Better the Devil You Know:

An Online Field Experiment on News Consumption

Donghee Jo*

March 18, 2018

[Download the latest version here]

Abstract

This paper investigates the causal link between the public's self-selective exposure to like-minded partisan media and polarization. I first present a parsimonious model to formalize a traditionally neglected channel through which media selection leads to reduced polarization. In a world where the media heavily distorts signals with its own partisan preferences, familiarity with media biases is vitally important. By choosing like-minded partisan media, news consumers are exposed to familiar news sources. This may enable them to arrive at better estimates of the underlying truth, which can contribute to an alleviation of polarization. The predictions of this model are supported by experimental evidence collected from a South Korean mobile news application that I created and used to set up an RCT. The users of the app were given access to curated articles on key political issues and were regularly asked about their views on those issues. Some randomly selected users were allowed to select the news source from which to read an article; others were given randomly selected articles. The users who selected their news sources showed larger changes in their policy views and were less likely to have radical policy views—an alleviation of polarization—in comparison with those who read randomly provided articles. The belief updating and media selection patterns are consistent with the model's predictions, suggesting that the mechanism explained in the model is plausible. The findings suggest that the designers of news curation algorithms and their regulators should consider the readers' familiarity with news sources and its consequences on polarization.

*Massachusetts Institute of Technology. Contact: djo@mit.edu. I truly appreciate the guidance of my advisors, Ben Olken, Daron Acemoglu, Esther Duflo, and Abhijit Banerjee, without whom this paper would not have been possible. Funding was provided by the George and Obie Shultz Fund and JPAL Incubator Funding. I am grateful to Manseok Jo, Dohyeon Kim, Young Jae Woo, and Jaehoon Yang, who helped build and operate a mobile application for the experiment. I also thank Enrico Cantoni, John Firth, Chishio Furukawa, Arda Gitmez, Nick Hagerty, Greg Howard, Peter Hull, Gabriel Kreindler, Matt Lowe, Rachael Meager, Yuhei Miyauchi, Arianna Ornaghi, Frank Schilbach, Cory Smith, Michael Wong, Muhamet Yildiz, and the participants of MIT Political Economy Seminar for helpful comments.

1 Introduction

"Increasingly we become so secure in our bubbles that we start accepting only information, whether it's true or not, that fits our opinions, instead of basing our opinions on the evidence that is out there."

– President Barack Obama, in his farewell address, January 10, 2017

Political polarization is on the rise, and many find this trend troubling.[1] Politicians and academics alike blame the media's adoption and dissemination of political biases, as well as the public's self-selected exposure to like-minded media, as the main drivers of political polarization. President Trump tweeted that some liberal media are "the enemy of the American people," for their coverage of "fake news." As shown in the opening quote, President Obama described people as being in their "bubbles" where we only accept information "that fits our opinions." Academics agree. People are allegedly in their own "echo chambers" or "filter bubbles," exposing themselves to the information sources that are likely to confirm their preexisting views (Sunstein, 2001; Prior, 2007; Pariser, 2011). As such, many experts advocate enhancing "chance encounters"—random and balanced media exposure—to counteract the accelerating trend of polarization (e.g., Sunstein, 2017).[2]

In this paper, I demonstrate a traditionally neglected channel through which selective exposure to partisan media can contribute to an alleviation—not exacerbation—of polarization. In a world where the media distort signals, familiarity with media biases is vitally important. By choosing like-minded partisan media, news consumers are exposed to familiar news sources through which they can arrive at better estimates of the underlying truth. In other words, selective exposure can contribute to an alleviation of polarization via facilitated learning.

I formalize this idea in a simple model and empirically investigate the causal link. I created a South Korean mobile news application (hereafter, "app") and used it to set up an online field experiment. The results of the experiment show that selective exposure alleviates partisan polarization of public views.

[1] For evidence on rising polarization, see Hetherington (2001); Jensen et al. (2012); Pew Research Center (2014); Iyengar and Westwood (2015); Gentzkow et al. (2016); Gentzkow (2016); McCarty et al. (2016); Boxell et al. (2017). Political polarization has increasingly been infiltrating into personal lives (YouGov, 2008; Iyengar et al., 2012; Iyengar and Westwood, 2015).

[2] Entrepreneurs and some news media organizations have been responding to these requests. For example, there are applications such as *Read Across the Aisle* or *Flipfeed* that help news readers receive balanced chance encounters. *The Wall Street Journal* provides a comparison between the Facebook news feeds of liberals and conservatives. *The New York Times* has a series, "Right and Left: Partisan Writing You Shouldn't Miss," that covers both left-wing and right-wing articles on a designated topic.

I present a parsimonious Bayesian model in which the agent's objective is to learn the true state of the world by selecting a news source for gaining information. The selected news source sends a contaminated signal: a mixture of the truth and the media bias of the source. Since the media bias of a familiar source is better understood by the agent, the model predicts a selection of a relatively familiar source. The most important and perhaps surprising prediction of the model is that getting a signal from a relatively familiar source, compared with getting it from a randomly chosen source, alleviates polarization. Specifically, the agent learns more about the truth as long as she has an approximately correct assessment of the sources' media biases.

I test some of the implications of the model using the experiments conducted in the aforementioned mobile app. The users of the app were given access to curated articles on key political issues in South Korea, and were regularly asked about their views on the issues. I explore two types of random variation of the user experience. First, every user ($N_1 = 1420$) was given randomly chosen articles during at least a partial period of the experiment. This allows random variation in the characteristics of the articles and their sources, which is used to establish the adequacy of the suggested model. Second, the users who remained in the app past five rounds of reading articles ($N_2 = 352$) were divided into several treatment groups, varied by two key aspects of the user experience. First, some users were allowed to choose the news source from which to read an article while others were given randomly selected articles. Second, when allowed to choose their sources, readers were shown either (a) the names of the news sources, (b) the news sources' average positions on the issue at hand, or (c) both.

I begin the empirical analysis by establishing the adequacy of the suggested model using the first type of random variation: random provision of news articles. The belief updating patterns observed in the data are consistent with the predictions of general Bayesian models with heterogeneous priors. The reader's view reported after reading an article is affected both by the article's position and the view reported prior to reading the article, and is less influenced by the article if she reported high confidence in her view prior to reading the article.

I then focus on the treatment groups that could select sources based on their names. They choose news sources that were likely to represent the party's views that the reader supported. This result suggests that the readers might associate partisan alignment with familiarity in the context of this experiment. Here, as in many other contexts, a partisan divide is the defining feature of political polarization. However, note that the insight of this paper can be applied to settings where other characteristics, such as religion or ethnicity, are more important. The amount of selective exposure found in the data is comparable to that of the US in terms of the isolation index defined in Gentzkow and Shapiro (2011).

Having established the existence of selective exposure to partisan media, I move on to the main question of this paper: the effect of selective exposure on polarization. To this end, I compare the evolution of beliefs held by the treatment groups. In particular, the main comparison is between users who could select sources based on their names (collectively called the "Source-Name Group."), and users who had random provision of articles (collectively called the "No-Choice Group").[3] Comparisons between the two groups identify the effect of unregulated media selection (Source-Name Group)—as opposed to random and balanced media exposure (No-Choice Group)—on political polarization.

Consistent with my model's implication about polarization, selective exposure to partisan media facilitates learning. The distance between the users' positions reported before and after reading an article is 17% larger (significant at the 5% level) for users who could select based on source names (Source-Name Group) compared with users who were given randomly selected articles (No-Choice Group). Users who could select based on source names also show a greater decline in their proportion of extreme policy views (a decline of 7 additional percentage points, statistically significant at the 5% level).[4] The results suggest that selective exposure to partisan media causes an alleviation of opinion polarization, contrary to the traditional view.

I find suggestive evidence that the channel described above—familiarity facilitates learning and reduces polarization—is the underlying mechanism of how media selection results in moderated policy views. Using ideological distance between each news source and each user as a proxy for unfamiliarity, I show that the policy views are affected more by the articles when the news sources are more familiar to the readers. This result is consistent with the model in which the information value is perceived to be higher when the article is written by more familiar media. I also empirically test a competing mechanism, in which the users are assumed to be "blind followers" of their favored parties, changing their views only in response to in-group news sources. In this case, selective exposure may induce convergence to the supported party's position, which can be interpreted as an exacerbation of polarization. I create several proxy measures of parties' positions and test whether the users' policy views are differentially—between treatment

[3]The comparison between these groups and a remaining group—users who could select only based on source positions on a particular issue ("Source-Position Group")—is also reported in the Results Section. However, I do not emphasize this comparison, which is mostly noisy due to the small size of this group. Merging this group into either one of the larger groups is unnatural for the following reasons. The Source-Position Group shows no sign of partisan selective exposure, which differentiates it from the Source-Name Group. On the other hand, the Source-Position Group is inherently different from the No-Choice Group, as the group could select their sources, albeit with limited information.

[4]Policy views are considered to be extreme if the reader reported a view that is farthest to the left or right of the policy view spectrum bar in the app. See Section 3 to see how readers reported their policy views.

groups—approaching the favored party's position. I do not find any evidence to support this claim.

Taken together, this paper suggests a simple but timely policy implication. When providing curated news articles, the mediator (or the mediating websites) should consider the receiver's familiarity with the news sources and its potential impact on polarization. In the short run, providing articles from the readers' familiar news sources could be better than providing balanced articles from unknown news sources. In the long run, an effort to build readers' familiarity with a small number of media can be beneficial. In a world where news curation services are ever more widespread (e.g., social network services such as Facebook and Twitter), this paper provides a simple insight relevant to both designers of curation algorithms and to regulators.

Related Literature. The results in this paper broadly speak to a rapidly growing literature that studies political polarization. Although there is abundant descriptive evidence on the increasing trend of polarization of political elites (Poole and Rosenthal, 1991; Jensen et al., 2012; Gentzkow et al., 2016; McCarty et al., 2016) and the public (Hetherington, 2001; YouGov, 2008; Iyengar et al., 2012; Pew Research Center, 2014; Iyengar and Westwood, 2015; Gentzkow, 2016; Boxell et al., 2017), there is little evidence on the causes of polarization. In particular, there has been speculation that selective exposure to like-minded media is one of the main drivers of political polarization (Sunstein, 2001), and partisan media are shown to have the potential to change the behaviors of voters (DellaVigna and Kaplan, 2007; Martin and Yurukoglu, 2016; Adena et al., 2015; Durante et al., 2015).[5] There is also abundant descriptive evidence on the prevalence of selective exposure (Iyengar and Hahn, 2009; Gentzkow and Shapiro, 2011). However, rigorous causal evidence linking selective exposure and the public's partisan polarization is limited.

The closest studies to this paper are Boxell et al. (2017) and Druckman et al. (2012). Boxell et al. (2017) reject the claim that the Internet, which makes selective exposure to like-minded media extremely easy, is a direct cause of rising polarization. The main evidence for their rejection is the fact that the age group who experienced the sharpest rise in opinion polarization comprises those older than 75, unlikely subscribers to information technology. In this paper, I provide more rigorous causal evidence on the relationship between selective exposure and opinion polarization. In a lab experiment, Druckman et al. (2012) provide evidence that deliberate selections of news articles contribute to the stability of preexisting opinions. In this paper, I explicitly introduce the notion of partisanship to directly study the impact of selective exposure on

[5]While media markets have been studied in the context of political bias, the focus has mostly been on the supply side(Mullainathan and Shleifer, 2005; Gentzkow and Shapiro, 2006, 2010) or on political elites (Campante and Hojman, 2013), rather than on consumers of news articles and the evolution of their opinions.

partisan polarization, an outcome that is not studied in Druckman et al. (2012).[6]

Theoretically, the model in this paper is an application of Sethi and Yildiz (2016a,b) on polarization. The authors build on the insight of Acemoglu et al. (2016), who demonstrate the importance of understanding the receiver's belief about the messenger. Sethi and Yildiz (2016a,b) focus on studying the tradeoff between expertise and familiarity in the selection of information sources, where familiarity is defined similarly as in this paper's model. Other related theoretical papers include Calvert (1985) and Suen (2004), where theoretical foundations of rational selective exposure are laid out. In these papers, the reason for selective exposure is to maximize the value of information, similar to the model in this paper. The difference is that the value of information in Calvert (1985) and Suen (2004) is driven by the discreteness of the actions in their models—the information is valuable only insofar as it can change the actions of the agents.

Methodologically, this study also relates to a burgeoning literature that conducts online field experiments. See Chen and Konstan (2015).

The remainder of the paper proceeds as follows. Section 2 presents a parsimonious model and its implications. Section 3 describes the setting, experimental design, and data. Section 4 provides the main results of this paper. Section 5 explores the mechanisms and Section 6 concludes.

2 Theoretical Framework

In this section, I provide a theoretical framework to formally demonstrate a potential mechanism through which unregulated media selection leads to reduced polarization. The model not only produces testable predictions but also sheds light on the generalizability of the results of this paper by providing precise conditions under which this mechanism can operate. The model purposefully simplifies the media environment to emphasize the essence of the mechanism.

2.1 Model

For each policy issue j, there is a common truth $\theta_j^* \in \mathbb{R}$, and the agents hold their own subjective beliefs about θ_j^* at the beginning of the period. This prior belief can be represented by a normal

[6] In the lab experiment that the authors conducted, an article with a particular framing is provided to the subjects at the beginning of the four-week experimental period, and another article that frames the issue differently is provided at the end. When allowed to select their sources of information based on article titles during the intermediate period, the subjects were more likely to choose titles that were consistent with the framing given to them at the beginning of the four-week period. As a result, at the end of the experimental period, the subjects who were allowed to select articles were more likely to lean toward the framing of the initial article, despite the freshness of the latest article.

distribution with support $\mathbb{R}$: $\theta_j \underset{i}{\sim} N\left(\theta_{ij0}, \frac{1}{\tau_{ij0}}\right)$. θ_{ij0} is the mean, and τ_{ij0} is the precision of the prior belief. Higher τ_{ij0} means that the agent has greater confidence in her prior belief.

For the supply side, assume that there are two news sources, $p \in \{F, A\}$. F represents a familiar news source (e.g., the *New York Times* for liberals or Fox News for conservatives), and A represents an alien one.

A signal that news source p sends to agent i on issue j is given by:

$$s_{pij} = \theta_j^* + I_{pj}^* + \varepsilon_{pij} \tag{1}$$

where $\varepsilon_{pij} \sim N\left(0, \frac{1}{\tau_s}\right)$ is the idiosyncratic noise that is added to the signal. τ_s is the precision of the signal, and it represents the journalistic competence of the source in collecting and conveying information. It also includes the agent's ability to understand the signal. For simplicity, these idiosyncratic noise terms are assumed to have the same precision (τ_s).

The signal that p sends (s_{pij}) is a sum of θ_j^*—the common truth—and I_{pj}^*. I_{pj}^* is the "media bias" added to the truthful signal. It is the position p advocates. For example, *Breitbart.com* can be thought of as having positive I_{pj}^*, adding right-wing biases to the articles it publishes. The addition of media bias (I_{pj}^*) can be achieved by omission of evidence or by careful control of tone or voice. I_{pj}^* is not directly observable to i, and each agent holds a subjective belief about it: $I_{pj} \underset{i}{\sim} N\left(I_{pji}, \frac{1}{\tau_{Ipji}}\right)$. I_{pji} is the mean, and τ_{Ipji} is the precision of the belief about I_{pj}^*. Higher τ_{Ipji} means that i has a better understanding of the bias p adds to the signal—the "style" of the news source is familiar to i.

In this section, I analyze a single-period belief updating[7] with a timeline as follows. An issue (j) is randomly selected by nature. The agent selects a news source, and she gets a contaminated signal—s_{pij} as in Equation (1)—from the chosen source. The agent updates her belief about θ_j^* using Bayes' rule, and chooses an action $a_{ij} \in \mathbb{R}$ to maximize utility:

$$U(a_{ij} \mid \theta_j^*) = -(a_{ij} - \theta_j^*)^2 \tag{2}$$

This is equivalent to a setting where the agent genuinely tries to learn the truth as accurately as possible. In the language of political psychology, agents have *accuracy* motives, instead of *directional* motives (Leeper and Slothuus, 2014).

