

# “Social Media and Political Persuasion: A Field Experiment”

## Pre-Registration

Ro’ee Levy\*

Feb 4, 2018

Note: This is an updated version of the pre-analysis plan, submitted after the baseline survey was conducted and before conducting the follow-up survey. Only Part 4, which discusses the analysis of the follow-up survey, and the appendix have been updated.

The main objective of the study is to understand the effects of news consumed through social media on political opinions and political behavior. The effects will be measured using a field experiment manipulating news content in subjects’ social media feeds. The following document details the pre-analysis plan for the experiment. The plan was submitted before the baseline survey took place and a revised version will be published between the baseline and follow-up surveys.

The study aims to answer two main research questions:

- **Does social media news consumption affect political opinions and behavior?**

According to a 2014 Pew Research Center survey “Among Millennials, Facebook is far and away the most common source for news about government and politics” (Pew, 2014). The consumption patterns of young Americans indicate a general change in news consumption habits as the share of Facebook users getting news through the site increased from 47% in 2013 to 67% in 2017 (Gottfried and Shearer, 2017). Furthermore, news consumed on social media may be different from traditional news sources and there is growing concern that the filter bubbles and echo chambers characterizing social media are polarizing the electorate.<sup>1</sup> I will measure whether and when people are persuaded by news they consume online and whether increased exposure to partisan news which matches the ideology of consumers, affects polarization.

- **How do consumers establish political opinions and what behavioral biases affect the learning process?**

There have been various experiments attempting to identify how people form beliefs based on information they receive, however the majority of these experiment take place in a lab.<sup>2</sup> While a lab provides a clean setting to investigate mechanisms, it also suffers from known limitations: usually short term effects are studied, the setting is often not natural, the subjects may be a captive audience and Hawthorne effects may bias the results. I propose exploiting social media’s infrastructure to naturally and gradually distribute news to subjects to understand how they process information to form political opinions.

---

\*Yale University. This project is supported by the Institute for Social and Policy Studies Field Experiment Initiative, by the Yale Economics Department and the Yale Program in Applied Economics and Policy.

<sup>1</sup>For example, in President Obama’s farewell address he mentioned the “splintering of our media into a channel for every taste” as a threat to democracy.

<sup>2</sup>Examples include: Arceneaux et al. (2012); Levendusky (2013); Guess and Coppock (2016); Nyhan and Reifler (2010).

# 1 Background

The study will compare theories explaining how individuals learn when consuming news from an outlet with an ideological slant. In each model consumers have limited information and a prior  $\theta^{C_0} \sim N(\theta^C, \frac{1}{h^C})$  on the state of the world, e.g. how is crime affected by immigration, how big of a threat is climate change, what is the quality of a candidate running for office. Media outlets receive a signal  $\theta^F = \theta + \varepsilon^F$  on the true state of the world,  $\theta$ , with the error term distributed normally around zero  $\varepsilon^F \sim N(0, \frac{1}{h^F})$ . Outlet  $j$  reports the state of the world with an ideological slant:  $r_j^F = \underbrace{\theta^F}_{\text{signal}} + \underbrace{s_j^F}_{\text{slant}}$ . Due to the slant, a conservative outlet will have a more positive report on a conservative candidate than a liberal outlet. The ideological slant of outlets can be explained by owner incentives (Anderson and McLaren, 2012), or by an attempt to maximize market share (Gentzkow and Shapiro, 2006).

I will compare three demand-side models that explain how consumers interpret the report,  $r^C(r^F)$ , and formulate a posterior.

- **Rational and informed consumers** - If consumers are rational and have information on the slant of the outlet, they will adjust any signal they receive from a media outlet according to the slant of the media outlet:  $r^C(r^F) = r^F - s_j^F = \theta^F$ . The consumer's posterior is the weighted average of her signal and the adjusted report:  $N(\frac{h^C\theta^C + h^F\theta^F}{h^C + h^F}, \frac{1}{h^C + h^F})$ . According to this model, the slant of an outlet has no effect since the consumer takes the slant into account and only learns from the signal the outlet actually observed. Therefore, it does not matter if households are exposed to more conservative or liberal outlets (assuming both outlets observe the same signal and that the only difference between the outlets is their slant).
- **Credulous Bayesians** (Glaeser and Sunstein, 2009) - Consumers may not be fully informed of the slant of the media outlets or fail to account for the slant. More generally, consumers may fail to account for correlated signals and repetition of information (DeMarzo et al., 2003). One simple way to model this behavioral failure is to assume consumers suffer from persuasion bias (DellaVigna and Kaplan, 2007).<sup>3</sup> In this model consumers individuals only partially adjust for the outlet's slant and infer that the signal received by the outlet is  $r^C(r^F) = r^F - (1 - \lambda)s_j^F$  where  $\lambda \in [0, 1]$  represents the persuasion bias. When  $\lambda = 0$ , consumers are fully aware of the outlet's slant and act according to the rational informed consumers model, and when  $\lambda = 1$ , they are completely unaware of the slant and treat each report of an outlet as being sampled from a distribution around the true state of the world. This model predicts that outlets with different slants lead to diverging opinions. For example, Fox News readers may become more conservative than *New York Times* readers.
- **Confirmation Bias** - A third option is that consumers suffer from confirmation bias and they uncritically accept news from outlets with a slant similar to their ideology and challenge news from opposing outlets (Taber and Lodge, 2006). One explanation for this behavior is that consumers' utility depends not only on the accuracy of the report they receive but also on reaching their desired conclusion or a position that strengthens their current opinion (Kunda, 1990; Mullainathan and Shleifer, 2005). If consumers suffer from persuasion bias, but they can partially overcome this bias with some cognitive cost,<sup>4</sup> they will trade-off the effort in overcoming the persuasion bias with the cost of changing their opinion.<sup>5</sup>

<sup>3</sup>Alternatively this type of behavior could be explained by lack of information or limited attention (Eyster and Rabin, 2010).

<sup>4</sup>The cognitive cost can be explained if disbelieving is a System 2 operation while believing is a system 1 operation (Gilbert, 1991).

<sup>5</sup>Confirmation bias can also be explained by the fact that people more easily and naturally recall memories of evidence supporting a current hypothesis, instead of evidence refuting it so perhaps even the costs of overcoming the persuasion bias are higher when consuming news from cross-cutting media outlets.

Consequently, consumers are more likely to be persuaded by signals that are closer to their current opinion. One way to model this behavior is:  $r^C(r^F) = \begin{cases} r^F & |r^F - \theta^C| \leq \tau \\ \theta^C & |r^F - \theta^C| > \tau \end{cases}$  In this interpretation, consumers ignore news if the distance between their opinion and the signal they receive is above a certain threshold,  $\tau$  which depends on the costs and benefits from changing one's opinion substantially, and naively believe news when the signal is close to their current opinion.

