we find null results. Even a post hoc analysis of the data which confirms the replication of the well-known effects of partisan cues fails to find evidence of the effect of clickbait headlines.

The experimental manipulation we implemented was relatively small—we only changed the headlines of the news story, keeping the text of the story constant—which should bias against finding treatment effects of clickbait. It is also possible that a larger (and in some ways more realistic) experimental setup that varied the body of the story to mirror the tenor of the headline might find a larger effect.

Our non-experimental analysis, however, allows for another pathway by which clickbait could affect American politics: by differentially changing the media diets of different

types of social media users. We find evidence of heterogeneously distributed preference for clickbait; in both of our samples, older respondents scored higher on this dimension while Democrats scored lower. We also found some coefficients that were significantly associated with PfCB in one sample but not the other.

We argue that the Facebook sample is more relevant because it is drawn from the population of interest: people opted into the sample because they clicked on our recruitment ad on Facebook. Among this sample, there are two particularly strong relationships: more frequent Facebook users have a higher PfCB, while more educated respondents have a lower PfCB.

The upshot of these descriptive findings is that the behavioral/attitudinal impact of social media use is always and everywhere *heterogeneous*.

The variable heterogeneity of the effects of different media technologies is well established in political science. The clearest example comes from Prior (2007), who conceptualizes two populations of television consumers: those with a high preference for entertainment (PfE), who will always chose to watch non-news programs, and those with low PfE. In the broadcast era, these groups were indistinguishable because of the lack of choice among the three broadcast providers. Broadcast television thus had a relatively *homogeneous* effect on viewers' political attitudes and information levels. With the advent of cable television, however, people with high PfE avoided news programs. The effect of cable television viewing was thus *heterogeneous* in the viewer's PfE; cable television led to a more polarized electorate as moderates became less politically engaged.

Changing the number of images simultaneously possible to view from 3 to 50 (broadcast to cable television) increased the heterogeneity of the effects of television. The internet and social media have made that number of possible images essentially infinite; you can never step in the same News Feed twice.

Heterogeneity should thus be central to any study of media or persuasive effects on social media. Average treatment effects on a representative population might well be expected to be zero.

There is growing evidence for this view in the context of political engagement. Using web-tracking data, Guess, Nyhan, and Reifler (2017) conclude that the average number of times US internet users viewed Fake News during the 2016 election was quite low. This average masks the fact that “almost six in ten visits to Fake News websites came from the 10% of Americans with the most conservative information diets.” It should be noted that this observational “preference for Fake News” finding closely mirrors our

results about the preference for clickbait.

## References

- Abramowitz, Alan I, and Kyle L Saunders. 2006. “Exploring the bases of partisanship in the American electorate: Social identity vs. ideology.” *Political Research Quarterly* 59 (2): 175–187.
- Abramowitz, Alan I, and Steven Webster. 2016. “The rise of negative partisanship and the nationalization of US elections in the 21st century.” *Electoral Studies* 41: 12–22.
- Arceneaux, Kevin, and Martin Johnson. 2013. *Changing minds or changing channels?: Partisan news in an age of choice*. University of Chicago Press.
- Baum, Matthew A. 2002. “Sex, lies, and war: How soft news brings foreign policy to the inattentive public.” *American Political Science Review* 96 (1): 91–109.
- Brewer, Robin, Meredith Ringel Morris, and Anne Marie Piper. 2016. Why would anybody do this?: Understanding older adults’ motivations and challenges in crowd work. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*. ACM pp. 2246–2257.
- Coppock, Alexander. 2018. “Generalizing from survey experiments conducted on mechanical Turk: A replication approach.” *Political Science Research and Methods* pp. 1–16.
- Ember, Sydney. 2018. *New York Times Co. Subscription Revenue Surpassed 1 Billion in 2017*. New York Times.
- Flaxman, Seth, Sharad Goel, and Justin M Rao. 2016. “Filter bubbles, echo chambers, and online news consumption.” *Public opinion quarterly* 80 (S1): 298–320.
- Guess, Andrew, Brendan Nyhan, and Jason Reifler. 2017. “Inside the Fake News Bubble? Consumption of online fake news in the 2016 U.S. election.” *Working Paper* .
- Guess, Andrew, Jonathan Nagler, and Joshua A. Tucker. 2018. “Who’s Clogging Your Facebook Feed? Ideology and Age as Predictors of Fake News Dissemination During the 2016 U.S. Campaign.” *Unpublished manuscript* .