Understanding the selection of the news source requires backward induction—let us begin by assuming that news source p is selected by i to get a signal about issue j. It takes tedious, albeit

[7] A long-run extension of this model is discussed in Section 2.2.

straightforward, algebraic maneuvers to derive[8] the posterior distribution after getting a signal, s_{pij}:

$$\theta_j \,|\, s_{pij} \underset{i}{\sim} N\left(\theta_{ij1}\,,\,\sigma^2_{ij1}\right) \tag{3}$$

where

$$\theta_{ij1} = \frac{\tau_s\tau_{ij0} + \tau_{ij0}\tau_{Ipji}}{\tau_s\tau_{ij0} + \tau_{ij0}\tau_{Ipji} + \tau_{Ipji}\tau_s}\theta_{ij0} + \frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0} + \tau_{ij0}\tau_{Ipji} + \tau_{Ipji}\tau_s}\left(s_{pij} - I_{pji}\right) \tag{4}$$

$$\sigma^2_{ij1} = \frac{\tau_s + \tau_{Ipji}}{\tau_s\tau_{ij0} + \tau_{ij0}\tau_{Ipji} + \tau_{Ipji}\tau_s} \tag{5}$$

As can be seen in Equation (4), the mean of the posterior belief, θ_{ij1}, is a convex combination of the prior mean (θ_{ij0}) and the bias-deducted signal ($s_{pij} - I_{pji}$). $\frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s}$ is the weight on the signal, and the remaining weight is on the prior. Propositions 1 and 2 describe this and provide comparative statics about the weight on the signal.

Proposition 1. *(Bayesian Updating) The posterior mean is a function of the prior mean (θ_{ij0}) and the bias-deducted signal ($s_{pij} - I_{pji}$). The weight on the signal is decreasing in τ_{ij0}.*

Proof. For the first sentence, see Equation (4). It is then trivial to show that $\frac{\partial}{\partial \tau_{ij0}}\left(\frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s}\right) < 0$. □

Proposition 2. *(Familiarity and updating) The weight on the signal, used to form the posterior mean, is increasing in τ_{Ipji}.*

Proof. It is trivial to show that $\frac{\partial}{\partial \tau_{Ipji}}\left(\frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s}\right) > 0$. □

The prediction that the posterior mean is a function of the prior and the signal is consistent with typical Bayesian models. The prediction that the weight on the signal decreases in prior confidence (τ_{ij0}) is also consistent with typical Bayesian models. The prediction on τ_{Ipji} (Proposition 2) is a unique feature of this model. The signal has greater influence on the posterior belief if it is from a source that the agent understands better. The predictions in Propositions 1 and 2 are empirically verified in later sections of this paper.

Now that we understand the final stage—belief updating—we can analyze the preceding stage: the agent's source selection. For this, I make the following assumption.

Assumption 1. *(familiarity) $\forall j$, $\tau_{IFji} > \tau_{IAji}$*

[8] See Appendix A for the derivation. I assume that the prior distribution of I_{pj} and θ_j are independent. This implies that the prior information that i has on issue j is not from source p. Relaxing this assumption does not affect the main results of this paper—although it complicates the algebra.

This assumption indicates that the agent has a more accurate assessment of the bias of the familiar source than of the alien one. Theoretically, this assumption has been explored in Sethi and Yildiz (2016a,b). Empirically, it has been shown that people have more nuanced knowledge about the biases of in-group news sources compared to that of out-group sources (Stroud et al., 2014).[9] The assumption implies that a signal from the familiar source has better information value for i, which makes her choose the familiar source as described in Proposition 3.

Proposition 3. *(Selective Exposure) Suppose Assumption 1 is true. The familiar source will be chosen.*

Proof. See Appendix A. □

In the mobile app experiment described in this paper, some subjects were allowed to choose their news sources and are predicted to choose the familiar source (Proposition 3). The remaining subjects who could not choose their news sources can be regarded as encountering the familiar source and the alien source with equal probability. The propositions below predict the difference in the evolution of beliefs between these two groups. For Proposition 4, I also assume that i's assessment of the media bias of p, I_{pji}, is approximately correct.

Assumption 2. *(Approximately Correct Bias Assessment)* $I_{pji} = I^*_{pj} + \nu_{pji}$, *where* $\left|\nu_{pji}\right| < M_j \frac{1}{\gamma_{pij}}$ *for some* $M_j > 0$, *where* $\gamma_{pij} \equiv \frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s}$.

Assumption 2 implies that the bias assessment of the agents is not too erroneous, and the absolute value of the error ($\left|\nu_{pji}\right|$) is bounded by a constant that is proportional to the weight (γ_{pij}) that the agent puts on the signal when forming the posterior mean (Equation (4)). When the agent has a greater familiarity with the media bias (e.g., higher τ_{Ipji}), the error of her bias assessment is subject to a tighter bound. For example, a liberal is expected to have a smaller error in her assessment on the media bias of *the New York Times* compared with her assessment on the bias of Fox News. I discuss a potential consequence of a violation of this assumption below (Proposition 5).

[9] Specifically, people appear to understand the media-specific bias of in-group sources, whereas they do not know the differences in biases between out-group sources. This is often called "out-group homogeneity." Formally, this can be directly linked to Assumption 1 if we take a Bayesian hierarchical model in which an individual media outlet has a separate bias I_{mj}, which is known to be a realization of $N(I_{Fj}, \tau)$ or $N(I_{Aj}, \tau)$—depending on whether the source is in-group (F) or out-group (A). If i has taken more signals from familiar sources ($m \in F$), she must have a better idea about the distribution of I_{mj}. On the other hand, if i has taken virtually no signals from alien sources, when asked about her best estimation of I_{mj} ($m \in A$), she will simply say I_{Aj} regardless of m, making her appear to have no understanding of the nuanced differences between out-group sources. This gives Assumption 1. Stroud et al. (2014) show that there is a positive correlation between media familiarity and the nuanced understanding of the media's bias.

For notational convenience, let us define the error of the prior as $\eta_{ij0} \equiv \theta_{ij0} - \theta_j^*$ or, equivalently, $\theta_{ij0} = \theta_j^* + \eta_{ij0}$. I put no structure on the distribution of η_{ij0}, except that it is independent of ν_{pi} and ε_{pij}.

Proposition 4 shows that those who are allowed to select their news source learn more on average (i.e., the distance between the prior mean and the posterior mean is higher when source selection is allowed).

Proposition 4. *(Learning) Suppose Assumptions 1 and 2 are true with some M_j. There exists $M_j > 0$ such that allowing source selection, as opposed to equal-chance exposure to either F or A, facilitates learning ($|\theta_{ij0} - \theta_{ij1}|$ is larger) on average.*

Proof. See Appendix A. □

Proposition 4 implies that an agent learns more and approaches the truth (θ_j^*) faster when she is allowed to choose her news source. In order to derive this, we needed to assume that the agents have approximately correct assessments of the media's biases, where the approximation is more accurate for the familiar news sources (Assumption 2). In the proposition below, I explore what could go wrong by taking an extreme case where the agent perceives that the familiar news source adds no bias to the signals.

Proposition 5. *(Convergence to Group's Position) Suppose Assumption 1 is true. Instead of Assumption 2, assume that $I_{Fji} = 0$ (a degenerate distribution), and $E\left[I_{Aji}\right] = I_{Aj}^*$. Also assume that the agents are not biased on average at the beginning ($E\left[\eta_{ij0}\right] = 0$). Allowing the source selection (as opposed to an equal-chance encounter) results in a faster convergence to the media bias of the source, I_{Fj}^*, on average.*

Proof. See Appendix A. □

Propositions 4 and 5 provide contrasting implications about the effect of selective exposure on opinion polarization. Insofar as the agents have an approximately correct assessment of the bias of the news sources (Assumption 2), selective exposure facilitates convergence to the truth. This means the beliefs of two people who are in different groups become closer to each other (because they are both closer to the truth) after reading an article when source selection is allowed. The reduced distance between them is clearly an alleviation of polarization.

On the other hand, if they have a systematically wrong perception about the biases of the news sources (e.g., imagine a liberal firmly believing that *the Huffington Post* is telling the truth, the whole truth, and nothing but the truth), then exposure to the familiar news source results in faster convergence to the source's media bias. As a result, the agent's group membership is easier to identify based on their opinions. This can be interpreted as an exacerbation of polarization.[10]

The "group," used above to define polarization, can take many different forms depending on the specific context. The empirical sections of this paper emphasize a partisan divide, because that appears to be the most salient determinant of media selection in the context of this paper. In other contexts, groups can be divided by, for example, religious beliefs, ethnic identity, or regional boundaries, as long as such group identity is most relevant for media selection and political polarization in that context.

2.2 Discussions on the model

The highly stylized model provided in Section 2.1 warrants a discussion. First and foremost, it is not reasonable to expect that such a parsimonious model can comprehensively represent real-life news consumption. However, the model provides important insights that the empirical results of this paper corroborate, while maintaining a standard structure of Bayesian models. The model clarifies the conditions under which the main mechanism—unregulated media selection facilitates learning due to the agent's familiarity with the biases of selected media—can operate. First, the agents need to have reasonably close "ideal points" for the policy issue in question. Although I assume common truth in the model above, the citizens' ideal points can vary in reality. We can extend the model to incorporate this by adding a term, b_{ij}, that represents individual characteristics to the utility function:

$$U(a_{ij} \mid \theta_j^*, b_{ij}) = -(a_{ij} - \theta_j^* - b_{ij})^2$$

Then the ideal point for i becomes $\theta_j^* + b_{ij}$. The condition for alleviation of polarization in this case becomes that b_{ij} has reasonably small variance across population relative to the distribution of the initial beliefs, θ_{ij0}. Second, the agents have approximately correct assessments on media biases (Assumption 2). Otherwise, facilitated learning is not guaranteed, and media selection can cause an exacerbation of polarization (Proposition 5). The conditions shed light on the generalizability of the main mechanism. Given this assumption, we may expect an alleviation of polarization among a particular subset of the population, which is explored in Section 5.

[10]Convergence of policy views to the favored party's position is one of the most salient forms of polarization of public opinions in the last two decades in the US. (Gentzkow, 2016). Baldassarri and Gelman (2008) argue that this trend is not because the general public changed their opinions (and is because the public switched their supporting parties to have better aligned policy view and parties adjusted their policy positions to better serve the supporters), consistent with the results of this paper—see Section 6.

In a model where media can choose to be silent about certain issues, the predictions can be different (Banerjee and Somanathan, 2001). In the context of this paper, this is an unlikely scenario, because each issue has been covered in almost all of the media appearing in the mobile app.

Note that selection of familiar sources can be driven by other reasons than the model above demonstrates. In the model, the familiar news source is selected because the agent hopes to learn the truth. This objective is better achieved when the agent acquires information from news sources whose bias is more familiar to the agent, thereby making their signals more informative. It is important to acknowledge that there are other potential motives for selective exposure that are not included in the suggested model. For example, people may get direct utility from the act of reading articles from familiar media. Empirically distinguishing the motives of selective exposure is beyond the scope of this paper. Although it is theoretically concise to have the same factor—familiarity to media bias—driving both the selection of familiar news sources and facilitated learning, it is not strictly necessary for the main results. The agent needs to be sophisticated to the extent that she takes the media biases into account, but she does not need to choose the familiar news sources solely due to the desire of facilitated learning.

Although I have chosen a simplification by describing only a one-shot Bayesian updating, learning is a multi-stage process in reality. Even if the agent begins with the wrong perception of her familiar media's bias—as in Proposition 5—she may learn its true bias through receiving sufficient amounts of signals in the long run. The process of learning true biases will be facilitated if she attempts to learn deeply about a particular issue, which makes it optimal to eventually take signals from other available sources (Liang et al., 2017).[11] In the long run, therefore, the assumption of a correct bias assessment (Assumption 2) may be more reasonable.

Due to my focus on consumers of news articles, this paper does not explore the strategic motives of the media. In particular, if the true biases of the sources are eventually revealed to the readers, why would the news media bother to add media biases to the signals? There are many possible explanations. For example, the news producers might get pleasure from adding media biases to the articles. It is also possible that a large enough part of the population is yet to reach the long-run steady state in which media biases are precisely known. In this case, adding biases can be beneficial for news producers either because it makes the readers biased or because it lets

[11]This is intuitive because the signals from one source are inevitably correlated. See Liang et al. (2017) for theoretical exposition. The authors also prove that, under some regularity conditions, a myopic agent and a forward-looking agent select approximately (or exactly, with additional assumptions) the same news sources. This may justify the use of myopic utility maximization in the model above.

the media sustain its reputation of providing high-quality information.[12] Finally, the pleasure derived from reading like-minded news may be large enough for readers to choose like-minded news sources.

2.3 Predictions

The model in Section 2.1 gives the following predictions, listed in order of their appearance in the empirical sections below. The Results Section—Section 4—provides evidence for predictions [1]-[3]. Predictions [4] and [5] are presented in Section 5 where I discuss mechanisms.

[1] *Belief updating patterns are consistent with typical Bayesian updating. The posterior mean is affected both by the prior mean and the signal, and it is affected less by the signal when the confidence on the prior is higher. (Proposition 1)*

[2] *Selective exposure exists (i.e., familiar news sources are selected when given a choice). (Proposition 3)*

[3] *When allowed to choose news sources (as opposed to getting randomly selected articles), readers change their views to a greater degree as long as they can properly deduct the biases of the articles (Proposition 4).*

[4] *Readers are more affected by the article's position when the articles are from familiar sources. (Proposition 2).*

[5] *If the media biases of news sources are not correctly assessed, selective exposure to familiar media can result in faster convergence to supporting group's position (Proposition 5).*

3 Setting, Experimental Design, and Data

3.1 Background

South Korea has a representative democracy with a presidential system. Since the end of military dictatorship in 1988, two parties have led the executive branch in turn. A right-wing party (called

[12]When the target population is biased, it may be optimal for the media to provide biased information—see Gentzkow and Shapiro (2006).This paper does not take a stand on whether the bias of the media is supply- or demand-driven. For an empirical analysis on this, see Gentzkow and Shapiro (2010). The appendix in Chiang and Knight (2011) discusses the theoretical foundation of both channels.

the “Saenuri Party” at the time of this experiment), which represents neoliberalism and social conservatism, brought four out of seven presidents to power. The remaining three presidents were from a center-left party (the Democratic Party), which advocates economic equality and social liberalism. In addition, there were two smaller parties, the People’s Party (center) and the Justice Party (left) during the period of the experiment.[13]

Political polarization is one of the most serious sources of societal conflict in South Korea (Korea Institute of Public Administration, 2015; p.165). There is survey evidence from 2004 and 2008 (Lupu, 2015) demonstrating that the partisan polarization of South Korea had been more severe than that of the US. Like many other countries, news sources in South Korea are believed to be politically biased, and are accused of being the main driver of political polarization. In a survey (Korea Institute of Public Administration, 2015), more than half (61%) of the respondents said that “the press is unfair” (p.142).

Traditional media, such as TV channels and newspapers, remain the primary news sources in South Korea. According to a recent survey (Korea Press Foundation, 2016), South Koreans on average spent 50% of their daily news consumption time (83 minutes) on TV, and 40% on newspapers (including online newspapers). The rapid rise of news consumption via mobile phone is noteworthy and made this experiment possible. In 2016, South Koreans spent 17 minutes per day reading newspaper articles on their phones, a 148% increase from 2011.

3.2 Setting: Mobile Application

To study the impact of media consumption on policy views, I developed an iOS mobile application and distributed it to the general public of South Korea from February to November, 2016. The app, “Spoon,” is primarily a news curation service. I collected articles daily on eight politically disputed issues from a wide spectrum of news sources spanning left to right. At each point in time, there were 10-20 articles on each issue and the articles were consistently replenished on a weekly basis. Figure 1 presents the list of issues covered in this application.[14] All eight issues had been of great interest to the general public throughout the experimental period. Therefore,

[13]President Park of the Saenuri Party led the executive branch during most of the experimental period. The Saenuri Party also had a parliamentary majority for the first three months of the experiment, but they lost it in the April, 2016 General Election. Shortly before the end of the experiment, there was an outbreak of one of the largest political scandals in the history of South Korea, which resulted in the impeachment of then-president Park. The experiment was mostly unaffected by this scandal.