## 2 Design Overview

Subjects will be recruited to the study using Facebook ads. Individuals who decide to begin the survey will login using their Facebook account and complete a baseline survey. In order to randomize news exposure, towards the end of the survey subjects will be given an option to “like” four Facebook news pages. “Liking” a page on Facebook is similar to subscribing to a (free) newspaper, and once a page is “liked” its content may start appearing in the subjects’ feed (the words “like” and “subscribe” will be used interchangeably throughout the rest of the document). Subjects will be assigned to a conservative, liberal or control treatment. In the conservative treatment four pages of conservative-leaning media outlets will be offered (e.g. Fox News); and in the liberal treatment the four pages offered will be of liberal-leaning media outlets (e.g. the *New York Times*). In the control treatment, subjects will not be offered any pages. This is an encouragement design and in each treatment, subjects will choose whether to “like” each of the pages.

Approximately 6-8 weeks after the baseline survey, subjects will be invited to a follow-up survey to measure any self-reported changes in their opinions and political knowledge.

### 2.1 Sample

Facebook users living in the US who are over 18 years old will be recruited to the study using Facebook ads. Two main ad sets will be used: one set of ads will emphasize the survey and gift card lottery, while the second set of ads will also mention that the survey may be of interest to people who follow politics. Several ads will be used to ensure that there is a large sample of subjects completing the survey on desktop and mobile. After clicking an ad, potential subjects will be referred to the online survey. Only individuals who login to Facebook, provide permission to access the list of pages they “like” and reach the final section of the survey, will be included in the sample.

### 2.2 Randomization

Each participant will be assigned to a control, conservative or liberal treatment, based on a randomized block design by participants’ baseline ideology. Since the sample frame arrives over time, the entire sample will not be stratified in advance, but will be split gradually into groups, according to subjects’ ideology.

At the beginning of the survey, respondents will be asked where they position themselves ideologically on a 7-point ideological scale from very liberal to very conservative, with an additional option of “I haven’t thought about it much.” Subjects will be assigned to a treatment based on where they position themselves on the scale and when they answered the question. Each block (stratum) will be composed of three sequential subjects who chose the exact same answer among the eight options in the ideology scale survey question. The first three subjects who gave the same answer in the ideology question will be randomly split to the conservative, liberal and control

groups. The next three individuals with the same ideology will be also be split into the three groups, and the process will continue with each additional respondent. As a result of this process, there will be approximately the same number of liberal, conservatives and moderates in each treatment group. In practice, since subjects arrive gradually, the first subject in each stratum will be randomly assigned to one of the three treatments, the second subject will be randomly assigned to one of the two remaining treatments and the third subject will be assigned to the final remaining treatment.<sup>6</sup> If there are any technical issues in determining the strata for a particular subject, the subject will be randomly assigned to one of the three treatments.

Responses from the baseline survey will be used to test for balance across treatment assignments. Subjects will be compared across treatments according to their sources of news consumption, number of Facebook “likes”, answers to several questions on ideology (approval of President Trump, party identification, ideology) and standard demographic variables (age, education, gender, income and race). The questions on income, race and education will be asked after the treatment and could theoretically be influenced by the treatment.<sup>7</sup> Therefore, as a robustness check, balanced will also be tested by imputed income and education, based on the subject’s location, which is not affected by the treatment.

## 2.3 Setting

The experiment will take place in the United States in February, 2018. The following pages will be offered:

- Conservative page: Fox News, Wall Street Journal, Washington Times, The National Review
- Liberal pages: The New York Times, Slate.com, Huffington Post, MSNBC

If a subject is already subscribed to one of the liberal or conservative pages being offered in the survey, an alternative page will be offered instead. The following backup pages will be offered in case a subject already likes one of the pages:

- Conservative backup pages: The Daily Caller, The Western Journal, Washington Examiner, Townhall, The Conservative Tribune, The Blaze, Newsmax, Breitbart
- Liberal backup pages: Washington Post, Salon, Daily Kos, NPR, Mother Jones, The Atlantic, The New Yorker, PBS

The news outlets were chosen according to several criteria. First of all, they have a relatively clear ideological slant. Secondly, preference was given to popular outlets (Fox News and the *New York Times* are the second and third most popular news pages on Facebook). Finally, I attempted to include outlets of varying quality and extremity of opinions in order to allow subjects more variety when choosing which page to “like” and in order to increase the likelihood that subjects engage with at least one of the pages offered.

---

<sup>6</sup>The main risk with this randomization design is that the treatment for some subjects is per-determined and this may create selection bias. However, I believe this potential issue is not a cause for concern since subjects do not know the randomization method, even if they knew the method, they would not know what is the expected treatment they will receive at a given time (since participants are being recruited constantly throughout the US), and they do not have an incentive to receive a specific treatment.

<sup>7</sup>While it is possible to include all the questions before the experiment, I believe it makes more practical sense to end the survey with the demographic questions, since some subject may stop completing the survey when presented with these questions. In addition, I prefer that the intervention (being offered to like a Facebook page) will not be the very last thing subjects are asked to do in the survey, in order to decrease the association of the survey with this particular question.

### 3 Data

The following data sets will be used:

- Baseline survey - The survey will include questions on news consumption habits, political opinions and demographics. Subjects will also be asked what they thought is the ideological slant of various outlets.
- Follow-up survey - Subjects will be asked about their political opinions on issues and figures which were recently discussed in the news (e.g. tax reform), in addition to questions related to partisan hostility and various news outlets.
- Facebook posts and likes - Subjects will be asked to provide permissions to access Facebook data on the pages they “liked” and their posts. The pages subjects’ “like” on Facebook will be used to measure which pages each subject subscribed to in the experiment and whether and when subjects unsubscribed from pages. Facebook posts will be used to analyze the effect of the treatments on political behavior.
- Browsing and Facebook feed data - Subjects who completed the baseline survey using Google Chrome on a desktop computer will be offered a small reward in exchange for installing a Chrome extension which will provide data on their news-related browsing behavior and Facebook feed. Any estimates using this data will also be analyzed while excluding data from the first couple of days after the extension was installed, since Hawthorne effects are less likely to affect browsing behavior as subjects become less conscious of the extension.

### 4 Empirical Analysis

This project aims to understand the chain of events which may shape news consumers’ opinions. Therefore, I will focus on the following outcome domains:

#### First Stage Outcomes

- Selective Exposure - Did participants comply with the treatment? Were participants more likely to subscribe to pages that match their opinion?
- News Exposure - Were participants exposed to more news as a result of the treatment?
- News Consumption - Did participants consume more news from the pages offered? Which pages were they more likely to consume?