- Huff, Connor, and Dustin Tingley. 2015. “Who Are These People?” Evaluating the Demographic Characteristics and Political Preferences of MTurk Survey Respondents.” *Research and Politics* 2 (1): 1–12.
- Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. “Affect, not ideology a social identity perspective on polarization.” *Public opinion quarterly* 76 (3): 405–431.
- Ladd, Jonathan M. 2011. *Why Americans hate the media and how it matters*. Princeton University Press.
- Leeper, Thomas J. 2016. “How does treatment self-selection affect inferences about political communication?” *Journal of Experimental Political Science* .
- Mason, Lilliana. 2016. “A Cross-Cutting Calm How Social Sorting Drives Affective Polarization.” *Public Opinion Quarterly* p. nfw001.
- Mason, Lilliana. 2018. *Uncivil agreement: How politics became our identity*. University of Chicago Press.
- Mullinix, Kevin J, Thomas J Leeper, James N Druckman, and Jeremy Freese. 2015. “The generalizability of survey experiments.” *Journal of Experimental Political Science* 2 (2): 109–138.
- Murtha, Jack. 2015. *What its like to get paid for clicks*. Columbia Journalism Review.
- Mutz, Diana C. 2015. *In-your-face politics: The consequences of uncivil media*. Princeton University Press.
- Prior, Markus. 2007. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.
- Sanders, Sam. 2017. “Upworthy Was One Of The Hottest Sites Ever. You Won’t Believe What Happened Next.” *NPR* June 20, 2017.
- Snowberg, Erik, and Leeat Yariv. 2018. “Testing the Waters: Behavior across Participant Pools.” .
- Stroud, Natalie Jomini. 2011. *Niche news: The politics of news choice*. Oxford University Press on Demand.

Zhang, Baobao, Matto Mildenerger, Peter D. Howe, Jennifer Marlon, Seth Rosenthal, and Anthony Leiserowitz. 2018. "Quota Sampling Using Facebook Advertisements." *Political Science Research and Methods* pp. 1–16.



## A Fully Interacted Models

Table 5: Preference for Clickbait: Fully Interacted Models

	<i>Dependent variable:</i>	
	PfCB	PfR
Facebook sample	0.057 (0.077)	-0.002 (0.087)
Facebook use	0.003 (0.003)	0.004 (0.003)
Twitter use	0.003 (0.002)	-0.001 (0.003)
Internet use	0.009 (0.007)	0.002 (0.008)
Age	0.001** (0.0004)	0.001 (0.0005)
Education	-0.005 (0.006)	-0.001 (0.007)
Offline news consumption	0.009*** (0.003)	-0.0001 (0.004)
Online news consumption	0.006 (0.004)	-0.002 (0.005)
Democrat	-0.052*** (0.014)	-0.129*** (0.016)
Lean Democrat	-0.032* (0.017)	-0.082*** (0.019)
Lean Republican	0.067*** (0.016)	0.112*** (0.018)
Republican	0.036** (0.016)	0.135*** (0.018)
Facebook use X Facebook sample	0.012* (0.007)	0.002 (0.008)
Twitter use X Facebook sample	-0.001 (0.003)	0.0001 (0.004)
Internet use X Facebook sample	-0.017* (0.009)	0.007 (0.011)
Age X Facebook sample	0.0003 (0.001)	-0.001 (0.001)
Education X Facebook sample	-0.009 (0.008)	-0.006 (0.009)
Offline news X Facebook sample	-0.008* (0.004)	-0.001 (0.005)
Online news X Facebook sample	-0.0005 (0.005)	-0.010* (0.006)
Democrat X Facebook sample	0.032* (0.018)	0.033 (0.021)
Lean Democrat X Facebook sample	-0.0002 (0.024)	-0.007 (0.027)
Lean Republican X Facebook sample	-0.055** (0.025)	-0.009 (0.029)
Republican X Facebook sample	-0.015 (0.023)	0.030 (0.026)
Constant	0.308*** (0.056)	0.450*** (0.064)
Observations	26	4,145
R <sup>2</sup>	0.037	0.149
Adjusted R <sup>2</sup>	0.032	0.145

Note:

\*p<0.1; \*\*p<0.05; \*\*\*p<0.01