[14]There were five additional issues (“temporary issues”) covered in the mobile app for a relatively short period of time, and I provided fewer articles on these issues (3-5) at each point in time. These issues are excluded from the analysis in this paper.

every issue had been covered by almost all the news sources that are featured in the app during the experimental period.

The users had additional incentives to use the app: I provided several types of aesthetically pleasant summary statistics on, for example, how their beliefs evolved over time and which articles and news sources have been most influential to them. These features contributed not only to attracting and retaining users, but also to incentivizing truthful revelations of the users' beliefs. There was no monetary compensation to use the app, except for 37 users (2.6% of the sample) who responded to the intensive follow-up e-mails.[15] Most of the users connected voluntarily to the app, presumably to read well-written articles and to see summaries of their policy views.

The app was distributed to the general public for free; anyone who had access to an iPhone could participate in the study. Facebook advertisements were the primary source of user inflow. A total of 2,627 people installed the app; 1,420 of them finished reading the first article and were included in the sample.

3.3 Experimental Design

I exploit two sources of randomness in this experiment. First, an article was randomly selected for each user in a subset of the experimental period. All 1,420 users in the sample are included in the analyses that exploit this random provision of articles. Second, each user was assigned to one of seven treatment groups. Only those who completed the five-day grace period (352 users) were included in this randomization.

As shown in Figure 2, upon installation and baseline survey, each user had a five-day grace period. During the grace period, a randomly selected article about a randomly selected issue was provided each day.[16] The grace period was introduced to alleviate the attrition problem by screening out those who were going to drop out early. See Schilbach (2015) for the use of a grace period in a different context.

[15] This intensive follow-up is not exploited in this paper due to the limited response to it. For those who are curious about the exact procedure: there were low- and high-intensity groups in this follow-up attempts. The low-intensity group could choose between (i) a 4,000 Won ($\approx$ \$4) reward with 100% probability and (ii) a 50,000 won ($\approx$ \$50) reward with 10% probability. The high-intensity group had the same choice except the reward amounts were doubled.

[16] The users were not required to use the app every day; they chose whether to use it on a particular day. The users could read multiple articles per day, and starting from her second article a user could select the next issue of interest. The next article was also randomly selected for the chosen issue. Since I allowed people to read multiple articles per day, the median term between rounds on an issue was 5.8 days, which is shorter than 8 to 9 days—the expected term between rounds if everyone read only one article per day.

I assigned the 352 users who finished the grace period into one of seven treatment groups—G1-G7 in Figure 2. The everyday experience of the treatment groups was slightly different, as described in Figure 2.[17] First—a common step to every treatment group—the issue of the day for each user was randomly selected by the app. The first three groups were allowed to select the news source from which to read an article about the issue of the day. The remaining four groups were given randomly selected articles. When making source selections (G1-G3), the first two groups (G1-G2) could see the names of the news sources while G3 could not (Panel (a) of Figure 3 versus Panel (b)). The average positions of the sources[18] on the issue at hand were shown only to G1 and G3, and not to G2. For those who were not allowed to select the news source (G4-G7), there was a screen (Panel (c)), immediately before reading the randomly chosen article, showing relevant information about the news source that published the article. The information about the news source shown to the readers varied, as described in Figure 2. Although the name of the news source was not salient to G3, G6, and G7 before reading an article, everyone could eventually determine the news source because the articles always indicated the source's name at the end of the text. Most articles also indicated the source's name at the beginning and in the middle of the text.

In most of the analyses below, I aggregate the groups into three larger groups based on their source selection patterns and supplement the analysis with subgroup comparisons in the Appendix. The Source-Name Group includes G1 and G2—they were allowed to select the news source based on source names. Behaviors of G1 and G2 are qualitatively very similar—additional information about sources' positions did not make a notable difference. The Source-Position group consists of only G3, which could not see the names (thus it was harder for them to identify the media that are likely to advocate the supported party's view), but could see source positions (thus it was easier for them to identify the prior-confirming media). The remaining groups (G4-G7) were not allowed to choose news sources and were given randomly selected articles. They were merged together as the No-Choice Group—the subgroups of the No-Choice Group show similar qualitative patterns. All group comparisons in the Results Section are based on three larger groups. The analysis by finest subgroups is provided in an Appendix whenever necessary. The Source-Name Group and the Source-Position Group are not merged together because they show substantially different selection patterns, as shown in Section 4.2.

[17] There was a cross-randomization where some of the users had access to the distribution of the positions of other users. This treatment arm is not the focus of this paper.

[18] The average position of a news source is the average of the positions of the articles written by the news source on the issue at hand. The average position of each article is calculated based on users' reports after reading the article—see Panel (e) of Figure 3.

3.4 Data Collection

All the data that I use in this study were collected in the app. After the (random or deliberate) selection of a news source, the users read an article from that source. Given the news source, an article was randomly selected among the pool of articles that were written by the source. The users then reported their positions on the issue and the level of confidence in their positions (Panel (d), Figure 3), and evaluated the quality and the position of the article (Panel (e), Figure 3). The beliefs and positions are continuous measures between 0 and 1. For beliefs, the app asks the user to "move the scroll bar to denote your view on issue <*IssueName*>." The app provided a rough cardinal benchmark to the positions on the horizontal belief bar. As the user moves the scroll bar, the app explained what each position meant in terms of specific policies or attitude on the issue. To be more specific, the horizontal belief bar was divided into five segments where each segment had a corresponding explanation. See Figure 4 for an example of such explanations. The confidence (Panel (d), Figure 3) and the article quality (Panel (e), Figure 3) are discrete "star" measures, where half stars are allowed. In the empirical results below, the confidence measures are also transformed into $[0, 1]$ for easier interpretation.

I also measured the total time that each user spent reading articles. In addition to the time spent on reading the article, the measure includes the time spent on the selection screen (Panels (a) and (b)), the media information screen (Panel (c)), and the evaluation screens (Panels (d) and (e)). Unfortunately, I do not have separate measures for the time spent on each screen.

3.5 Summary Statistics and Experimental Validity

Appendix Table 1 provides sample statistics. As expected for typical app users, the sample is younger (31 years) than the South Korean population (41 years) on average. It is also more liberal, less likely to watch TV news (56% in my sample vs. 83% for the population), and slightly more likely to use the Internet for news consumption (85% vs. 74%). All population figures are from the Korea Press Foundation (2016). One should be careful in generalizing the results of this paper to a wider population or to other nations.

The median time each person spent on the app each day was 2.2 minutes. This is 3% of the self-reported news consumption time of South Koreans (Korea Press Foundation, 2016). However, considering the facts that (i) only 30% of the total news consumption was on political news and (ii) self-reported news consumption time is exaggerated by a factor of 7-8 for the app's target population (Prior, 2009), 60-70% of the total political news consumption can be accounted for by the news consumption in the app. Many of the subjects reported extreme policy views at the baseline—37% of the reported policy views are either 0 or 1, the most extreme views in either direction that could be reported in the app.

Appendix Table 2 provides the randomization check for the random provision of articles, and Appendix Table 3 provides the check for random assignment to treatment groups. It appears to be balanced for both. Note that article pool fixed effects are included in the regressions of Appendix Table 2. The article position is exogenous only after including these fixed effects, because the news articles are removed and replenished over time to keep the app up-to-date. The position of the article is crowd-sourced—it is the average of the user-evaluated article positions.

There is a concern about attrition for the analyses that explore the random assignment to treatment groups. When I compare the experimental groups to identify the causal effect of selective exposure on polarization, I use user$\times$issue combinations with at least one article read after the group assignment. As shown in Appendix Table 4, there is 39% attrition for this comparison, though encouragingly, the baseline characteristics or assignment to treatments are uncorrelated with the attrition.

In comparison to the users who drop out during the grace period, the users who completed the grace period marked high scores in a test of political knowledge (Appendix Table 5). They are also more liberal and spend more time on news consumption according to their baseline self-reports. The group assignment was made among the users who completed the grace period, so this is not an internal validity concern. However, we should take this sample selection into account when discussing the generalizability of the empirical results obtained from the comparison between the treatment groups.

4 Results

Subsection 4.1 below establishes that the user's belief updating patterns are consistent with typical Bayesian models, giving confidence in using a simple Bayesian model to analyze the mechanism. Subsection 4.2 investigates the news-source-selection patterns. Subsection 4.3 provides the main results of this paper: Unregulated media selection, as opposed to random exposure, facilitates learning and alleviates polarization.

4.1 The belief updating patterns: consistent with typical Bayesian models (Prediction [1])

I begin the empirical analysis by establishing the adequacy of the suggested model. In this subsection, I use the subsample that was given randomly selected articles. This subsample includes everyone during the grace period and the users in the No-Choice Group after the grace period.

Because the provision of articles was random, there is no spurious correlation between the article position and the reader's view reported after reading an article.

I test some basic predictions of typical Bayesian models: (i) the reader's view reported after reading an article is affected by the article's position and the view reported prior to the reading, and (ii) the article has lower influence if the reader reported high confidence in her view prior to reading the article.

I first estimate:

$$Belief^{post}_{ijr} = \alpha_j + \alpha_r + \alpha_A + \gamma_1 \cdot Belief^{prior}_{ijr} + \gamma_2 \cdot s_{ijr} + \mathbf{X_i}\beta + \varepsilon_{ijr} \tag{6}$$

where $Belief^{post}_{ijr}$ is the posterior belief on j reported by i after article-reading round r, $Belief^{prior}_{ijr}$ is the belief reported prior to reading this article, s_{ijr} is the position of the article (signal), and $\mathbf{X_i}$ is the vector of user controls. Prior beliefs are recorded either from a post-reading survey in the previous round, or a random survey that often appears after reading an article about other issues. The position of the article is crowd-sourced—it is the average of the user-evaluated article positions. The pool from which an article is randomly drawn changed over time—I add article pool fixed effects (α_A) to deal with this. α_j is the issue fixed effect, and α_r is the round fixed effect.

Columns 1-3 of Table 1 provide the results. Each column comes from a separate OLS regression with a different combination of fixed effects and user controls as specified. Standard errors are clustered by user. Since the results are similar across the article-reading rounds, I pool them together.

User beliefs are affected by the article's position. The posterior belief consists of approximately 7-8% of the signal, 60-70% of the prior, and other unexplained factors such as pure statistical noise. Results are similar with and without fixed effects and user controls.

One may wonder whether the effect is highly transient, and the knowledge or attitude will be quickly dissipated as time goes by. This is not the case. As can be seen in Appendix Table 6, the effect of the first article persists even after reading the second article (median 2.3 days later) and after reading the third article (median 12.9 days later). The result above is also robust to (i) dropping the observations where the users hold extreme prior views ($\in \{0,1\}$) before reading an article (Appendix Table 7) and (ii) keeping only the first round observations (Appendix table 8) to minimize the concern about attrition—excluding the grace period (Columns 5 and 6 of Appendix Table 8), the attrition rate after starting reading an article is about 1.5%, and therefore the results are not likely driven by the attrition.

Taking the theory more seriously, one needs to take into account the bias deduction of the agents (see Equation (4)). To control for this, I add news source fixed effect (Column 2 of Appendix Table 9) and its interaction with the dummy variable indicating the most trusted party at the baseline (Column 3 of Appendix Table 9). The results are similar—user beliefs are positively affected by the position of the article and the magnitude of the effect is similar.

Typical Bayesian models also predict less influence of the signal when the agent has high confidence on the prior belief. In Column 4 of Table 1, I interact both the prior and the signal with self-reported confidence on the prior. I control for the confidence, and the interaction of article pool fixed effect and $Dist_{pi}$ is also included. The result is consistent with the prediction, where the highest confidence (= 1) means almost zero coefficient on the signal ($0.195 - 0.160 \cdot 1 \approx 0$).

The results in this subsection give confidence in using a Bayesian model to understand the mechanism of the main empirical results of this paper.

4.2 Selective exposure to partisan news sources (Prediction [2])

In this subsection, I report the observed media selection patterns. The first step is to establish the existence of selective exposure—selection of media that are likely to represent the favored political group's position. The political group can be defined either by party affiliation, ethnicity, religion, regional boundaries, or other characteristics that are most relevant for media selection and political polarization in the particular context. In the context of this experiment, a partisan divide is the most relevant factor, and below I show the result that the readers choose media that are likely to represent the supported party's position. This suggests that the readers might associate partisan alignment with familiarity in the context of this experiment.

I first report the source selection of the group that could select based on source names (Source-Name Group). It was easy for them to identify and choose the media that are expected to represent the supported party's position if they wanted to do so, because the source names were shown on the selection screen (see Panel (a) of Figure 3). I estimate:

$$C_{ij} = \alpha_j + \alpha_r + \alpha_A + \gamma_1 \cdot y_{ijr}^{prior} + \gamma_2 \cdot P_{ij} + \mathbf{X_i}\beta + \varepsilon_{ijr} \tag{7}$$

where C_{ij} is a proxy for the chosen source's expected position on issue j, and P_{ij} is the position on issue j of the i's most trusted party at the baseline, which is calculated by taking the average of the baseline positions on issue j of other users who share the most trusted party with user i and have the baseline confidence ($\in [0,1]$) of 0.8 or larger. The user's level of trust for each party was surveyed at the baseline. Also included are issue FE, round FE, and user controls. The pool of articles changed over time—I add article pool fixed effect (α_A) to deal with this. Standard errors are clustered by user. $\gamma_2 = 0$ indicates the scenario where there is no selective exposure to the partisan media.

In Columns 1-3 of Table 2, the position of the article that was read by the user is used as a proxy for the source's expected position. The chosen source's expected position is strongly associated with the supported party's position, indicating that the news sources that are expected to represent the supported party's position were selected. The association is strong even after controlling for the prior belief (Column 2): Partisan news sources are selected rather than the news sources that are likely to confirm the preexisting views of the readers. In Columns 3 and 4 of Appendix Table 10, I show that even when the prior belief is relatively far from the party position, the users choose their partisan media.

The result is similar in Column 3 of Table 2, where an alternative measure of expected source position is used. "Expected position of the chosen source" is the average position on an issue of the articles written by the news sources that share the same favorite party as the news source at hand. The favorite party of the news source is the party that has the smallest average position distance across issues with the news source. In Columns 4 and 5, I take only the first-round selection after the grace period for each user×issue combination to minimize the concerns of sample selection due to dropouts. The results are similar. See Section 5 and Appendix Table 10 for subsample analysis of the Source-Name Group's media selection patterns.

The magnitude of selective exposure found here is comparable to the level found in Gentzkow and Shapiro (2011). I calculate the isolation index defined in the paper, deliberately coarsening the information that I have on users to match what the authors report in their paper. As described in Gentzkow and Shapiro (2011), the isolation index equals "the average conservative exposure of conservatives minus the average conservative exposure of liberals."[19] I find an isolation index of 0.076 to 0.143, depending on how I match the parties with the binary variable indicating ideology. If I categorize the two most conservative parties as "conservative" and the two remaining as "liberal," I get 0.076. If I instead categorize the Saenuri Party as the sole conservative party, I get 0.143. I do not directly ask about the user's ideology in the survey. The standard error, calculated from a non-parametric bootstrapping over users, is 0.06. Although the measure is noisier than Gentzkow and Shapiro (2011) due to the smaller sample size, the magnitude is comparable to the

[19] See Appendix A of Gentzkow and Shapiro (2011) for the precise method used by the authors. I closely follow their method, including leave-one-out estimation to adjust for the small sample. The only notable adjustment that I make to their method is that I put higher weight on the conservative users, so that the measure is comparable to the nationally representative sample. Specifically, I normalize so that the conservative sample as a whole receives the same weight as the liberal sample as a whole.

isolation index they find for the Internet (0.075) or national newspapers (0.104). It is encouraging to observe a similar level of selective exposure in the context of this experiment—the effect of selective exposure found later in this paper may be applicable to other contexts, such as online news consumption in the US.