#### Media Effect Outcomes

- Persuasion - Were participants’ opinions and behavior affected as a result of the treatment? This participant will mostly compare subjects assigned to the conservative treatment with subjects assigned to the liberal treatment.
- Polarization - Did participants’ become more polarized as a result of the treatment? This analysis will mostly compare subjects who were assigned news matching their baseline ideology with subjects who were assigned news from opposing sources
- Knowledge - Did the treatment increase participants’ political knowledge?

## 4.1 Pre-analysis strategy

For each domain I specify below the main specification and the primary outcomes. I will create one index out of all the primary outcomes in each of the media effect domain. This procedure has the advantage of reducing the number of hypothesis tests within a domain into one test, and therefore there is no need for multiple hypothesis testing adjustments. The index will be created by first converting the outcome variables such that a higher value always has the interpretation (more knowledge / more conservative / more polarized). Each outcome will then be normalized by demeaning the outcome and dividing it by its control group standard deviation. The index will be created by assigning weights to maximize the information provided by the index, according to the method suggested by Anderson (2008). In addition to testing the effect on the standardized primary outcomes, I will present the estimates for each of the individual outcomes within the domain. In addition to raw p-values, I will report p-values that are corrected by the Benjamini-Hochberg procedure, within domains of outcome variables.

Below and in the appendix, secondary and first-stage outcomes are specified as well. These outcomes will be used to better understand how individuals learn from news they are exposed to. I do not plan to adjust the p-values for secondary and first stage outcomes. These outcomes are more exploratory and I do not have as strong priors about some of them. They will mostly be used to better understand the mechanisms, to explore other possible effects and to rescale any effects found on the main outcomes. It is likely that additional exploratory analysis will be conducted based on the final data set, in order to better explore possible mechanisms and further test theories.<sup>8</sup>

## 4.2 Selective Exposure

The first set of results will analyze which subjects complied with the treatment and “liked” at least one of the pages offered. A model such as the following will be used:

$$Y_i = \beta_1 T_i^L I_i^L + \beta_2 T_i^L I_i^C + \beta_3 T_i^C I_i^L + \beta_4 T_i^C I_i^C + \alpha X_i + \varepsilon_i, \text{ where:}$$

- $T_i^L \in \{0,1\}$  is whether subject  $i$  was assigned to a liberal treatment, and  $T_i^C \in \{0,1\}$  is whether subject  $i$  was assigned to a conservative treatment.
- $I_i^L \in \{0,1\}$  is whether subject  $i$  has a liberal ideology, and  $I_i^C \in \{0,1\}$  is whether subject  $i$  has a conservative ideology. The ideology of each subject will be determined according to where subjects place themselves on the ideological scale. For moderate subjects or subjects who do not know where to place themselves, the party they identify with will be used to determine their ideology, and if they do not lean towards either party, their ideology will be determined according to the candidate they supported in the 2016 elections. I predict that by using this method more than 90% of individuals will be assigned an ideology.
- $Y_i$  is whether the subject complied with the treatment and “liked” at least one page. As a robustness test I will run an identical regression measuring compliance two months after the initial treatment. The result for this outcome could be different if subjects “unliked” some of the pages they initially subscribed to in the experiment. In the follow-up survey, participants will be offered an option to unlike the pages they originally “liked”. I will run a similar regression to measure who unlikes pages before the follow-up survey and during the follow-up survey.
- $X_i$  are standard demographics: gender, age and squared age. In addition,  $X$  will also control for the set of pages offered.

---

<sup>8</sup>Olken (2015) provides a detailed discussion on the advantages and limitations of pre-analysis plan and exploratory analysis in economics.

The main coefficient of interest is the difference between the effect of being offered a page that matches your ideology and being offered a page with an opposing ideology. I will test if  $\beta_1 = \beta_3$  and  $\beta_2 = \beta_4$

To better understand the choice for each specific page, the following model can be used:

$Y_{ij} = \beta_1 Know_{ij} + \beta_2 Know_{ij} * (Ideology_i - Slant_{ij}) + \alpha X_i + \varepsilon_{ij}$  where:

- $Y_{ij}$  is whether individual  $i$  “liked” page  $j$
- $Know_{ij} = \begin{cases} 0 & \text{Subject } i \text{ knows the slant of page } j \\ 1 & \text{Otherwise} \end{cases}$
- $Slant_{ij}$  is the perceived ideology of outlet  $j$  according to subject  $i$  on a  $[-3,3]$  scale, where  $-3$  is most liberal and  $+3$  is most conservative
- $Ideology_i$  is the self-reported ideology of the individual on the same scale.
- $\beta_2$  measures how the ideological distance from a page affects the willingness to subscribe to it (“selective exposure”).

Note that this regression is only meant to provide a simple model for how people choose outlets, but the relationship is not necessarily causal, since it does not rely on the exogenous variation generated by the intervention. This regression can be analyzed with and without page fixed effects.

A separate measure of selective exposure is whether participants choose to unsubscribe from pages after they have been exposed to them. I will measure who unliked the offered pages during the study period, and in the follow-up survey I will offer participants an opportunity to unlike the page they original subscribed to.

### 4.3 News exposure

Two main specifications will be used to measure whether the experiment had an effect on participants’ Facebook feeds. In the first specification I will compare the treatment groups with the control group.

$Y_i = \beta_0^L T_i + \alpha X_i + \varepsilon_i$  where:

- $T_i = \begin{cases} 1 & \text{Participant } i \text{ is assigned to the liberal or conservative treatment} \\ 0 & \text{Otherwise} \end{cases}$
- $Y_i$  is the average number of pages participant  $i$  has seen in her news feed. The value will be calculated using a survey question asking participants how many times they saw different Facebook pages in their news feed in the past week. For the treatment groups  $Y_i$  will be calculated only according to the appearance of the pages the participants were randomly offered and will be divided by the number of pages offered (usually four). By definition, the control group was not offered any pages, so  $Y_i$  will be calculated according to the appearance of all eight pages that could have been offered if the participant would have been assigned to the conservative or liberal group.
  - For the subset of participants who installed the chrome extension this outcome will also be calculated based on the actual number of pages that appeared in their Facebook feed.
- $X_i$  control for the number of pages the participant was exposed to in the past, based on the baseline survey.

Similarly to the selective exposure section, more detailed regressions can be analyzed to measure the effect of liberal and conservative pages or matching and opposing pages separately.

Since for each participant four pages are randomly offered and four pages are randomly not offered, an alternative specification is to estimate this effect only in the treatment groups using individual fixed effects. In this specification, instead of exploiting variation between individuals, the variation exploited is between pages, within an individual:

$$Y_{ij} = \beta_0 O_{ij} + \gamma_i + \varepsilon_{ij} \text{ where:}$$

- $O_{ij}$  is whether page  $j$  was offered to individual  $i$ .

Since in practice the pages are split into two bundles,  $j$  can denote a group of outlets, such that:

$$j = \begin{cases} C & \text{Conservative Pages} \\ L & \text{Liberal Pages} \end{cases} \text{ and } O_{ij} = \begin{cases} 1 & (j = C \text{ and } T_i^C = 1) \text{ OR } (j = L \text{ and } T_i^L = 1) \\ 0 & \text{Otherwise} \end{cases}$$

- $Y_{ij}$  is the average outcome at the individual\*page level.
- $\gamma_i$  controls for individual fixed effects.

Finally, exposure will also be measured using a question in the follow-up survey asking people how diverse their news feed is.

#### 4.4 News consumption

News consumption data will only be observed for individuals who installed the chrome extension. Three main outcomes can be used to measure consumption:

- Direct effect on consumption - Browsing data will be used to test whether individuals were more likely to consume news through Facebook, from pages they were randomly offered, compared to the control group and to pages not offered. Past consumption data may be used to control for baseline news consumption. Similarly, the effect of exposure on consumption can be estimated, using pages offered as the instrument. One challenge with using this measure is identifying when consumers accessed a news article through Facebook.
- Effect on consumption through other means - If individuals start consuming news from the Facebook pages which were offered to them, they may start consuming articles from the pages' corresponding websites regardless of whether the article appears in their Facebook feed. Alternatively, it is possible that the news consumed through Facebook will replace the consumption of the same content through other means. Therefore, the same analysis described in the previous step will be conducted on news which was not accessed through Facebook, and on all news consumed.
- Indirect effects on other news - If indeed the exposure and consumption of the assigned news content increase, I will explore how exposure and consumption of other content are affected. It is possible that the exposure to other news content will not change (e.g. news offered in the experiment crowds out exposure of entertainment sources instead of other news), it may increase (if news articles are complements) or decrease (if news articles are substitutes). Furthermore, it is important to understand whether the slant of other news consumed was affected by the experiment. For example, if an individual was randomly offered to like Fox News in the treatment and started consuming content from that outlet instead of other conservative outlets

(such as the *Wall Street Journal*) we would expect a different effect than in a case where Fox News crowded-out liberal content (such as the *New York Times*). I will analyze if the total number of conservative and liberal outlets consumers were exposed to and consumed changed as a result of the treatment. Exposure will be analyzed using posts appearing in the subjects' feeds and consumption will be based on news consumed through Facebook.

## 4.5 Persuasion

The following intention-to-treat regressions will be used to measure the average effect of the intervention.

$Y_i = \beta_0^L T_i^L + \beta_0^C T_i^C + \alpha X_i + \varepsilon_i$  where:

- $Y_i$  is the outcome, where a higher  $Y_i$  indicates a more conservative position. The following outcomes will be used:
  - Opinions - Opinion on topics and figures currently discussed in the news are the primary outcome which will be studied. The final list of outcomes is detailed in the appendix. I expect the effect to be larger for outcomes where participants did not have strong priors and which were discussed in the news often during the study period (such as the favorability ratings of individuals that happened to be in the news). The effect on opinions regarding concrete policies will probably be weaker, especially if participants already had strong opinions on the policies, and I expect any effect on party identification or ideology to be very weak and probably undetectable.
  - Behavior - The slant of each article or post shared by an individual will be determined according to a database based on articles shared on Facebook (Bakshy et al., 2015).<sup>9</sup> I will measure whether the treatment had an effect on the difference between the number of conservative and liberal posts. I will also test whether individuals were more likely to share posts specifically from the pages that they were randomly offered. This outcome will measure whether the treatments had any effect on actual political behavior (sharing partisan articles with friends), in addition to the effect on self-reported political opinions.  
Behavior will also be estimated using two survey measures. In the followup survey, participants will be asked to view three Facebook posts and state if they would "like" or share the posts. For a subset of consumers who grant permissions to observe their Facebook posts, I will be able to observe whether the posts were actually shared. Participants will be offered a conservative, liberal and moderate post. In addition, participants will be asked to choose between two possible gift cards, a gift card for DICK's Sporting Goods and a gift card for Target. I will estimate if recent news regarding DICK's Sporting Goods new gun policy had an effect on participants' choice.
- $X_i$  is a set of covariates from the baseline survey. Since the effects of the experiment are expected to be relatively small, and since the baseline coefficients will have high predictive power, it is important to include coefficients to increase precision. The covariates included will be the strata, self-reported ideology, party affiliation and the approval level of President Trump. Additional covariates will be included for outcomes where a lagged variable exists (for example, in regressions where the outcome is the slant of subjects' posts, the model will also control for the baseline slants of posts). Using a nested F-Test I will check if a model which includes these covariates along with age, age squared and gender, is significantly different than a model without the demographic covariates when predicting the outcomes of policy questions in the baseline

---

<sup>9</sup>A database created by BuzzFeed may be used to identify the slant of additional outlets: <https://github.com/BuzzFeedNews/2017-08-partisan-sites-and-facebook-pages>

survey. If demographics add valuable information they will be included in the main analysis as well. After conducting the baseline survey, it became apparent that some participants did not complete the survey due to a technical issue. The issue occurred only among participants completing the survey with iPhone and iPads. Therefore, I will also control for whether the participant took the survey using an iOS operating system.

- $\beta_0^C, \beta_0^L$  are the coefficient of interests, which measure the effect of being offered conservative vs. liberal news pages. For example,  $\beta_0^C > 0$ , will imply that a conservative treatment made subjects more conservative.

All the regressions will use robust standard errors. The main estimations will not be weighted.

The primary test is whether  $\beta_0^L < \beta_0^C$ , that is whether participants randomly offered liberal news become more liberal than participants randomly offered conservative news. This will determine if consumers suffer from persuasion bias and tend to change their opinions according to an outlet they are randomly exposed to, or whether they are rational and adjust for the slant of the outlet. Testing whether  $\beta_0^L < 0$  and  $0 < \beta_0^C$  may provide further evidence on which subjects were persuaded by the slant of the outlet, by comparing each treatment arm to the control group. However, these tests require more power and therefore the probability that an effect will be detected is smaller. In addition, these tests are problematic in cases where subjects hold extreme opinions and as a result exposure to news that matches their opinion makes them more moderate compared to the control group.

If subjects are indeed persuaded by outlets, the model will also be estimated only for subjects who “liked” pages for which they were able to estimate the slant of the outlet. When subjects are exposed to news from outlets they are not familiar with, they may be affected by the slant of the outlet since they are not aware of the slant. However, if subjects suffer from persuasion bias also when they are familiar with the outlet, the results will suggest that persuasion bias cannot be explained only as a consequence of missing information.