The Source-Name Group—comprising those expected to show selective exposure—does indeed choose like-minded partisan media, as shown above. Table 3 compares the selection patterns of the three treatment groups. Columns 1 and 2 in Table 3 are benchmarks that correspond to the group that could select based on source names (Source-Name Group), a repetition of Table 2. Columns 3 and 4 in Table 3 show that selection of partisan media is effectively shut for the group that could select only based on source positions (Source-Position Group), presumably because it is hard to identify the media's party affiliations without seeing their names. Columns 5 and 6 present a placebo test. The observations that were given randomly selected articles should have no correlation between the positions of the chosen sources and the predetermined variables, and this is verified in the Table. Appendix Table 11 compares the media selection patterns of the finest subgroups. The selection patterns of subgroups within the three bigger treatment groups are similar.

4.3 Unregulated source selection facilitates learning and reduces extremism (Prediction [3])

In the last subsection, I show that the readers select news sources that are expected to represent their favored party's position (and not prior-confirming ones). I also observe that the three treatment groups showed quite different selection patterns. The group that could select based on source names (Source-Name Group) showed selective exposure, whereas the group that could select only based on source positions (Source-Position Group) revealed effective shutdown of selective exposure.[20] The group that was given randomly selected articles (No-Choice Group) did not have any choice and thus mechanically showed no selective exposure.

In this subsection, I present the main findings of this paper—the impact of media selection on evolution of beliefs. I compare the beliefs of three treatment groups that show different patterns of selective exposure as described above. Although I show the results of all three treatment groups for completeness, I focus on the comparison between the group that could select based

[20]Merging this group with either one of the remaining groups is unnatural for different reasons. The Source-Position Group shows no sign of partisan selective exposure, which is different from the Source-Name Group. On the other hand, the Source-Position Group is inherently different from the No-Choice Group as the group could select the source, albeit with limited information.

on source names (Source-Name Group) and the group that was given randomly selected articles (No-Choice Group). The comparison between the two groups identifies the effect of unregulated media selection—a setting that resembles the real world—(Source-Name Group) as opposed to random and balanced media exposure (No-Choice Group). Comparison between the group that could selective based only on source positions (Source-Position Group) and other groups is not highlighted because the results are generally noisy, as the group has the fewest subjects.

The model in Section 2 demonstrates that selective exposure enables the readers to select sources that they perceive to have high information value. The readers who were allowed to select familiar news sources will experience facilitated learning—the beliefs of the readers will change to a greater degree after reading articles (Prediction [3]).

In Table 4, I estimate:

$$\left|Belief_{post} - Belief_{pre}\right| = \gamma_0 + \gamma_1 \mathbf{I}\left(\text{SourceName}\right)_i + \gamma_2 \mathbf{I}\left(\text{SourcePosition}\right)_i + \mathbf{X_i}\beta + \alpha_j + \varepsilon_{ij} \quad (8)$$

where $\mathbf{I}(\cdot)$ is the indicator function, i is individual, and j is issue. $\mathbf{X_i}$ is the vector of user controls, and α_j is issue fixed effect. The level of observations is user$\times$issue, and the standard errors are clustered by user. The omitted group is the group that was given randomly selected articles (No-Choice Group).

I observe a 17% larger (significant at the 5% level) absolute distance between prior and posterior beliefs for the group that could select based on source names (Source-Name Group), compared with the group that was given randomly selected articles (No-Choice Group; the omitted category in the regression). In other words, unregulated source selection actually facilitates learning. The result is robust to the addition of user control variables (Column 2) and issue fixed effects (Column 3). This result is consistent with Prediction [3]. The result is not driven by any single issue—the effect is robust to dropping any one of the issues (Appendix Table 12).

If the larger distance between prior and posterior beliefs is driven by facilitated learning about the underlying true states of the world—as demonstrated by the model in Section 2.1—and if we are willing to believe that the ideal positions of the readers are probably not at the extreme points, we should observe larger movements in beliefs for those who have extreme baseline beliefs. Therefore, I next focus on the users who reported having extreme views at the baseline. A user is categorized as having an extreme opinion on an issue if he has an opinion of 0 or 1, which are the endpoints of the continuous horizontal scroll bar that the readers used to record their opinions in the app (See Panel (d) of Figure 3).

In Table 5, I investigate whether the proportion of people holding extreme opinions is affected by selective exposure. I run two different types of regressions: first difference (FD) and lagged dependent variable (LDV). Randomized controlled trials with good balance should, in principle, give similar results even though they make different kinds of assumptions on data generation process. Because there are concerns about the baseline balance for the group that could select based only on source positions (Source-Position Group) as can be seen in Column 4 of Table 5, I ran both regressions, as recommended by Angrist and Pischke (2008) and Wooldridge (2010). The regression equations are as follows.

$$y_{ij,post} - y_{ij,pre} = \gamma_0^{FD} + \gamma_1^{FD}\mathbf{I}(SourceName)_i + \gamma_2^{FD}\mathbf{I}(SourcePosition)_i + \mathbf{X_i}\beta + \alpha_j + \varepsilon_{ij} \tag{9}$$

$$y_{ij,post} = \gamma_0^{LDV} + \gamma_1^{LDV}\mathbf{I}(SourceName)_i + \gamma_2^{LDV}\mathbf{I}(SourcePosition)_i + \gamma_3^{LDV}y_{ij,pre} + \mathbf{X_i}\beta + \alpha_j + \varepsilon_{ij} \tag{10}$$

where $\mathbf{X_i}$ is the vector of user controls, and α_j is issue fixed effect. "*pre*" indicates baseline, and "*post*" is after reading an article in the experimental period. The outcome variable (y_{ij}) in Table 5 is an indicator variable that equals 1 if the belief of i on issue j is extreme (i.e., $Belief \in \{0,1\}$). The hypothesis that we want to test is whether free choice of media (Source-Name Group) facilitates a reduction of extremism compared to providing randomly selected articles (No-Choice Group). In terms of regression coefficients, we want to test whether $\gamma_1^{FD} = 0$ and $\gamma_1^{LDV} = 0$.

Regardless of the treatment, extremism is alleviated on average after reading an article—there was an 8 percentage point decline in the proportion of extreme opinions (0.368 at the baseline, and 0.284 after the first round of reading in the experimental period). Columns 1-3 show a robust negative effect of unregulated media selection on extremism—the group that could select based on source names (Source-Name Group) has 7 percentage points fewer people with extreme beliefs compared with the group that was given randomly selected articles (No-Choice Group; the omitted category in the regressions). The result for the group that could select only based on source positions (Source-Position Group) is noisy and inconsistent.

The results are not driven by a particular issue—they are robust to dropping any one issue (Appendix Table 13). Also, as shown in Appendix Table 14, similar results are found when I instead use a continuous measure of extremism: $(Belief - 0.5)^2$.

The results also can be seen graphically in the histograms in Figure 5 in which the prior and posterior belief distributions are contrasted for each group.

The results in this subsection indicate that selective exposure to partisan media may play a role in making the opinions less polarized. The readers learn more when they are allowed to choose the news sources, and their opinions get more moderated.

5 Mechanisms

In the Results Section above I find that unregulated media selection contributes to moderated policy views, an alleviation of polarization. The model in Section 2 explains a potential mechanism of this result. Readers select familiar media with biases about which they have a more sophisticated understanding. This facilitates learning.

In this section, I provide evidence that supports this mechanism. I first show that the belief updating patterns are consistent with prediction [4] of the model in Section 2—the policy view is affected more by the article if the article is written by a news source that is ideologically similar to the user, which I use as a proxy for familiarity. Second, I examine the heterogeneity of the treatment effect in political knowledge score. Third, I provide a series of suggestive evidence to demonstrate that the moderation of extreme policy views is not driven by users' changing their opinions toward the views of the like-minded partisan media. Finally, I examine whether learning is facilitated through increased attention paid by the users.

More Familiar News Source: Articles' Influence is Greater. The model in Section 2.1 predicts that a more familiar news source has greater influence in the readers' beliefs. According to the model, readers perceive that the articles from more familiar news sources have higher information value. In Bayesian language, readers put more weight on articles produced by such media (Proposition 2).

To test this, I again use the sample that was given randomly selected articles as in Subsection 4.1. I use an empirical specification similar to Equation (6), except now the prior and the article position are interacted with a proxy measure of familiarity—$Dist_{pi}$.

$$\begin{aligned} Belief_{ijr}^{post} = \alpha_j + \alpha_r + \alpha_A + \gamma_1 \cdot Belief_{ijr}^{prior} + \gamma_2 \cdot Belief_{ijr}^{prior} \cdot Dist_{pi} + \\ \gamma_3 \cdot s_{ijr} + \gamma_4 \cdot s_{ijr} \cdot Dist_{pi} + \gamma_5 \cdot Dist_{pi} + \alpha_A \cdot Dist_{pi} + \mathbf{X_i}\beta + \varepsilon_{ijr} \end{aligned} \tag{11}$$

where the ideological distance $Dist_{pi}$ is the average belief distance—across the app's eight issues—between the news source and the user at the baseline. This is a proxy measure of how unfamiliar the news source is to i. As before, $Belief_{ijr}^{post}$ is the posterior belief on j reported by i after article-reading round r, $Belief_{ijr}^{prior}$ is the belief reported prior to reading the article, and s_{ijr} is the position of the article (signal). In addition to the usual set of fixed effects, the interaction of article pool fixed effect with $Dist_{pi}$ is also included ($\alpha_A \cdot Dist_{pi}$).

Proposition 2 predicts that i is affected less by the signal when the news source is ideologically farther from the reader's belief on average (i.e., γ_4 is negative). The results are presented in Column 2 of Table 6 and are consistent with the model's prediction. When the source's and the user's ideologies are perfectly aligned, the article's coefficient is comparable to 40% of the coefficient put on the prior belief. However, when the ideological distance is in its 95 percentile (≈ 0.54), the article has virtually no influence ($0.244 - 0.54 \cdot 0.440 \approx 0$).

There is suggestive evidence that $Dist_{pi}$ is a good proxy for unfamiliarity. In Appendix Table 15, I show that the users tend to more accurately evaluate the article's position when the news source is ideologically closer to the user (low $Dist_{pi}$). This statistically significant association is true even after controlling for the prior belief's distance to the article position (Column 2). Since higher accuracy in article position evaluations is arguably associated with greater familiarity with media biases, I interpret this result as evidence supporting the use of $Dist_{pi}$ as a proxy for unfamiliarity.

$Dist_{pi}$ is still not a perfect proxy, and therefore one may formulate other explanations for the heterogeneity by devising a variable that is correlated with $Dist_{pi}$. One prominent alternative is the prior belief's distance to the article's position, $\left|s_{ijr} - Belief_{ijr}^{prior}\right|$. This is correlated with $Dist_{pi}$, and there is a reasonable argument for such heterogeneity: When the article is too far from the prior position, the reader may discount the quality of the article, taking the article less seriously (Gerber and Green, 1999; Gentzkow and Shapiro, 2006). This also may be the reason why the news sources that are likely to produce articles with low $\left|s_{ijr} - Belief_{ijr}^{prior}\right|$ can be selected. In this scenario, the group that was allowed to select based on source names (Source-Name Group) may learn more because they receive articles that they take more seriously.[21]

However, this hypothesis has no empirical support. First, as shown in Column 3 of Table 6, the interaction of s_{ijr} and $Dist_{pi}$ is still significant after controlling for the interaction of s_{ijr} and $\left|s_{ijr} - Belief_{ijr}^{prior}\right|$. Moreover, as shown in Column 1 of Table 7, the articles read by the group that was allowed to select based on source names (Source-Name Group) are not any closer to their prior views than articles read by the group that was given randomly selected articles (No-Choice Group). This is consistent to the findings of Subsection 4.2—the readers do not select prior-confirming news sources (Tables 2 and 3). In fact, the Source-Name Group selects partisan media even if they are likely to show articles that contradict their preexisting views (Columns 3 and 4 in Appendix Table 10). Combining the evidence, the mechanism is not likely to be

[21] On the other hand, note that in this scenario the users discount the value of information from articles that are far from their prior beliefs, which prohibits large changes in their opinions.

the selection of the prior-confirming media and the greater influence of articles written by such media.

Heterogeneity in Political Knowledge As discussed in Subsection 2.2, unregulated media selection leads to reduced polarization only when the readers have approximately correct assessment of media biases. The theory therefore predicts that the effect of media selection will be more pronounced among the subsample with sufficient sophistication. On the other hand, if a reader is highly sophisticated, it is expected that she is already near her ideal points, limiting the scope of the treatment. Given this tradeoff, we would expect the association between the treatment effect and the political sophistication to be inverse U-shaped.

I proxy the sophistication of the users with the "political knowledge score," a baseline measure of political knowledge. At the baseline, I ask the users to associate five political words[22] with one of the two major parties in South Korea. The "political knowledge score" is the number of correct associations.

In Table 8, I run the similar regressions in Equations (8) and (10), but this time interacting the treatment dummy variables with the political knowledge score. As before, I focus on the comparison between Source-Name Group and the No-Choice Group—the first six rows—where I have a sufficient number of users in each subsample of political knowledge score groups. As expected, the effect is driven by the users with middle-level scores (score $\in \{2,3\}$) for both columns 1 and 2.

Blind Followers? An alternative mechanism explaining the moderated policy views is the theory of "blind followers": The readers may take the articles at face value and update their beliefs toward the view of the articles without taking potential biases into account. Under this theory, if the group that was allowed to select based on source names (Source-Name Group) read less extreme articles on average, I would have found the same result—moderated policy views. This claim is not empirically supported. As shown in Columns 2 and 3 of Table 7, the articles read by the users who could select based on source names (Source-Name Group) are not any more extreme than the ones read by the users who were given randomly selected articles (No-Choice Group). In Column 2, I test whether there is any imbalance in the level of *relative* extremism, which equals $Sign(Belief_{pre} - 0.5) \times (Article\ Position - Belief_{pre})$. This measure is positive if the article is pooling the user's belief toward the extreme position of the user's current side, and it is negative if the article is pooling the user's belief toward the other side. In Column

[22] (i) Small government, (ii) Universal welfare, (iii) Neo-liberalism, (iv) Trickle-down effect, (v) *Eobeoi Yeonhap* (a famous neo-liberal organization of South Korea)

3, I test the imbalance in the level of *absolute* extremism, which equals $|Article\ Position - 0.5|$.

A more sophisticated "blind follower" may only follow the views of their partisan news sources. Suppose the partisan biases added to the signals are relatively moderate. Then—even if people are "blindly following" their partisan media's views—moderation of policy views can be achieved, consistent with the main results of this paper. This does not undermine the fact that media selection results in moderated policy views, a clear alleviation of polarization. However, in this scenario unregulated media selection, as opposed to random provision of articles, also implies a convergence to the partisan ideology, as demonstrated in Proposition 5 (Prediction [5]). After reading an article, people who share a favored political group will have more similar policy views, which can be interpreted as an exacerbation of polarization (Duclos et al., 2004). If this is the case, the overall impact of media selection on polarization is ambiguous.

Inconsistent with this hypothesis, users' beliefs are not differentially approaching the supported party's positions. In Table 9, I test this using the regressions in Equation (9) and Equation (10), except now the outcome variable is the distance between the user's policy view and the supported party's position. As shown in Table 9, there is no statistically distinguishable difference between the group that was given randomly selected articles (No-Choice Group) and the group that could select based on source names (Source-Name Group). I also try other definitions of party's positions, including the average baseline position of (1) the users who share the same most-trusted party and have trust in that party of over 0.5 ($\approx$ median of the trust in the most-trusted party) and (2) all the users who share the same most trusted party at the baseline. As shown in Appendix Table 17, there is no evidence of differential convergence regardless of the measure. I conclude that I find no evidence of increased polarization induced by selective exposure.[23]

Attention. As shown above, news articles have a greater influence on the reader's opinion if they are from media that are ideologically closer to the reader's beliefs. How does the "high influence" manifest to actual reading behaviors? For example, do people pay more attention to those articles and spend more time reading them?