#### 4.5.1 IV Regression

Any result from the intention-to-treat estimate is relevant for a policy offering people exposure to new opinions, but will not estimate the effect of actually subscribing to news outlets. Therefore, I will use the following IV regressions based on two separate measures of compliance:

$$Y_i = \beta_0^L C_i^L + \beta_0^C C_i^C + \alpha X_i + \varepsilon_i \text{ with } T_i^L, T_i^C \text{ as instruments and}$$

$$Y_i = \beta_0^L E_i^L + \beta_0^C E_i^C + \alpha X_i + \varepsilon_i \text{ with } T_i^L, T_i^C \text{ as instruments. Where:}$$

- $C_i^L, C_i^C$  is whether subject  $i$  initially complied with the treatment and “liked” liberal or conservative pages.
- $E_i^L, E_i^C$  is whether subject  $i$  was exposed to the liberal/conservative pages as a result of the treatment. Exposure to news pages will be defined based on self-reported answers in the follow-up survey. For the subset of users who install the chrome extension, exposure can be defined according to actual news feed data. This measure is important, since subjects may “like” a page but not see its content in their feed, either because they “unliked” it, Facebook’s algorithm determines that the content offered by the page would not interest the subjects or the subjects do not engage with the page and its content stops appearing as a result. It is unlikely that a page will have an effect on opinions if it does not appear in the subject’s feed.
- $\beta_0^L, \beta_0^C$  are the LATE estimators, the average effect of conservative/liberal pages among individuals who subscribe to these pages when they are offered, or among individuals who are exposed to articles from those pages as a result of the pages being offered.

These regressions will estimate local average treatment effects, the average effects for compliers. Therefore, it is important to understand who are the compliers. Since persuasion in the experiment occurs over a roughly two-month period, it is possible that some subjects would have complied with the treatment without being treated (“always-takers”). For example, in the weeks following the experiment subjects could have “liked” one of the pages they were offered in the experiment even if the page would not have been offered. Thus, it is necessary to disentangle the characteristics of compliers from always-takers. The treatment groups will be used to identify the proportion of never-takers, and data from the control group will be used identify the share of always-takers in order to formally identify the characteristics of compliers (Abadie, 2003).

#### 4.5.2 Predicted compliance

One challenge with an intention to treat design is that the average effect is made much weaker due to participants who do not comply with the treatment. This substantially increases the power required to detect an effect. In this experiment, many covariates are collected before the randomization occurs. I will attempt to use this data (especially pages individuals previously liked) to predict who will comply with the experiment. For each participant a probability of compliance for the conservative and liberal treatments can be calculated. Using these estimates, I will re-estimate the ITT effect for individuals who were predicted to have a high probability of complying with the treatment.

#### 4.5.3 Motivated Reasoning and Confirmation Bias

One of the goals of the study is to determine how respondents are affected by counter-attitudinal news sources: are they persuaded, do they ignore content from opposing outlets or are their original positions reinforced (a “boomerang effect”). Therefore, the most important heterogeneous analysis will be based on the following regression model:

$$Y_i = \beta_1 T_i^L I_i^L + \beta_2 T_i^L I_i^C + \beta_3 T_i^C I_i^L + \beta_4 T_i^C I_i^C + I_i + \alpha X_i + \varepsilon_i$$

Only subjects whose ideology is identified ( $I^L = 1$  or  $I^C = 1$ ) will be included in this regression.

- The primary test for motivated reasoning is whether  $\beta_2 = \beta_3 = 0$ . This will test if individuals are affected by opposing news, i.e. are liberals persuaded by the conservative treatment and conservatives persuaded by the liberal treatment. In addition, I will test if  $\beta_1 = \beta_2$  and  $\beta_3 = \beta_4$  to determine if conservatives and liberals were affected similarly by the same treatment.<sup>10</sup>

#### 4.5.4 Additional heterogeneity analysis

Additional heterogeneous effects will be tested, such as the following:

- Subjects’ openness - The purpose of this analysis is to test whether people who are open to new political opinions are more likely to be persuaded. Openness will be defined according to two questions taken from the Ten Item Personality Inventory (Gosling et al., 2003). I may also be able to predict openness based on pages the subjects subscribed to before the experiment (Kosinski et al., 2013). Similarly, I will test if people who have changed their minds in the past according to something that appeared in their feed and who are less certain in their opinions are more affected by the treatment.

<sup>10</sup>If there is only a small difference between the effect of matching news and non-matching news, measuring the interaction effects may require more power and the estimates will be imprecise. One possibility to increase the precision of these estimates is to assume that the effect of conservative and liberal treatments is similar (in opposite direction) and set  $T_i^L = -T_i^C$

- Echo chamber - I will test for heterogeneous effects according to whether subjects are more likely to be exposed to similar opinions in their news feed (whether their news feed resembles an echo chamber). This effect will be measured according to a question in the baseline survey and based on extension data on actual posts appearing in the subjects' news feed.
- Absolute opinion - Using self-reported ideology, I will test whether ideological partisans are less likely to be persuaded than moderates.<sup>11</sup>
- Sophistication - Two questions in the baseline survey will be based on the cognitive reflection test (CRT) which attempts to distinguish between two types of cognitive processes: intuitive, quick processes, compared to slower and reflective processes (Shane, 2005). We would expect less sophisticated subjects to suffer from persuasion bias, since they may think less critically on the outlet's slant. On the contrary, more sophisticated subjects are likely to fit the rational consumer model, or even act according to motivated reasoning theory, since they can more easily challenge news they do not agree with (Pennycook and Rand, 2017).

All the analysis of heterogeneous effect will control for the pages the subjects were offered in order to prevent an omitted variable bias. Such a bias is possible since individuals who were offered backup pages have different characteristics than individuals who were offered the primary pages, and they were also exposed to different news as a result of the treatment.

## 4.6 Polarization

Models like the following will be used to measure the effect of variation in news exposure on polarization:

$$Y_i = \beta_4^P P_i + \beta_4^A A_i + \alpha X_i + \varepsilon_i$$

- $P_i$  measures whether the participant was randomly offered pro-attitudinal outlets:

$$P_i = \begin{cases} 1 & (T_i^L = I_i^L = 1) \text{ OR } (T_i^C = I_i^C = 1) \\ 0 & \textit{otherwise} \end{cases}$$

- $A_i$  measures whether the participant was randomly offered counter-attitudinal (anti-attitudinal) outlets:

$$A_i = \begin{cases} 1 & (T_i^L = I_i^C = 1) \text{ OR } (T_i^C = I_i^L = 1) \\ 0 & \textit{otherwise} \end{cases}$$

This model is based on the same randomization used in the previous model, but the focus is not on whether the treatment is conservative or liberal but on whether it matched or opposed the subject's baseline opinion.

- $Y_i$  is the outcome, where a higher  $Y$  indicates greater polarization. Three main outcomes will be used in this model (the detailed list of outcomes appears in the appendix):
  - Absolute opinions (attitudinal-polarization): absolute opinions will be measured according to a similar set of variables measuring the persuasion effect, but the outcomes will be determined by the strength of opinions instead of their position on a left-right scale. For example, I will test whether the *absolute* slant of posts increased when subjects were assigned to receive matching news pages, compared to opposing news pages. In other words, this variable will measure whether people become more extreme when being exposed to information matching their ideology.

<sup>11</sup>Since a very small share of conservatives initially completed the survey, an additional set of ads was used to target only conservatives and moderates. It is possible that individuals who were recruited through these ads were more ideological, but also were different in other dimensions. Therefore, as a robustness test, I will control for ad fixed effects when testing for heterogeneity by absolute opinions.

- Partisan hostility (affective-polarization): partisan hostility will be measured according to the difference between the feelings of subjects towards the party they identify with and their feelings towards the opposing party, based on a feeling thermometer question, and according to how subjects state they would feel if their son/daughter married an individual identifying with the opposing party. Affective polarization will also be measured based on political empathy, which will be defined according to an index of two survey questions (Reit et al., 2017).
  - Perceived polarization: Participants will be asked how liberal and conservative a typical Democrat/Republican is, the difference will be defined as perceived polarization.
  - Behavior: Participants will be asked whether they would “like” and share a moderate Facebook post stating “In seeking truth, you have to get both sides of a story”. For participants who provide permissions to access their Facebook posts, I will measure whether the post was actually shared.
- Coefficients of interest
    - $\beta_4^P, \beta_4^A$  measure whether exposure to opposing news mitigates polarization as compared to matching news. For example,  $\beta_4^P > 0$  implies that exposure to more pro-attitudinal news increases polarization. The most important test is whether  $\beta_4^A < \beta_4^P$ : do subjects become more polarized when exposed to matching news compared to opposing news. As usual, it is possible to compare the effect of matching and opposing news to the control group by testing whether  $\beta_4^A < 0, 0 < \beta_4^P$ .

My prior for most of these outcomes is that exposure to opposing news will make people less polarized. However, a backlash effect is also possible and it is important to investigate if such an effect occurs. Therefore, all tests will be two-sided.

## 4.7 Knowledge and Engagement

While this study focuses on persuasion and polarization, I will also take advantage of the experiment to test whether exposure to random news increases political knowledge and engagement, compared to the control group. More interestingly, it is possible that exposure to opposing news will provide subjects with greater knowledge on specific issues that are mostly debated by one side of the aisle. Political knowledge will be measured with the following outcomes:

- Self reported familiarity with individuals and events - Participants will be asked if they have a favorable opinion on several individuals recently discussed in the news, and will have an option to state that they have never heard of the individual. I will test whether the treatments affected name recognition. Similarly, I will ask individuals how much they have heard about several recent events. While this question could be used to measure news consumption, it also provides an indirect measure of knowledge since it is safe to assume that individuals are uninformed of events they have not heard of.
- Accurate perceptions - Participants will be asked how many people they think have a favorable opinion on President Trump, and will be asked whether four news stories are correct or not.
- Engagement - Participants will be asked how often they follow the news and whether they care about the 2018 elections.

All estimates in this section will control for how often people follow the news in the baseline survey.

## 5 Implications

New technology generally, and social media specifically, have enabled individuals to consume more news from media outlets that match their ideology. This study can show whether exposure to outlets with clear ideological slants polarizes readers, and which policies may decrease polarization. If individuals act as Credulous Bayesians and are persuaded by the outlets they read ( $\beta_0^L < \beta_0^C$ ), then social media, which is associated with increased segregation in news consumption, is expected to increase polarization. However, perhaps polarization could be mitigated by designing mechanisms that expose individuals to more cross-cutting news (if  $\beta_4^A < 0$ ). Such policies have already been suggested (Sunstein, 2017), but they should first be tested empirically. This experiment provides a nudge that encourages people to receive random news and thus tests how serendipity encounters affect consumers.

# Appendix - Followup Survey Outcomes of Interest

Category	Measure	Purpose of measure	Prior	Notes
<b>First Stage / Mechanisms</b>				
News exposure	How many times have you seen each of the following news sources in your Facebook feed over the past week?	First Stage	Match > Opposing > Control	Control for baseline measure. This effect can be measured in two ways: With individual fixed effects by comparing pages offered with pages not offered, and by comparing the treatment and control groups.
	Thinking about the opinions you see about government and politics on Facebook, how often are they in line with your own views?	Manipulation Check / Mechanisms	Match/Control > Opposing	Control for baseline measure
Media perception	In general, how much trust and confidence do you have in the mass media when it comes to reporting the news fully, accurately, and fairly?	Mechanisms	Match > Opposing	
	Where would you place each of the following news outlets on the following scale from very liberal to very conservative?	Manipulation Check / Mechanisms	Participants in the treatment groups may form an opinion that is more accurate on outlets they were randomly exposed to. Participants who were offered opposing pages may believe that pages they previously liked are more biased.	Control for baseline measure. This effect can be measured in two ways: With individual fixed effects by comparing pages offered with pages not offered, and by comparing the treatment and control groups.
Behavior	If you liked news page(s) in the previous survey, would you like an option to unlike the page(s) at the end of this survey?	Selective exposure	Opposing > Match. Most people will not unlike the page	

<b>Category</b>	<b>Measure</b>	<b>Purpose of measure</b>	<b>Prior</b>	<b>Notes</b>
	Did participants who expressed interest in unliking a page, actually unlike it at the end of the survey	Selective exposure	Opposing > Match Most people will not unlike the page	The data for this outcome will be available for participants who granted permissions to access their likes data

Category	Measure	Purpose of measure	Prior	Notes
<b>Persuasion</b>				
Political Sentiment: Individuals Favorability Ratings	Donald Trump (job approval)	Primary Outcome	Conservative > Control > Liberal	Controls for baseline measure
	Andrew McCabe	Primary Outcome	Liberal > Control > Conservative	
	David Hogg	Primary Outcome	Liberal > Control > Conservative	
	Robert Mueller	Primary Outcome	Liberal > Control > Conservative	
	Stormy Daniels	Primary Outcome	Liberal > Control > Conservative	
	John Bolton	Primary Outcome	Conservative > Control > Liberal	
	Michael Cohen	Primary Outcome	Conservative / Control > Liberal	
	Scott Pruitt	Primary outcome	Conservative / Control > Liberal	
	Hillary Clinton	Primary outcome	Liberal / Control > Conservative	
	Nikki Haley	Placebo	Do not expect an effect	
Political Sentiment: Group Favorability Ratings	The NRA	Primary Outcome	Conservative > Control > Liberal	
	The March for Our Lives and National School Walkout	Primary Outcome	Liberal > Control > Conservative	
	The Government of California	Primary Outcome	Liberal > Control > Conservative	
	The FBI	Primary Outcome	Liberal > Control > Conservative	
	Illegal / Undocumented Immigrants	Primary Outcome	Liberal > Control > Conservative	Additional control: Concern about illegal immigration in baseline
Political Sentiment: Feeling Thermometer	Feeling thermometer: Donald Trump	Primary Outcome	Conservative > Control > Liberal	Control for identical question in baseline
	Feeling thermometer: Republican Party	Secondary Outcome	Conservative > Control > Liberal	Control for identical question in baseline
	Feeling thermometer: Democratic Party	Secondary Outcome	Liberal > Control > Conservative	Control for identical question in baseline
Political Sentiment	How much do you blame each of the following for mass shootings in the United States? The NRA	Secondary outcome	Liberal > Control > Conservative	