In Table 10, I compare the number of minutes that treatment groups spent on reading articles as a proxy measure of the attention users paid to the articles. There are several caveats to mention about this proxy measure because it is not a pure measure of the amount of time readers spent

[23] Note, however, that the result in this subsection is not the rejection of an incorrect assessment of media bias. For example, the readers may start with extremely biased positions, very close to the party ideology. In this case, even without any bias reduction, news consumption from partisan media can take readers farther away from the party ideology and closer to the truth, as long as the signal retains some truth.

on articles. First, it includes time spent on selecting news sources (Source-Name and Source-Position Group; see Panel (a) and (b) of Figure 3) and on the news-source information screen (No-Choice Group; see Panel (c)). If we believe selection takes more time than simply getting information about the randomly selected source, users who were allowed to choose their news sources will have mechanically higher numbers for this measure. Second, users often temporarily move away from the app, but the time spent away from the app is still included in this measure. If we belief being away from the app is a form of inattention, we should interpret a recorded number that is too high as less attention—contradictory to its natural interpretation that more time means more attention. Third, I only measure the quantity, not the quality, of time spent on the app. It is possible that greater attention translates into less time spent to understand an article.

Keeping these caveats in mind, Table 10 shows that the group that could select based on source names (Source–Name Group) spent a median of 20 more seconds on the app or 12% more time when we consider only those who spent less than 10 total minutes to complete the reading process—those who probably stayed in the app the whole time—(Columns 2 and 3). The Source-Name Group is also more likely to abandon reading the article, but the difference is not statistically significant (Columns 4 and 5). All in all, the results are mixed.

6 Conclusion

This paper theoretically and empirically demonstrates that selective exposure can contribute to an alleviation of partisan polarization. By selecting partisan media, news consumers are exposed to familiar news sources. This facilitates learning and reduces polarization. The theory of this paper predicts that the reduction is more likely when (i) the ideal positions of the population is sufficiently close and (ii) the news consumers are sufficiently sophisticated to the extent that they take into account the media biases when updating beliefs.

Over the last 20 years, we have experienced increases in the scope of media selection, mainly due to the introduction of the Internet and the success of many social network websites. At the same time, polarization increased. The insight in this paper is not necessarily in contrast to these recent trends and might be a key to understanding them. The rapid increase in the sheer number of potential news sources may have eroded media loyalty and reduced media familiarity. The fame of news sources comes and goes quickly nowadays, to the extent that even famous news aggregators often fail to keep up with the popular demands. According to the mechanism introduced in this paper, reduced media familiarity can result in less learning and exacerbated polarization.

The result in this paper provides timely policy implications. Large social network websites have immense influence on how news articles are curated to their users: Their algorithms determine how news links are sorted and filtered on the users' news feeds. These companies are actively devising ways to regulate how the news articles are shared, in reaction to the recent emergence of fake news (e.g., see CyberScoop, 2017). At the same time, there is a society-wide discussion on how to regulate these intermediaries (e.g., see Fortune, 2017). According to the mechanism described above, in the short run, it may be beneficial to provide news articles from the users' familiar news sources instead of trying to provide balanced articles from relatively unfamiliar media. In the long run, an effort build media familiar can pay off.

How do we reconcile the mechanism introduced in this paper with the traditional wisdom that selective exposure drives partisan polarization (Sunstein, 2001)? Although selecting news sources among the established media is a very important decision, news consumers can make other choices. Readers may choose between the traditional, established media and other forms of new media, including pundits on Twitter or Facebook, fake news websites, and online communities. The selection can be across policy issues—readers can choose to stop paying attention to certain issues (Golman et al., 2017) or to political issues altogether (Kim and Kim, 2017). These "extensive-margin" selections may be the main drivers of political polarization. It is beyond the scope of this paper to compare the consequences of these different types of selections, and is an area I hope to address in future research.

It is important to note that the net social cost of partisan polarization may not necessarily be positive. In a model where the party and the reader have a set of shared underlying values that determine the "correct position" for each political issue, taking the arguments of the partisan media at face value could be welfare enhancing, especially when the information acquisition and the verification of the argument is costly to the general public (as in Canes-Wrone et al. 2001, Maskin and Tirole 2004). On the other hand, if these traits of partisan readers are strategically exploited by elites, the readers could be manipulated. Coordinated responses by non-partisan voters are required to counteract such under-informed partisans (Feddersen and Pesendorfer, 1996), which is a difficult goal to achieve (Battaglini, 2017). It is plausible that the misalignment in policy positions on *most of the issues* may cause deepened conflict and misunderstanding across the aisle, potentially even affecting personal matters such as marriages or hiring decisions (Iyengar and Westwood, 2015).

Until further studies are conducted in different contexts, it is an open question whether the results in this study will generalize. The model presented in Section 2 clarifies the conditions under which the channel can operate. Note that South Korea does not have a particularly harmonizing media environment. In fact, survey evidence shows that the level of partisan polarization is similar to or more severe than that of the US (Lupu, 2015). Also, as shown in Section 4.2, the level of selective exposure found in this study is similar to the level in the US, giving hope for wider applicability of the results. The theoretical insight can certainly be generalized to other contexts where a partisan divide is not the defining characteristic of polarization. For example, the idea can be applied to situations where there is a great divide in policy views between ethnic or religious groups.

References

Acemoglu, Daron, Chernozhukov, Victor, and Yildiz, Muhamet. Fragility of asymptotic agreement under bayesian learning. *Theoretical Economics*, 11(1):187–225, 2016.

Adena, Maja, Enikolopov, Ruben, Petrova, Maria, Santarosa, Veronica, and Zhuravskaya, Ekaterina. Radio and the rise of the nazis in prewar germany. *The Quarterly Journal of Economics*, 130(4):1885–1939, 2015.

Angrist, Joshua D and Pischke, Jörn-Steffen. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press, 2008.

Baldassarri, Delia and Gelman, Andrew. Partisans without constraint: Political polarization and trends in american public opinion 1. *American Journal of Sociology*, 114(2):408–446, 2008.

Banerjee, Abhijit and Somanathan, Rohini. A simple model of voice. *Quarterly Journal of Economics*, pages 189–227, 2001.

Battaglini, Marco. Public protests and policy making. *The Quarterly Journal of Economics*, 132 (1):485–549, 2017.

Boxell, Levi, Gentzkow, Matthew, and Shapiro, Jesse M. Is the internet causing political polarization? evidence from demographics. Working paper 23258, National Bureau of Economic Research, 2017.

Calvert, Randall L. The value of biased information: A rational choice model of political advice. *The Journal of Politics*, 47(02):530–555, 1985.

Campante, Filipe R and Hojman, Daniel A. Media and polarization: Evidence from the introduction of broadcast tv in the united states. *Journal of Public Economics*, 100:79–92, 2013.

Canes-Wrone, Brandice, Herron, Michael C, and Shotts, Kenneth W. Leadership and pandering: A theory of executive policymaking. *American Journal of Political Science*, pages 532–550, 2001.

Chen, Yan and Konstan, Joseph. Online field experiments: a selective survey of methods. *Journal of the Economic Science Association*, 1(1):29–42, 2015.

Chiang, Chun-Fang and Knight, Brian. Media bias and influence: Evidence from newspaper endorsements. *The Review of Economic Studies*, page rdq037, 2011.

CyberScoop. Facebook admits 'false news' is a problem, pledges to curb 'information operations', April 2017. URL https://www.cyberscoop.com/facebook-admits-fake-news-problem-pledges-curb-information-operations/. [Written by Chris Bing Online; posted 27-April-2017].

DellaVigna, Stefano and Kaplan, Ethan. The fox news effect: Media bias and voting. *The Quarterly Journal of Economics*, pages 1187–1234, 2007.

Druckman, James N, Fein, Jordan, and Leeper, Thomas J. A source of bias in public opinion stability. *American Political Science Review*, 106(02):430–454, 2012.

Duclos, Jean-Yves, Esteban, Joan, and Ray, Debraj. Polarization: concepts, measurement, estimation. *Econometrica*, 72(6):1737–1772, 2004.

Durante, Ruben, Pinotti, Paolo, and Tesei, Andrea. The political legacy of entertainment TV. DP 10738, CEPR Discussion Papers, 2015.

Feddersen, Timothy J and Pesendorfer, Wolfgang. The swing voter's curse. *The American economic review*, pages 408–424, 1996.

Fortune. Facebook and google need to be regulated, says british news industry, March 2017. URL http://fortune.com/2017/03/09/facebook-google-regulated/. [Written by Mathew Ingram Online; posted 09-March-2017].

Gentzkow, Matthew. Polarization in 2016. *Mimeo*, 2016.

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

Gentzkow, Matthew and Shapiro, Jesse M. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71, 2010.

Gentzkow, Matthew and Shapiro, Jesse M. Ideological segregation online and offline. *The Quarterly Journal of Economics*, 126:1799–1839, 2011.

Gentzkow, Matthew, Shapiro, Jesse M, and Taddy, Matt. Measuring polarization in high-dimensional data: Method and application to congressional speech. Working paper 22423, National Bureau of Economic Research, 2016.

Gerber, Alan and Green, Donald. Misperceptions about perceptual bias. *Annual review of political science*, 2(1):189–210, 1999.

Golman, Russell, Hagmann, David, and Loewenstein, George. Information avoidance. *Journal of Economic Literature*, 55(1):96–135, 2017.

Hetherington, Marc J. Resurgent mass partisanship: The role of elite polarization. *American Political Science Review*, 95(3):619–631, 2001.

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

Iyengar, Shanto and Westwood, Sean J. Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science*, 59(3):690–707, 2015.

Iyengar, Shanto, Sood, Gaurav, and Lelkes, Yphtach. Affect, not ideologya social identity perspective on polarization. *Public opinion quarterly*, 76(3):405–431, 2012.

Jensen, Jacob, Naidu, Suresh, Kaplan, Ethan, Wilse-Samson, Laurence, Gergen, David, Zuckerman, Michael, and Spirling, Arthur. Political polarization and the dynamics of political language: Evidence from 130 years of partisan speech [with comments and discussion]. *Brookings Papers on Economic Activity*, pages 1–81, 2012.

Kim, Eunji and Kim, Jin Woo. Switching on and off: Rethinking partisan selective exposure. *Working Paper*, 2017.

Korea Institute of Public Administration. Korea social integration survey, 2015.

Korea Press Foundation. Media audience survey, 2016.

Leeper, Thomas J and Slothuus, Rune. Political parties, motivated reasoning, and public opinion formation. *Political Psychology*, 35(S1):129–156, 2014.

Liang, Annie, Mu, Xiaosheng, and Syrgkanis, Vasilis. Optimal learning from multiple information sources. 2017.

Lupu, Noam. Party polarization and mass partisanship: A comparative perspective. *Political Behavior*, 37(2):331–356, 2015.

Martin, Gregory J and Yurukoglu, Ali. Bias in cable news: Persuasion and polarization. 2016.

Maskin, Eric and Tirole, Jean. The politician and the judge: Accountability in government. *The American Economic Review*, 94(4):1034–1054, 2004.

McCarty, Nolan, Poole, Keith T, and Rosenthal, Howard. *Polarized America: The dance of ideology and unequal riches*. mit Press, 2016.

Mullainathan, Sendhil and Shleifer, Andrei. The market for news. *American Economic Review*, pages 1031–1053, 2005.

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

Pew Research Center. Political polarization in the american public. 2014.

Poole, Keith T and Rosenthal, Howard. Patterns of congressional voting. *American journal of political science*, pages 228–278, 1991.

Prior, Markus. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press, 2007.

Prior, Markus. The immensely inflated news audience: Assessing bias in self-reported news exposure. *Public Opinion Quarterly*, 73(1):130–143, 2009.

Schilbach, Frank. Alcohol and self-control: A field experiment in india. 2015.

Sethi, Rajiv and Yildiz, Muhamet. Communication with unknown perspectives. *Econometrica*, 84(6):2029–2069, 2016a.

Sethi, Rajiv and Yildiz, Muhamet. Culture and communication. 2016b.

Stroud, Natalie Jomini, Muddiman, Ashley, and Lee, Jae Kook. Seeing media as group members: An evaluation of partisan bias perceptions. *Journal of Communication*, 64(5):874–894, 2014.

Suen, Wing. The self-perpetuation of biased beliefs. *The Economic Journal*, 114(495):377–396, 2004.

Sunstein, Cass. *Republic. com*. Princeton, NJ: Princeton University Press, 2001.

Sunstein, Cass R. *# Republic: Divided Democracy in the Age of Social Media*. Princeton University Press, 2017.

Wooldridge, Jeffrey M. *Econometric analysis of cross section and panel data*. MIT press, 2010.

YouGov. Anglo-saxon attitudes: A survey of british and american views of the world. 2008.

Table 1. Belief updating patterns: consistent with typical Bayesian models

Subsample:	Random provision			
	(1)	(2)	(3)	(4)
	Belief $_{post}$			
Belief $_{pre}$	0.670***	0.660***	0.616***	0.259***
	(0.0145)	(0.0135)	(0.0151)	(0.0396)
Belief $_{pre}$ × Prior confidence				0.453***
				(0.0524)
Article position	0.0760***	0.0760***	0.0759***	0.195***
	(0.00954)	(0.00964)	(0.00972)	(0.0450)
Article position × Prior confidence				-0.160***
				(0.0530)
Prior Confidence				0.117
				(0.163)
Constant	0.0523***	0.0718***	0.133***	-0.0499
	(0.00628)	(0.00590)	(0.0140)	(0.144)
Round FE	N	Y	Y	Y
User controls	N	N	Y	Y
Article Pool FE (includes Issue FE)	Y	Y	Y	Y
Article Pool FE × Prior Confidence				Y
# of users	1,420			
Observations (user×issue×round)	7,794			

Note: Each column in this table originates from a separate OLS regression of the belief after reading an article on respective regressors. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. Included are the observations where users were provided randomly selected articles: that is, all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. The article pool from which an article is randomly drawn changes over time—article pool FE is added to the regression. "Article position" is the average position of the article reported by all the users after reading it. "Prior Confidence" ($\in [0,1]$) is the reported confidence in the prior belief. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 2. Readers select their partisan media (and not prior-confirming ones)

Subsample:	Source-Name Group (source selection allowed; with source names shown)			Source-Name Group, 1st round only	
	(1)	(2)	(3)	(4)	(5)
	Position of the article read	Position of the article read	Expected position of the chosen source	Position of the article read	Expected position of the chosen source
Prior belief		0.0103	0.00520	0.0183	0.00889
		(0.0279)	(0.0220)	(0.0523)	(0.0414)
Party position of the reader's most-trusted party on the issue	0.282***	0.273***	0.181***	0.298**	0.218**
	(0.0634)	(0.0650)	(0.0466)	(0.121)	(0.0882)
Constant	0.366***	0.363***	0.422***	0.429***	0.496***
	(0.0330)	(0.0324)	(0.0277)	(0.0543)	(0.0547)
Article Pool FE (includes Issue FE)	Y	Y	Y	Y	Y
Round FE	Y	Y	Y	N/A	N/A
User Controls	Y	Y	Y	Y	Y
# of users	84	84	84	84	84
Observations (user×issue×round)	1,659	1,489	1,489	438	438

Note: Each column in this table originates from a separate OLS regression of position of the chosen article on the latest position users held before making the choice (prior) and the party position of the reader's most-trusted party on the issue. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added when applicable. "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "Expected source position" is the average position on an issue of the articles written by the sources that share the same favorite party as the source at hand. The favorite party of the source is determined by the party that has the smallest squared position distance across issues with the source. The article pool from which an article is randomly drawn changes over time—article pool FE is added to the regression. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 3. Source selection: comparison between treatment groups