Category	Measure	Purpose of measure	Prior	Notes
	How much do you blame each of the following for mass shootings in the United States? Law enforcement	Secondary outcome	Conservative > Control > Liberal	Expect a stronger effect for participants who took the baseline survey in Feb, 2018
	How much do you blame each of the following for mass shootings in the United States? Security in schools	Secondary outcome	Conservative > Control > Liberal	Expect a stronger effect for participants who took the baseline survey in Feb, 2018
Political Beliefs / Opinions	Do you think McCabe was fired because... (select all that apply)? Improper action / A way to damage McCabe / Revenge	Primary Outcome	Conservative (improper action) > Control > Liberal	
	Do you think the FBI investigation into Trump campaign officials' contacts with Russian government officials is...A serious attempt to find out what really happened / A politically-motivated attempt to embarrass Donald Trump / Both	Primary Outcome	Liberal (serious attempt) > Control > Conservative	Control for baseline measure on whether the investigation is fair
	Do you believe the following is true: President Trump has attempted to derail or obstruct the investigation into the Russian interference in the 2016 election	Primary Outcome	Liberal > Control > Conservative	Control for similar baseline measure
	How likely do you think it is that a trade war will develop between the United States and foreign countries in the next year?	Primary Outcome	Liberal > Control > Conservative	
	Do you think tariffs on imported products help or hurt the U.S. economy?	Secondary Outcome	Conservative (help) > Control > Liberal	
	Some people think that the way people talk needs to change with the times to be more sensitive to people from different backgrounds. Others think that this has already gone too far and many people are just too easily offended. Which is closer to your opinion?	Secondary outcome	Liberal (the way people talk needs to change) > Control > Conservative	
	In the past two months, have you modified your views about a political or social issue because of something you saw on social media?	Manipulation Check / Mechanisms	Liberal / Conservative > Control	Control for baseline measure

Category	Measure	Purpose of measure	Prior	Notes
	Overall, do you think the economy is getting better or worse?	Secondary Outcome	Conservative > Control > Liberal	
	Would you say things in the country are going in the right direction?	Secondary Outcome	Conservative > Control > Liberal	
Policy	Do you support: Banning assault-style weapons	Primary Outcome	Liberal > Control > Conservative	
	Do you support: Repealing the second amendment	Secondary outcome	Liberal > Control > Conservative	
	Do you support: Allowing teachers and school officials to carry guns on school grounds	Secondary outcome	Conservative > Control > Liberal	
Perceptions	Who do you think will win the majority of seats in the US House of Representatives?	Secondary outcome	Conservative (Republican Party) > Control > Liberal	
	In your opinion, what is the percentage of people who approve of the job Donald Trump is doing as President?	Secondary outcome	Conservative > Control > Liberal	Control for baseline measure
	Do you think the main purpose of the March for Our Lives movement is? (Passing gun control laws / Repealing the 2nd amendment / Both)	Secondary outcome	Conservative (2nd amendment) > Control > Liberal	
Political Identification	Generally speaking, do you think of yourself as a Democrat, a Republican or an Independent? If Republican/Democrat: Would you call yourself a strong Republican/Democrat or a not very strong Republican/Democrat? If Independent: Do you think of yourself as closer to the Republican or Democratic Party?	Secondary outcome	Conservative (Republican) > Control > Liberal	Control for baseline measure (I would be surprised if there is any detectable effect on party identification or ideology)
	Where would you place yourself on a scale from very liberal to very conservative?	Secondary outcome	Conservative > Control > Liberal	Control for baseline measure
Behavior	If an election for U.S. Congress were being held today, who would you vote for in the district where you live?	Secondary outcome	Conservative (Republican candidate) > Control > Liberal	Control for baseline measure
	Did you participate in a demonstration or protest in the past two months? Gun Rights	Secondary outcome	Conservative > Control > Liberal	

<b>Category</b>	<b>Measure</b>	<b>Purpose of measure</b>	<b>Prior</b>	<b>Notes</b>
	Did you participate in a demonstration or protest in the past two months? Gun Control	Secondary outcome	Liberal > Control > Conservative	
	Conservative post: Would you "like" the following post?	Secondary outcome	Conservative > Control > Liberal	
	Liberal post: Would you "like" the following post?	Secondary outcome	Liberal > Control > Conservative	
	Actually share conservative post	Secondary outcome	Conservative > Control > Liberal	
	Actually share liberal post	Secondary outcome	Liberal > Control > Conservative	
	WTP to pay for a DICK's sporting goods gift card compared to a Target gift card	Secondary outcome	Liberal > Control > Conservative	I will also test whether randomly asking participants if they had heard of DICK's Sporting Goods new gun policy affects this outcome

Category	Measure	Purpose of measure	Prior	Notes
<b>Polarization</b>				
Attitudinal-Polarization	Absolute opinions on favorability outcomes	Secondary outcome	Match (more extreme) > Control > Opposing	The favorability variables will be transformed to measure how extreme the participants' opinions are, where the highest value is for very favorable/very unfavorable opinion
Affective Polarization / Political Empathy	Difference between good ideas in the party the individual is closer to and good ideas in the other party	Primary outcome	Match > Control > Opposing	
	How upset would you feel if your son or daughter married a Republican/Democrat (asked for other party)	Primary outcome	Match (more upset) > Control > Opposing	
	Feeling thermometer: Difference between party the individual is closer to and the other party	Primary outcome	Match > Control > Opposing	
	Difference between finding it difficult to see things from other party's point of view and finding it difficult to see things from the party the individual is closer to point's of view	Primary outcome	Match > Control > Opposing	Control for measure in baseline survey
	Difference between finding it important to consider the perspective of the party the individual is closer to and finding it important to consider the perspective of the other party	Primary outcome	Match > Control > Opposing	
Behavior	Moderate post: Would you "like" the following post?	Primary outcome	Opposing > Control > Match	
	Moderate post: Would you share the following post?	Primary outcome	Opposing > Control > Match	
	Moderate post: Actually share moderate post	Primary outcome	Opposing > Control > Match	The data for this outcome will only be available for participants who granted permissions to access timeline data
Perceived Polarization	Absolute difference between the perceived ideology of Republicans and Democrats	Secondary outcome	Match > Control > Opposing	Control for baseline measure