Subsample:	*Source-Name Group*: identifying partisan media is *easy* (selection allowed; with source names shown)		*Source-Position Group*: identifying partisan media is *hard*, identifying prior-confirming media is *easy* (selection allowed; with only source positions shown)		Placebo: articles are randomly selected	
	(1)	(2)	(3)	(4)	(5)	(6)
	Position of the article read	Expected position of the chosen source	Position of the article read	Expected position of the chosen source	Position of the article read	Expected position of the chosen source
Prior belief	0.0103	0.00520	-0.0461	-0.0528	-0.00311	-0.00226
	(0.0279)	(0.0220)	(0.0471)	(0.0431)	(0.00918)	(0.00734)
Party position of the reader's most-trusted party on the issue	0.273***	0.181***	-0.00159	-0.0194	0.0218	0.0164
	(0.0650)	(0.0466)	(0.127)	(0.0752)	(0.0225)	(0.0163)
Constant	0.363***	0.422***	0.470***	0.445***	0.454***	0.456***
	(0.0324)	(0.0277)	(0.0582)	(0.0453)	(0.0108)	(0.00804)
Pool, Round FE, User Controls	Y	Y	Y	Y	Y	Y
# of users	84	84	38	38	1305	1305
Observations	1,489	1,489	680	680	7,348	7,347

Note: Each column in this table originates from a separate OLS regression of position of the chosen article on the latest position users held before making the choice (prior) and the party position of the reader's most-trusted party on the issue. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added when applicable. "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "Expected source position" is the average position on an issue of the articles written by the sources that share the same favorite party as the source at hand. The favorite party of the source is determined by the party that has smallest squared position distance across issue with the source. The article pool from which an article is randomly drawn changes over time—article pool FE is added to the regression. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 4. More learning caused by selective exposure

Subsample:	User × issue combination w/ at least one article read after grace period		
	(1)	(2)	(3)
	$\lvert Belief_{post}$ - $Belief_{pre} \rvert$		
Source-Name Group	0.0233**	0.0220**	0.0221**
(selection allowed; with source names shown)	(0.0112)	(0.0103)	(0.0103)
Source-Position Group	-0.00871	-0.00365	-0.00207
(selection allowed; with only positions shown)	(0.0158)	(0.0160)	(0.0163)
Constant	0.135***	0.188***	0.138***
(Omitted: *No-Choice Group* – random provision of articles)	(0.00544)	(0.0193)	(0.0220)
p-value: *Source-Name* = *Source-Position*	0.0731	0.141	0.175
User Controls	N	Y	Y
Issue FE	N	N	Y
# of users	336	336	336
Observations (user×issue)	1,790	1,790	1,790

Note: Each column in this table originates from a separate OLS regression of the absolute value of change in belief on the treatment group dummies. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. $Belief_{pre}$ ([0,1]) is the reported belief at the baseline. $Belief_{post}$ ([0,1]) is the reported belief after reading an article in the experimental period (not grace period). Extreme opinion is the belief that equals to 0 or 1. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 5. Moderation of extreme beliefs

Model:	Basic	First Difference (FD)	Lagged Dependent Variable (LDV)	Baseline balance
	(1)	(2)	(3)	(4)
	$Extreme_{post}$	$Extreme_{post}$ - $Extreme_{pre}$	$Extreme_{post}$	$Extreme_{pre}$
Sample Mean	0.284	-0.0838	0.284	0.368
Source-Name Group	-0.0656*	-0.0740**	-0.0687**	0.00835
(selection allowed; with source names shown)	(0.0335)	(0.0353)	(0.0286)	(0.0387)
Source-Position Group	0.0325	-0.0832	-0.0102	0.116*
(selection allowed; with only positions shown)	(0.0585)	(0.0533)	(0.0467)	(0.0662)
$Extreme_{pre}$			0.370***	
			(0.0266)	
Constant	0.381***	0.00766	0.243***	0.373***
(Omitted: *No-Choice Group* – random provision of articles)	(0.0584)	(0.0588)	(0.0501)	(0.0662)
p-value: *Source-Name* = *Source-Position*	0.105	0.872	0.226	0.129
User controls, Issue FE	Y	Y	Y	Y
# of users	336	336	336	336
Observations (user×issue)	1,790	1,790	1,790	1,790

Note: Each column in this table originates from a separate OLS regression. $Extreme_{post}$ is a dummy variable indicating whether the belief ([0,1]) equals to 0 or 1 after reading an article in the experimental period (not grace period). $Extreme_{pre}$ is an analogous dummy variable for the baseline belief. “User Controls” include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 6. Article has larger influence if it is written by familiar media

Subsample:	Random provision		
	(1)	(2)	(3)
		Belief $_{post}$	
Belief $_{pre}$	0.616***	0.600***	0.582***
	(0.0151)	(0.0296)	(0.0323)
Belief $_{pre}$ × "Source's ideology distance to the reader" (95th Percentile: 0.54)		-0.00812	-0.0325
		(0.0933)	(0.116)
Belief $_{pre}$ × |Article position - Belief $_{pre}$|			0.0433
			(0.0813)
Article position	0.0759***	0.244***	0.266***
	(0.00972)	(0.0287)	(0.0300)
Article position × "Source's ideology distance to the reader"		-0.440***	-0.331***
		(0.0854)	(0.114)
Article position × |Article position - Belief $_{pre}$|			-0.0913
			(0.0674)
"Source's ideology distance to the reader"		0.0643	0.298***
		(0.0722)	(0.109)
|Article position - Belief $_{pre}$|			-0.205***
			(0.0793)
Constant	0.133***	0.0471	0.0390
	(0.0140)	(0.0362)	(0.0328)
Round FE, Article Pool FE, User controls	Y	Y	Y
Article Pool FE × Source's ideology distance		Y	Y
Article Pool FE × |Article position - Belief $_{pre}$|			Y
# of users	1420	1420	1420
Observations (user×issue×round)	7,794	7,794	7,794

Note: Each column in this table originates from a separate OLS regression of belief after reading an article on respective regressors. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. Included are the observations where users were provided randomly selected articles: that is, I include all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. Article pool from which an article is randomly drawn is changing over time – article pool FE is added to the regression. Article position is the average position of the article reported by all the users after reading it. "Source's ideology distance to the reader" ($\in [\mathbf{0}, \mathbf{1}]$) is the mean absolute distance between the baseline user belief and the source's average article position. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 7. Summary statistics of articles read

	(1)	(2)	(3)
	$\lvert$Article Position - Belief $_{pre}\rvert$	Article's relative extremism = sign(Belief $_{pre}$-0.5)× (Article Position - Belief $_{pre}$)	Article's absolute extremism = $\lvert$Article Position – 0.5$\rvert$
Sample Mean	0.316	-0.293	0.175
Source-Name Group	0.00255	-0.00379	-0.00154
(selection allowed; with source names shown)	(0.0141)	(0.0167)	(0.00569)
Source-Position Group	0.0299	-0.0375	0.0102
(selection allowed; with only positions shown)	(0.0200)	(0.0236)	(0.00714)
Constant	0.312***	-0.287***	0.175***
(Omitted: *No-Choice Group* – random provision of articles)	(0.00804)	(0.00952)	(0.00334)
p-value: *Source-Name* = *Source-Position*	0.207	0.188	0.135
User controls, Issue FE	Y	Y	Y
# of users	336	336	336
Observations (user×issue)	1,772	1,772	1,772

Note: Each column in this table originates from a separate OLS regression. Belief$_{pre}$ ([0,1]) is the reported belief at the baseline. Article position is the average position of the article reported by all the users after reading it. "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 8. Learning and moderation: driven by people with middle score

	(1)	(2)
	$\lvert \text{Belief}_{post} - \text{Belief}_{pre} \rvert$	Extreme_{post}
Sample Mean	0.140	0.284
Source-Name Group × Score = 0	-0.00457	-0.0227
	(0.0226)	(0.0632)
Source-Name Group × Score = 1	-0.0179	-0.0603
	(0.0317)	(0.0706)
Source-Name Group × Score = 2	0.0787***	-0.0772
	(0.0228)	(0.0669)
Source-Name Group × Score = 3	-0.00103	-0.127*
	(0.0211)	(0.0702)
Source-Name Group × Score = 4	0.0193	-0.0970
	(0.0268)	(0.0677)
Source-Name Group × Score = 5	-0.000865	-0.0525
	(0.0179)	(0.0612)
Source-Position Group × Score = 0	-0.0826***	0.255***
	(0.0208)	(0.0817)
Source-Position Group × Score = 1	0.0108	-0.256***
	(0.0463)	(0.0716)
Source-Position Group × Score = 2	-0.0138	-0.101
	(0.0377)	(0.0946)
Source-Position Group × Score = 3	-0.0110	-0.00135
	(0.0225)	(0.0823)
Source-Position Group × Score = 4	-0.000851	-0.0406
	(0.0362)	(0.0756)
Source-Position Group × Score = 5	0.0176	0.0475
	(0.0300)	(0.102)
User controls, Issue FE, Score FE, Constant	Y	Y
Control for: Extreme_{pre} × Score FE	N	Y
# of users	336	336
Observations (user×issue)	1,790	1,790

Note: Each column in this table originates from a separate OLS regression. “Belief_{pre}” ([0,1]) is the reported belief at the baseline. “Belief_{post}” ([0,1]) is the reported belief after reading an article in the experimental period (not the grace period). “Extreme_{post}” is a dummy variable indicating whether the belief ([0,1]) equals to 0 or 1 after reading an article in the experimental period (not grace period). “Extreme_{pre}” is an analogous dummy variable for the baseline belief. “Score” ([0,5]) is the political knowledge score, which is the number of correct answers to simple quiz conducted at the baseline. “User Controls” include baseline characteristics of the users: age, gender, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 9. No evidence on convergence to party's position

Model:	Basic	First Difference (FD)	Lagged Dependent Variable (LDV)	Baseline balance
	(1)	(2)	(3)	(4)
	$\lvert Belief_{post}$ - Party Position$\rvert$	$\lvert Belief_{post}$ - Party Position$\rvert$ - $\lvert Belief_{pre}$ - Party Position$\rvert$	$\lvert Belief_{post}$ - Party Position$\rvert$	$\lvert Belief_{pre}$ - Party Position$\rvert$
Sample Mean	0.166	-0.0153	0.166	0.181
Source-Name Group (selection allowed; with source names shown)	0.00132 (0.0115)	-0.0157 (0.00972)	-0.00487 (0.00956)	0.0170 (0.0108)
Source-Position Group (selection allowed; with only positions shown)	-0.00445 (0.0132)	-0.00109 (0.0124)	-0.00323 (0.0115)	-0.00337 (0.0123)
$\lvert Belief_{pre}$ - Party Position$\rvert$			0.364*** (0.0340)	
Constant (Omitted: *No-Choice Group* – random provision of articles)	0.154*** (0.0213)	0.0106 (0.0204)	0.102*** (0.0184)	0.143*** (0.0222)
p-value: *Source-Name* = *Source-Position*	0.713	0.288	0.903	0.141
User controls, Issue FE	Y	Y	Y	Y
# of users	316	316	316	316
Observations (user×issue)	1,682	1,682	1,682	1,682

Note: Each column in this table originates from a separate OLS regression. $Belief_{pre}$ ([0,1]) is the reported belief at the baseline. $Belief_{post}$ ([0,1]) is the reported belief after reading an article in the experimental period (not grace period). "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Table 10. Time spent on reading articles

Subsample:	Time spent < 10 minutes		Full sample		
Model:	Linear regression	Quantile regression (median)	Quantile regression (median)	Linear probability model	Linear probability model
	(1)	(2)	(3)	(4)	(5)
	Time spent (minutes)	Time spent (minutes)	Time spent (minutes)	Time spent > 10 minutes	Time spent > 1 hours
Sample Mean	2.394	2.394	1218	0.161	0.119
Sample Median	1.850	1.850	2.233	0	0
Source-Name Group	0.234	0.298*	0.332	0.0464	0.0225
(selection allowed; with source names shown)	(0.174)	(0.177)	(0.217)	(0.0297)	(0.0232)
Source-Position Group	-0.00313	0.146	0.0598	-0.0243	-0.0355
(selection allowed; with only positions shown)	(0.215)	(0.215)	(0.259)	(0.0311)	(0.0264)
Constant	3.705***	2.766***	3.545***	0.306***	0.244***
(Omitted: *No-Choice Group* – random provision of articles)	(0.344)	(0.336)	(0.447)	(0.0538)	(0.0454)
	0.344	0.519	0.374	0.0626	0.0575
User controls, Issue FE	Y	Y	Y	Y	Y
# of users	326	326	336	336	336
Observations (user×issue)	1,488	1,488	1,774	1,774	1,774

Note: Each column in this table originates from a separate regression. The model, sample, and dependent variable for each regression are indicated in the first few rows of the table. "Time spent" is minutes a person spent on reading an article; the measure also includes time spent on the selection screen (for Source-Position and Source-Name Group) or the source info screen (for No-Choice Group) and time spent on belief updating and article evaluation. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Figure 1. Issues covered in the mobile application

1. Foreign policy toward North Korea
2. Providing free lunch to every elementary school student
3. Single nationalized history textbook
4. Cash support to unemployed youth
5. Who should pay (central vs. regional government) for free early education for ages 3-5
6. Minimum wage
7. Law governing legislative process (minority protection vs. faster process)
8. Flexible labor market vs. job security

Figure 2. Randomization design

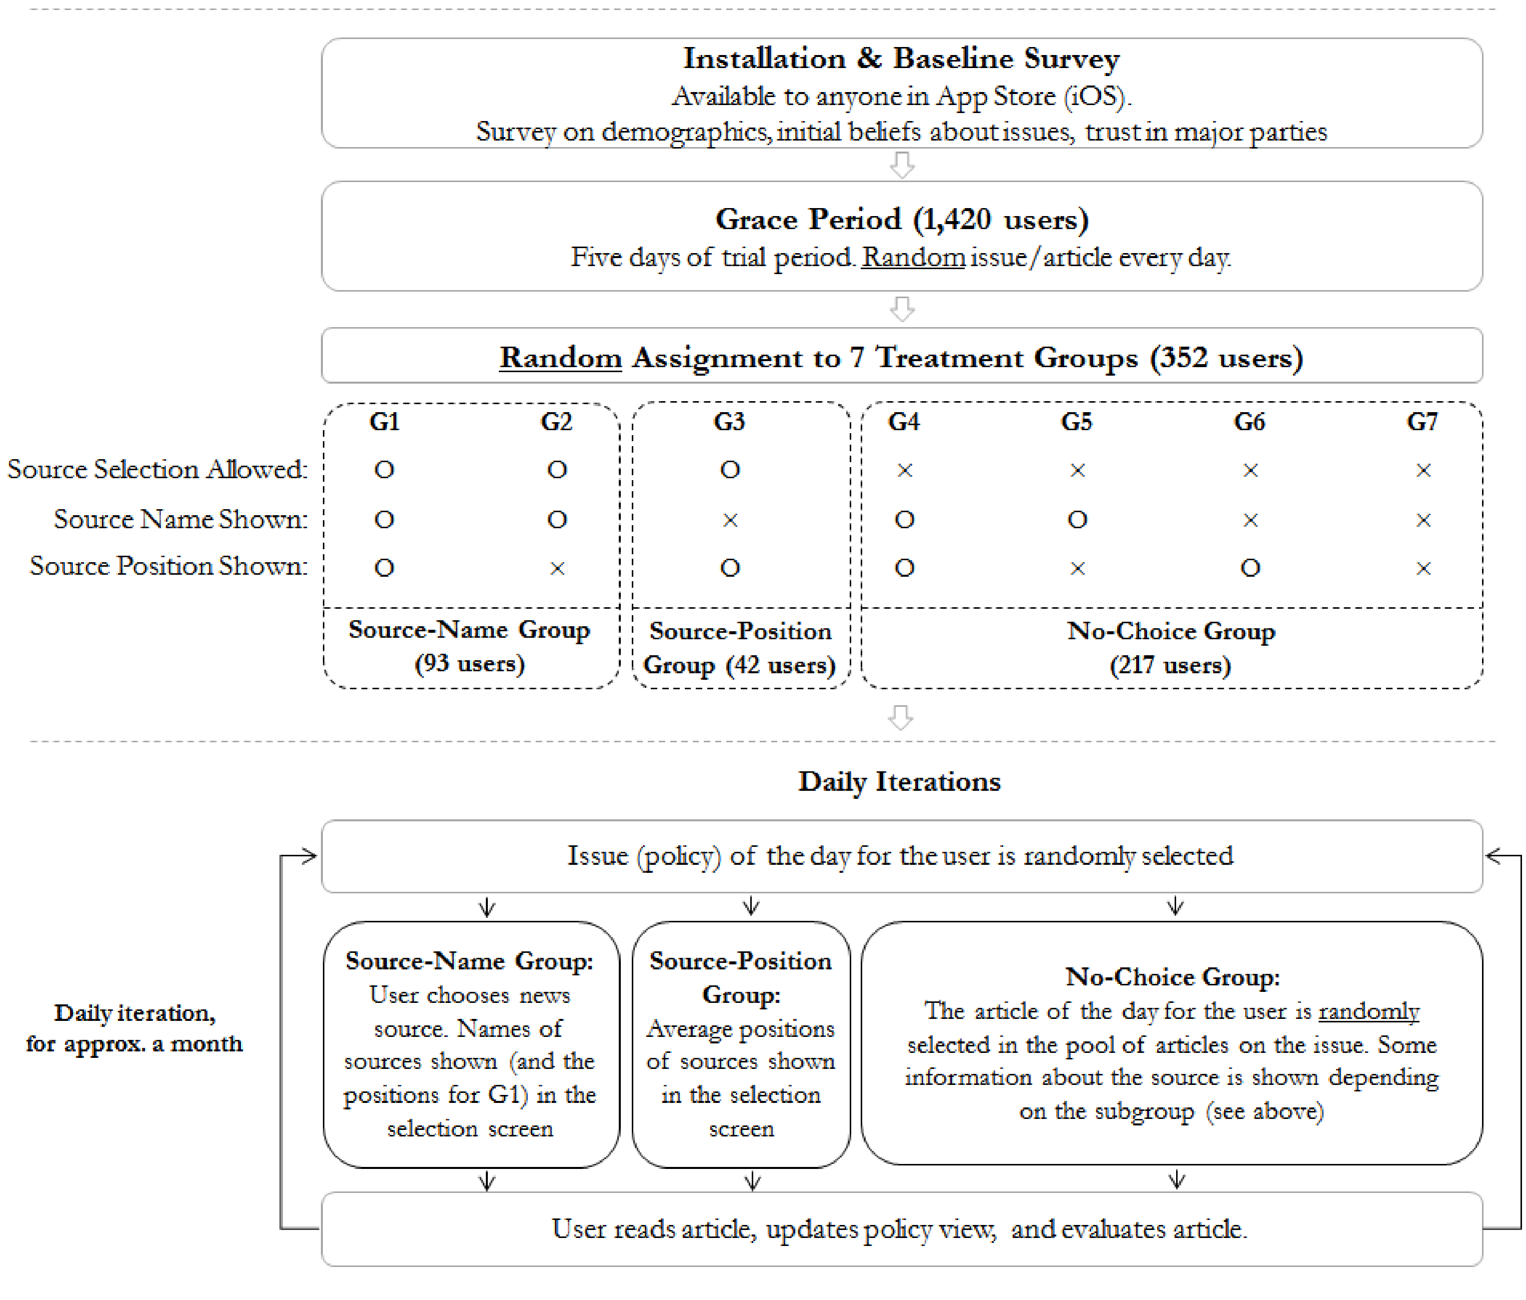

Figure 3. Screens of the app

(a) **Source-Name Group** chooses news source with names (circled) of the sources shown (and positions, if G1)

(b) **Source-Position Group** chooses source only based on source's positions (average position on the issue; circled); provider name is *not* shown

(c) **No-Choice Group** gets randomly selected article, but can also be informed about name and/or position of the source of the article

(d) After reading an article, user **updates** her **belief** on the issue, and her **confidence** in the position.

(e) The user also reports the **subjective quality** and **position** of the article.

Figure 4. Examples of cardinal description (foreign policy toward North Korea)

1. North Korea is our enemy; we should be aggressive in every front
2. Supporting North Korea can be considered, but only after nuclear weapons are destroyed
3. Both sides have their points; I am neutral
4. Communication and collaboration first, but aggression is also necessary
5. We should put all our effort on communication and collaboration

Figure 5. Histogram of beliefs, before and after article reading

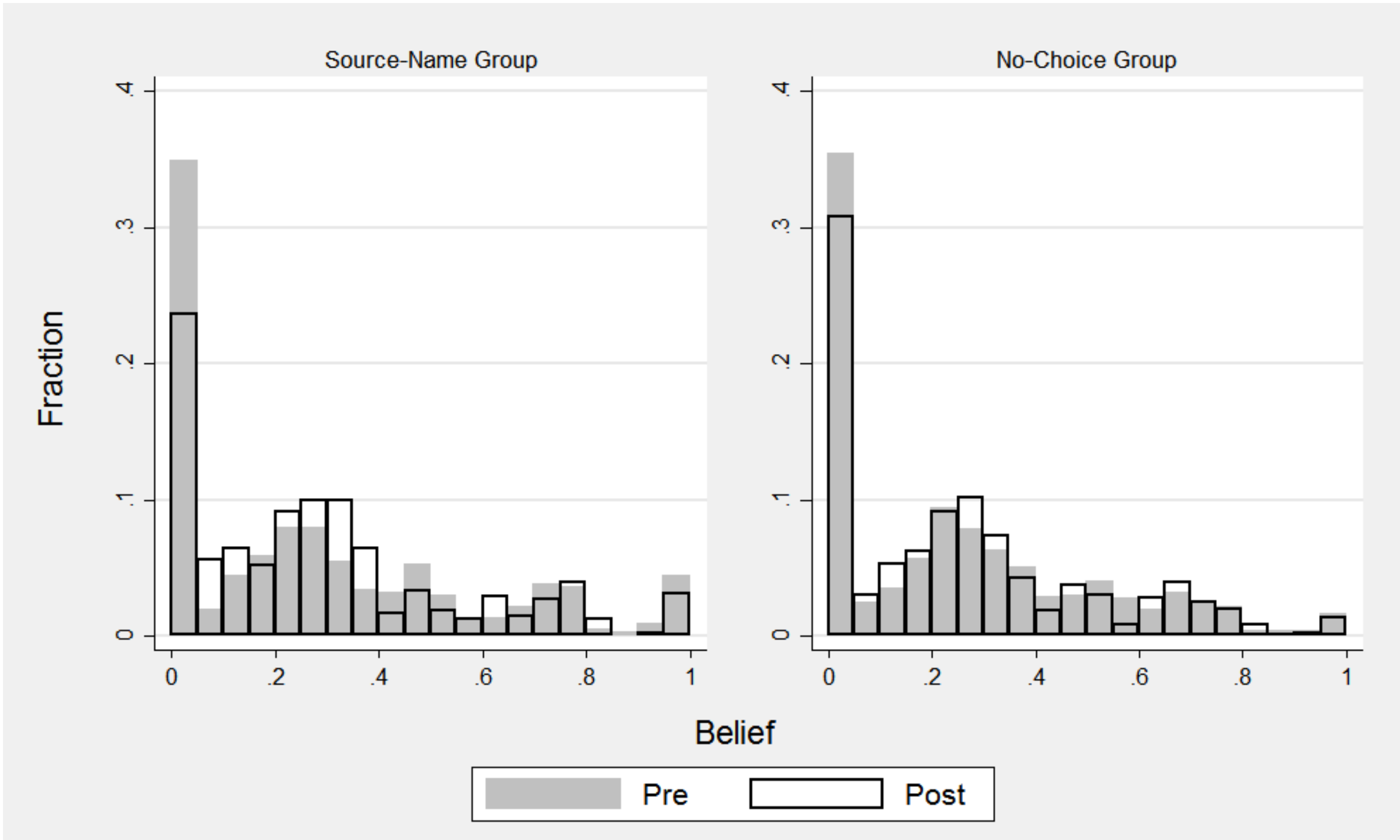

Note: "Pre" indicate the beliefs at the baseline, and "Post" is after reading an article in the experimental period. All the issues are pooled together.

Appendix

A Proofs

Derivation of Posterior Distribution

The density function of posterior belief after getting signal s_{pij} is proportional to the multiplication of the likelihood function, $l(s_{pij}|\theta_j)$, and the prior distribution, $f(\theta_j)$.

$$f(\theta_j|s_{pij}) \propto l(s_{pij}|\theta_j) \cdot f(\theta_j)$$

Since I_p is also unknown, we must add that component to get the relevant likelihood function.

$$\begin{aligned} f(\theta_j|s_{pij}) &\propto \int l(s_{pij}|\theta_j, I_{pj}) \cdot f(I_{pj}|\theta_j) dI_{pj} \times f(\theta_j) \\ &= \int l(s_{pij}|\theta_j, I_{pj}) \cdot f(I_{pj}) dI_{pj} \times f(\theta_j) \end{aligned}$$

where the final equality is due to the independence of the prior distribution of I_{pj} and θ_j. Plugging in the density functions and rearranging, we get:

$$\begin{aligned} f(\theta_j|s_{pij}) &\propto \int \exp\left(-\frac{\tau_s}{2}\left(s_{pij}-\theta_j-I_{pj}\right)^2 - \frac{\tau_{Ipji}}{2}\left(I_{pji}-I_{pj}\right)^2\right) dI_{pj} \times f(\theta_j) \\ &\propto \int \exp\left(-\frac{1}{2}\left[\left(\tau_s+\tau_{Ipji}\right)I_{pj}^2 - 2\left\{\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}\right\}I_{pj} + \tau_s\left(s_{pij}-\theta_j\right)^2\right]\right) dI_{pj} \times f(\theta_j) \\ &= \int \exp\left(-\frac{\tau_s+\tau_{Ipji}}{2}\left[I_{pj}^2 - 2\frac{\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}}{\tau_s+\tau_{Ipji}}I_{pj} + \left(\frac{\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}}{\tau_s+\tau_{Ipji}}\right)^2\right] - \frac{1}{2}\left[-\frac{\left(\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}\right)^2}{\tau_s+\tau_{Ipji}} + \tau_s\left(s_{pij}-\theta_j\right)^2\right]\right) dI_{pj} \\ &\quad \times f(\theta_j) \\ &\propto \int \sqrt{\frac{\tau_s+\tau_{Ipji}}{2}} \exp\left(-\frac{1}{2}\left(I_{pj} - \frac{\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}}{\tau_s+\tau_{Ipji}}\right)^2\right) dI_{pj} \times \exp\left(-\frac{1}{2}\left[-\frac{\left(\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}\right)^2}{\tau_s+\tau_{Ipji}} + \tau_s\left(s_{pij}-\theta_j\right)^2\right]\right) \times f(\theta_j) \\ &= \exp\left(-\frac{1}{2}\left[-\frac{\left(\tau_s\left(s_{pij}-\theta_j\right)+\tau_{Ipji}I_{pji}\right)^2}{\tau_s+\tau_{Ipji}} + \tau_s\left(s_{pij}-\theta_j\right)^2\right]\right) \times f(\theta_j) \end{aligned}$$

where the last equality is because the integrand is a density of a normal distribution. Plugging in

θ_j, rearranging, and dropping a term unrelated to θ_j, we get:

$$
\begin{aligned}
f(\theta_j|s_{pij}) &\propto \exp\left(-\frac{1}{2}\left[-\frac{\tau_s^2}{\tau_s+\tau_{Ipji}}\left(\theta_j - s_{pij}\right)^2 + \frac{2\tau_s\tau_{Ipji}}{\tau_s+\tau_{Ipji}} I_{pji}\left(\theta_j - s_{pij}\right) + \tau_s\left(\theta_j - s_{pij}\right)^2 + \tau_{ij0}\left(\theta_j - \theta_{ij0}\right)^2\right]\right) \\
&\propto \exp\left(-\frac{1}{2}\left[\frac{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}{\tau_s+\tau_{Ipji}}\theta^2 - 2\left\{\frac{\tau_s\tau_{Ipji}}{\tau_s+\tau_{Ipji}}\left(s_{pij} - I_{pji}\right) + \frac{\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}{\tau_s+\tau_{Ipji}}\theta_{ij0}\right\}\theta_j\right]\right) \\
&\propto \sqrt{\frac{\frac{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}{\tau_s+\tau_{Ipji}}}{2}} \exp\left(-\frac{1}{2}\left\{\theta_j - \frac{\tau_s\tau_{Ipji}}{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}\left(s_{pij} - I_{pji}\right) - \frac{\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}\theta_{ij0}\right\}^2\right)
\end{aligned}
$$

which is the density of normal distribution with mean of $\frac{\tau_s\tau_{Ipji}}{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}\left(s_{pij} - I_{pji}\right) - \frac{\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}\theta_{ij0}$ and precision of $\frac{\tau_s\tau_{Ipji}+\tau_s\tau_{ij0}+\tau_{Ipji}\tau_{ij0}}{\tau_s+\tau_{Ipji}}$.

Proof of Proposition 3

Proof. Minimizing the loss function is equivalent to minimizing the posterior variance.

$$
\begin{aligned}
\underset{p}{argmax} E_i\left[max_{a_{ij}} E_i\left[-(a_{ij} - \theta_j^*)^2 \mid s_{pij}\right]\right] &= \underset{p}{argmin}\left[Var_i\left(\theta_j^* \mid s_{pij}\right)\right] \\
&= \underset{p}{argmin}\frac{\tau_s+\tau_{Ipji}}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s}
\end{aligned}
\tag{12}
$$

According to Assumption 1, $\tau_{IFji} > \tau_{IAji}$. The result immediately follows. □

Proof of Proposition 4

Proof. Starting from Equation (4), it is trivial to derive:

$$
\left|\theta_{ij1} - \theta_{ij0}\right| = \gamma_{pij}\left|s_{pij} - I_{pji} - \theta_{ij0}\right|
$$

where $\gamma_{pij} \equiv \frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s} \in [0,1]$. Plugging in Equation (1) and rewriting I_{pji} and θ_{ij0} with the error terms, we get

$$
\begin{aligned}
\left|\theta_{ij1} - \theta_{ij0}\right| &= \gamma_{pij}\left|\theta_j^* + I_{pj}^* + \varepsilon_{pij} - I_{pj}^* - \nu_{pi} - \theta_j^* - \eta_{ij0}\right| \\
&= \gamma_{pij}\left|\varepsilon_{pij} - \nu_{pji} - \eta_{ij0}\right|
\end{aligned}
$$

Taking expectations over the continuum of agents,

$$
E\left[\left|\theta_{ij1} - \theta_{ij0}\right| \mid p = F\right] = \int\int\int\int \gamma_{pij}\left|\varepsilon_{pij} - \nu_{pji} - \eta_{ij0}\right| dF\left(\varepsilon_{pij}\right) dF\left(\nu_{pji}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Fij}, \gamma_{Aij}\right)
$$

We need to prove that $\frac{E\left[\left|\theta_{ij1}-\theta_{ij0}\right| \mid p=F\right]}{E\left[\left|\theta_{ij1}-\theta_{ij0}\right| \mid p=A\right]}$. Take M_j such that:

$$0<M_j<\frac{\int\int\int \gamma_{Fij}\left|\varepsilon_{Fij}-\eta_{ij0}\right| dF\left(\varepsilon_{Fij}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Fij},\gamma_{Aij}\right)-\int\int\int \gamma_{Aij}\left|\varepsilon_{Aij}-\eta_{ij0}\right| dF\left(\varepsilon_{Aij}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Aij},\gamma_{Aij}\right)}{2} \tag{13}$$

which exists because $\gamma_{Fij}>\gamma_{Aij}$ according to Assumption 1. Then,

$$\begin{aligned}
\frac{E\left[\left|\theta_{ij1}-\theta_{ij0}\right| \mid p=F\right]}{E\left[\left|\theta_{ij1}-\theta_{ij0}\right| \mid p=A\right]} &= \frac{\int\int\int\int \gamma_{Fij}\left|\varepsilon_{Fij}-\nu_{Fji}-\eta_{ij0}\right| dF\left(\varepsilon_{Fij}\right) dF\left(\nu_{Fji}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Fij},\gamma_{Aij}\right)}{\int\int\int\int \gamma_{Aij}\left|\varepsilon_{Aij}-\nu_{Aji}-\eta_{ij0}\right| dF\left(\varepsilon_{Aij}\right) dF\left(\nu_{Aji}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Fij},\gamma_{Aij}\right)} \\
&> \frac{\int \gamma_{Fij}\left|\varepsilon_{Fij}-\eta_{ij0}\right| dF(\cdot)-\int \gamma_{Fij}\left|\nu_{Fji}\right| dF(\cdot)}{\int \gamma_{Aij}\left|\varepsilon_{Aij}-\eta_{ij0}\right| dF(\cdot)+\int \gamma_{Aij}\left|\nu_{Aji}\right| dF(\cdot)} \\
&> \frac{\int\int\int \gamma_{Fij}\left|\varepsilon_{Fij}-\eta_{ij0}\right| dF\left(\varepsilon_{Fij}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Fij},\gamma_{Aij}\right)-M_j}{\int\int\int \gamma_{Aij}\left|\varepsilon_{Aij}-\eta_{ij0}\right| dF\left(\varepsilon_{Aij}\right) dF\left(\eta_{ij0}\right) dF\left(\gamma_{Aij},\gamma_{Aij}\right)+M_j} \\
&> 1
\end{aligned}$$

I used triangle inequality both on the denominator and the numerator for the first inequality. $dF(\cdot)$ is an abuse of notation for simplicity, indicating the distribution of all relevant variables. For the second inequality, I applied Assumption 2. The final inequality can be derived by applying Equation (13).