Category	Measure	Purpose of measure	Prior	Notes
<b>Knowledge and Engagement</b>				
Self-Reported Knowledge	Heard of: Louis Farrakhan	Primary Outcome	Conservative > Control / Liberal	Control for whether people remember name of congressmen in baseline for all name recognition outcomes
	Heard of: Michael Cohen	Primary Outcome	Liberal > Control / Conservative	
	Heard of: John Bolton	Secondary Outcome	Liberal / Conservative > Control	
	Heard of: Andrew McCabe	Secondary Outcome	Liberal / Conservative > Control	
	Heard of: Nikki Haley	Placebo	Don't expect an effect	
	How much have you heard about: Stephon Clark was shot and killed by police officers in Sacramento	Primary outcome	Liberal > Control / Conservative	
	How much have you heard about: Hillary Clinton suggested that many white women voted for Trump since they took their voting cues from their husbands	Primary outcome	Conservative > Control / Liberal	
	How much have you heard about: Attorney General Jeff Sessions fired Andrew McCabe two days before he was scheduled to retire as Deputy Director of the FBI	Secondary outcome	Conservative / Liberal > Control	
	How much have you heard about: DICK's Sporting Goods raised the minimum age for purchase of guns and ended sales of all assault-style rifles in its stores	Secondary outcome	Conservative / Liberal > Control (expect a stronger effect for participants who began the baseline survey before February 29, 2018)	Randomized to half the sample
Accurate Knowledge	Do you believe the following is true: Most people will receive an income tax cut, salary increase or bonus under the new tax reform law	Primary outcome	Conservative > Control > Liberal	Statement is probably true

<b>Category</b>	<b>Measure</b>	<b>Purpose of measure</b>	<b>Prior</b>	<b>Notes</b>
	Do you believe the following is true: The Russian government tried to influence the 2016 presidential election	Primary outcome	Liberal > Control > Conservative	Statement is probably true
	Do you believe the following is true: President Trump is a criminal target of Robert Mueller's investigation	Primary outcome	Liberal > Control > Conservative	Statement is probably false
	Do you believe the following is true: The US has recently started building a new border wall at the US-Mexico border	Primary outcome	Conservative > Control > Liberal	Statement is probably false
	In your opinion, what is the percentage of people in the country who approve of the job Donald Trump is doing as President? Difference between Trump's actual favorability when the respondent answered the question (based on <a href="http://fivethirtyeight.com">http://fivethirtyeight.com</a> ) and the respondent's guess.	Secondary outcome	Match (larger mistake) > Control > Opposing	Control for baseline measure.
Engagement	How much would you say that you personally care about the way the election to the U.S. House of Representatives comes out?	Secondary outcome	Conservative / Liberal > Control	Control for previous voting behavior in midterms, based on baseline survey
	How often do you pay attention to what is going on in government and politics?	Secondary outcome	Conservative / Liberal > Control	Control for baseline measure

## References

- Abadie, A. (2003). Semiparametric Instrumental Variable Estimation of Treatment Response Models. *Journal of Econometrics* 113(2), 231–263.
- Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects. *Journal of the American Statistical Association* 103(484), 1481–1495.
- Anderson, S. P. and J. McLaren (2012). Media Mergers and Media Bias with Rational Consumers. *Journal of the European Economic Association* 10(4), 831–859.
- Arceneaux, K., M. Johnson, and C. Murphy (2012). Polarized Political Communication, Oppositional Media Hostility, and Selective Exposure. *The Journal of Politics* 74(01), 174–186.
- Bakshy, E., S. Messing, and L. A. Adamic (2015). Exposure to Ideologically Diverse News and Opinion on Facebook. *Science* 348(6239), 1130–1132.
- DellaVigna, S. and E. Kaplan (2007). The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics* 122(3), 1187–1234.
- DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003). Persuasion Bias, Social Influence, and Unidimensional Opinions. *The Quarterly Journal of Economics* 118(3), 909–968.
- Eyster, E. and M. Rabin (2010). Naïve Herding in Rich-Information Settings. *American Economic Journal: Microeconomics* 2, 221–243.
- Gentzkow, M. and J. M. Shapiro (2006). Media Bias and Reputation. *Journal of Political Economy* 114(2), 280–316.
- Gilbert, D. T. (1991). How Mental Systems Believe. *American Psychologist* 46(2), 107–119.
- Glaeser, E. L. and C. R. Sunstein (2009). Extremism and Social Learning. *Journal of Legal Analysis* 1(1), 263–324.
- Gosling, S. D., P. J. Rentfrow, and W. B. Swann (2003). A Very Brief Measure of the Big-Five Personality Domains. *Journal of Research in Personality* 37(6), 504–528.
- Gottfried, J. and E. Shearer (2017). News Use Across Social Media Platforms 2017. *Pew Research Center*.
- Guess, A. and A. Coppock (2016). The Exception, Not the Rule? The Rarely Polarizing Effect of Challenging Information. *Working Paper*.
- Kosinski, M., D. Stillwell, and T. Graepel (2013). Private Traits and Attributes are Predictable from Digital Records of Human Behavior. *Proceedings of the National Academy of Sciences* 110(15), 5802–5805.
- Kunda, Z. (1990). The Case for Motivated Reasoning. *Psychological Bulletin* 108(3), 480–498.
- Levendusky, M. S. (2013). Why do Partisan Media Polarize Viewers? *American Journal of Political Science* 57(3), 611–623.
- Mullainathan, S. and A. Shleifer (2005). The Market for News. *American Economic Review* (February), 1031–1053.
- Nyhan, B. and J. Reifler (2010). When Corrections Fail: The Persistence of Political Misperceptions. *Political Behavior*, 32(2), 303–330.
- Olken, B. A. (2015). Promises and Perils of Pre-Analysis Plans. *Journal of Economic Perspectives* 29(3), 61–80.

- Pennycook, G. and D. G. Rand (2017). Who Falls for Fake News? The Roles of Analytic Thinking Motivated Reasoning, Political Ideology, and Bullshit Receptivity. *Working Paper*.
- Pew (2014). *Political Polarization and Media Habits*. Pew Research Center.
- Reit, E., R. Willer, and J. Zaki (2017). Causes and Consequences of Political Empathy. *Work in Progress, Stanford University*.
- Shane, F. (2005). American Economic Association. *The Journal of Economic Perspectives* 19(4), 25–42.
- Sunstein, C. (2017). *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Taber, C. S. and M. Lodge (2006). Motivated Skepticism in the Evaluation of Political Beliefs. *American Journal of Political Science* 50(3), 755–769.