$\frac{E\left[\left|\theta_{ij1}-\theta_{ij0}\right| \mid p=F\right]}{E\left[\left|\theta_{ij1}-\theta_{ij0}\right| \mid p=A\right]}>1$ implies that there will be more learning when the reader receives the signal from the familiar source than from the alien source. It then immediately implies that selective exposure to the familiar source will give, on average, a bigger movement in beliefs than an equal-chance encounter. □

Proof of Proposition 5

Proof. Taking expectations of Equation (4) over continuum of i, $E\left(\theta_{ij0}\right)=E\left(\theta_j^*+\eta_{ij0}\right)=\theta_j^*$. On the other hand, after taking the signal from F,

$$\begin{aligned}
E\left(E_i\left[\theta_j \mid s_{Fij}\right]\right) &= E\left[(1-\gamma_{Fij})\theta_{ij0}+\gamma_{Fij}\left(s_{Fij}-I_{Fji}-\theta_{ij0}\right)\right] \\
&= E\left[(1-\gamma_{Fij})\left(\theta_j^*+\eta_{ij0}\right)+\gamma_{Fij}\left(\theta_j^*+I_{Fj}^*+\varepsilon_{Fij}-0-\theta_j^*-\eta_{ij0}\right)\right] \\
&= (1-\gamma_{Fij})\theta_j^*+\gamma_{Fij}\cdot I_{Fj}^*
\end{aligned}$$

where $\gamma_{pij}\equiv\frac{\tau_{Ipji}\tau_s}{\tau_s\tau_{ij0}+\tau_{ij0}\tau_{Ipji}+\tau_{Ipji}\tau_s}\in[0,1]$.

When the selection is shut down, half of the signals come from A. Note that:

$$E\left(E_i\left[\theta_j \mid s_{Aij}\right]\right)=E\left[(1-\gamma_{Aij})\left(\theta_j^*+\eta_{ij0}\right)+\gamma_{Aij}\left(\theta_j^*+I_{Aj}^*+\varepsilon_{Aij}-I_{Aji}-\theta_j^*-\eta_{ij0}\right)\right]=\theta^*$$

Therefore, $E\left(E_i\left[\theta_j \mid s_{random,ij}\right]\right) = \frac{1}{2} \cdot E\left(E_i\left[\theta_j \mid s_{Fij}\right]\right) + \frac{1}{2} \cdot E\left(E_i\left[\theta_j \mid s_{Aij}\right]\right)$ will be closer to θ^* and farther from I^*_{Fj}, compared with $E\left(E_i\left[\theta_j \mid s_{Fij}\right]\right)$. $\square$

B Appendix Tables

See next page.

Appendix Table 1. Summary Statistics

Subsample:	Sample for random provision of articles			Sample for random assignment to experimental groups		
	Observations (1)	Mean (2)	Std. Dev. (3)	Observations (4)	Mean (5)	Std. Dev. (6)
Gender (1=f, 2=m)	1420	1.617	0.486	352	1.639	0.481
Age	1420	31.141	7.696	352	30.824	7.427
Political Knowledge Score ([0,5] scale)	1420	2.652	1.678	352	2.855	1.645
Trust in Democratic Party ([-1,1] scale)	1420	0.115	0.529	352	0.134	0.54
Trust in Saenuri party ([-1,1] scale)	1420	-0.701	0.434	352	-0.741	0.399
Trust in Justice party ([-1,1] scale)	1420	0.169	0.557	352	0.224	0.556
Whether spent time on any TV news	1420	0.563	0.496	352	0.574	0.495
Whether spent time on ground wave TV news	1420	0.397	0.489	352	0.389	0.488
Whether spent time on cable TV news	1420	0.308	0.462	352	0.318	0.466
Whether spent time on any Internet news	1420	0.849	0.358	352	0.886	0.318
Whether spent time on Internet news portal service	1420	0.77	0.421	352	0.801	0.4
Whether spent time on a newspaper website	1420	0.248	0.432	352	0.284	0.452
Whether spent time on reading paper news	1420	0.131	0.338	352	0.153	0.361
Time spent on news media last week (minutes)	1420	166.754	284.999	352	190.739	386.48

Note: This table provides summary statistics for all baseline characteristics that are observed.

Appendix Table 2. Balance: random provision of articles

	(1) Age	(2) Gender (1=f, 2=m)	(3) Political knowledge score: [0,5]	(4) Trust in Democrat [-1,1] scale	(5) Trust in Saenuri (Korean Republican) [-1,1] scale	(6) Trust in Justice Party (left-wing) [-1,1] scale	(7) Initial knowledge on issue [0,5] scale	(8) Belief reported before reading the article
Position of first article	-0.729*	-0.0264	0.0223	-0.0270	-0.00529	0.00220	0.0549	-0.00295
	(0.417)	(0.0249)	(0.0885)	(0.0283)	(0.0218)	(0.0305)	(0.0607)	(0.0132)
Constant	31.67***	1.626***	2.927***	0.155***	-0.754***	0.238***	1.927***	0.255***
	(0.552)	(0.0289)	(0.0946)	(0.0326)	(0.0220)	(0.0329)	(0.0760)	(0.0112)
Pool FE	Y							
# of users	1420							
Observations (user × issue)	7,794							

Note: Each column in this table originates from a separate OLS regression of baseline characteristics of the users on the position of the first article. User-issue level observations are pooled together. The “Position of the article” is the average position of the article reported by the users after reading it. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 3. Balance and attrition: between treatment groups

Subsample:	Users who finished the grace period			User × issue combination w/ at least one article read after grace period		
	(1)	(2)	(3)	(4)	(5)	(6)
	Sample mean	[Source-Name] – [No-Choice]	[Source-Position] – [No-Choice]	Sample mean	[Source-Name] – [No-Choice]	[Source-Position] – [No-Choice]
Age	30.82	-0.276	1.144	31.28	0.198	1.149
	[7.42]	[.895]	[1.339]	[7.54]	[1.104]	[1.529]
Gender (1=f, 2=m)	1.64	0.038	0.097	1.64	0.065	0.093
	[.48]	[.059]	[.077]	[.48]	[.067]	[.09]
Political Knowledge Score [0,5] scale	2.86	-0.183	0.272	2.97	-0.156	.519*
	[1.64]	[.205]	[.289]	[1.63]	[.226]	[.301]
Time spent on news media last week (min)	190.7	121.6*	55.92	202.9	137.9	14.63
	[386]	[70.98]	[40.30]	[439.1]	[96.35]	[26.87]
Whether spent time on ground wave TV news	0.39	0.022	-0.03	0.39	0.072	-0.041
	[.49]	[.061]	[.081]	[.49]	[.07]	[.092]
Whether spent time on cable TV news	0.32	.101*	0.048	0.33	0.086	0.057
	[.47]	[.059]	[.079]	[.47]	[.068]	[.093]
Whether spent time on internet News portal service	0.8	0.04	0.022	0.81	0.023	0.013
	[.4]	[.048]	[.067]	[.4]	[.055]	[.074]
Whether spent time on a newspaper website	0.28	0.015	-0.048	0.29	-0.01	-0.124
	[.45]	[.057]	[.073]	[.45]	[.066]	[.077]
Whether spent time on reading paper news	0.15	-0.068	-0.032	0.15	-0.036	-0.004
	[.36]	[.041]	[.06]	[.36]	[.048]	[.07]
Trust in liberal party [-1,1] scale	0.13	0.068	0.118	0.15	0.001	0.143
	[.54]	[.066]	[.096]	[.54]	[.076]	[.104]
Trust in conservative Party [-1,1] scale	-0.74	0.038	-0.05	-0.77	0.076	-0.046
	[.4]	[.052]	[.062]	[.38]	[.058]	[.054]
Trust in progressive party [-1,1] scale	0.22	0.001	0.05	0.27	-0.05	0.082
	[.55]	[.065]	[.108]	[.55]	[.077]	[.115]
Retention				0.61	0.038	0.055
				[.49]	[.042]	[.058]
# Users	352	352	352	322	322	322
Observations (user × issue)	2816	2816	2816	1711	1711	1711

Note: The Source-Name Group was allowed to select the news source, with the source names given. The Source-Position Group was allowed to select the news source, with the sources' positions on the issue. The No-Choice Group got randomly provided articles. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 4. Attrition is not explained by key variables

Subsample:	Users who finished the grace period			
	(1)	(2)	(3)	(4)
	Attrition dummy	Attrition dummy	Attrition dummy	Attrition dummy
Source-Name Group	-0.0380			-0.0364
(selection allowed; with source names shown)	(0.0422)			(0.0437)
Source-Position Group	-0.0548			-0.0519
(selection allowed; with only positions shown)	(0.0584)			(0.0559)
Extreme		-0.0194		-0.0688
		(0.0828)		(0.0785)
|Baseline Belief - Party Position|		0.0966		0.0409
		(0.0720)		(0.0698)
Confidence		-0.0179		0.0629
		(0.0626)		(0.0625)
Age			-0.00424	-0.00355
			(0.00275)	(0.00290)
Gender (1=f, 2=m)			0.0411	0.0418
			(0.0424)	(0.0448)
Political Knowledge Score ([0,5] scale)			-0.0180	-0.0190
			(0.0122)	(0.0128)
Time spent in news media last week (min)			-2.95e-05	1.13e-05
			(3.27e-05)	(8.49e-05)
Trust in liberal party ([-1,1] scale)			0.00606	-0.00114
			(0.0432)	(0.0441)
Trust in conservative party ([-1,1] scale)			0.0290	0.0250
			(0.0526)	(0.0577)
Trust in progressive party ([-1,1] scale)			-0.0571	-0.0794*
			(0.0442)	(0.0460)
Constant	0.409***	0.389***	0.551***	0.513***
	(0.0242)	(0.0495)	(0.0938)	(0.106)
Other baseline characteristics	Y	Y	Y	Y
p-value for F-stat	0.505	0.572	0.150	0.224
# Users	352	352	352	352
Observations (user × issue)	2,816	2,664	2,816	2,664

Note: “Other baseline characteristics” include five dummy variables indicating whether the user spent time on (i) ground wave TV news (ii) cable TV news (iii) Internet news portal service (iv) newspaper websites, or (v) reading paper news (omitted for brevity, no statistical significance). “Baseline Belief” ([0,1]) is the reported belief at the baseline. “Confidence” ([0,1]) is the reported confidence in the belief at the baseline. “Extreme” is a dummy variable indicating whether the belief ([0,1]) equals 0 or 1 at the baseline. “Party Position” is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 5. Users who completed grace period vs. dropouts

	(1) Age	(2) Gender (1=f, 2=m)	(3) Political knowledge score: [0,5]	(4) Trust in Democrat [-1,1] scale	(5) Trust in Saenuri (Korean Republican) [-1,1] scale	(6) Trust in Justice Party (left-wing) [-1,1] scale	(7) Whether spent time on ground wave TV news	(8) Whether spent time on cable TV news	(9) Whether spent time on Internet news portal service	(10) Whether spent time on a newspaper website	(11) Whether spent time on reading paper news	(12) Time spent on news media last week (mins)
Finished Grace Period	-0.421	0.0297	0.270***	0.0240	-0.0540**	0.0736**	-0.0106	0.0129	0.0408	0.0481*	0.0298	31.89*
	(0.473)	(0.0299)	(0.103)	(0.0325)	(0.0266)	(0.0342)	(0.0301)	(0.0284)	(0.0258)	(0.0265)	(0.0207)	(17.50)
Constant	31.25***	1.610***	2.585***	0.109***	-0.687***	0.150***	0.400***	0.305***	0.760***	0.236***	0.124***	158.8***
	(0.236)	(0.0149)	(0.0512)	(0.0162)	(0.0133)	(0.0170)	(0.0150)	(0.0141)	(0.0129)	(0.0132)	(0.0103)	(8.714)
Observations (user)						1420						

Note: Each column in this table originates from a separate OLS regression of baseline characteristics of the users on whether the user finished the grace period. User-issue level observations are pooled together. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 6. The effect of article: not transient

Subsample:	Random provision		
	after reading first article	after reading second article	after reading third article
	(1)	(2)	(3)
	Belief reported after reading the last article		
Belief reported before reading the first article	0.616***	0.575***	0.554***
	(0.0151)	(0.0245)	(0.0314)
Article position [Current]	0.0759***		
	(0.00972)		
Article position [Lag 1]		0.0428***	
		(0.0139)	
Article position [Lag 2]			0.0304**
			(0.0151)
Constant	0.133***	0.00673	0.0457
	(0.0140)	(0.0382)	(0.0361)
Round FE, user controls	Y	Y	Y
Article Pool FE (includes Issue FE)	Y	Y	Y
# of users	1420	599	280
Observations (user×issue×round)	7,794	3,883	2,453

Note: Each column in this table originates from a separate OLS regression of belief after reading an article on respective regressors. “User Controls” include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Included are the observations where users were provided randomly selected articles: that is, I include all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. Article pool from which an article is randomly drawn is changing over time – article pool FE is added to the regression (for lagged regressions in Columns 2 and 3, lagged article pool FE is added). Article position is the average position of the article reported by all the users after reading it. “Prior Confidence” ($\in [0,1]$) is the reported confidence in the prior belief. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 7. Belief updating patterns: dropping obs. w/ extreme priors

Subsample:	Random provision w/o observations with Belief $_{pre}$ ∈ {0,1}			
	(1)	(2)	(3)	(4)
	Belief $_{post}$			
Belief $_{pre}$	0.646***	0.639***	0.599***	0.250***
	(0.0170)	(0.0168)	(0.0183)	(0.0531)
Belief $_{pre}$ × Prior confidence				0.489***
				(0.0757)
Article position	0.102***	0.103***	0.102***	0.190***
	(0.0141)	(0.0144)	(0.0145)	(0.0561)
Article position × Prior confidence				-0.137*
				(0.0751)
Prior Confidence				0.178
				(0.291)
Constant	0.0601***	0.0752***	0.140***	-0.105
	(0.00874)	(0.00914)	(0.0182)	(0.279)
Round FE	N	Y	Y	Y
User controls	N	N	Y	Y
Article Pool FE (includes Issue FE)	Y	Y	Y	Y
Article Pool FE × Prior Confidence				Y
# of users	1,130			
Observations (user×issue×round)	5,062			

Note: Each column in this table originates from a separate OLS regression of belief after reading an article on respective regressors. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. Included are the observations where users were provided randomly selected articles: that is, I include all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. Article pool from which an article is randomly drawn is changing over time – article pool FE is added to the regression. Article position is the average position of the article reported by all the users after reading it. "Prior Confidence" (∈ [0,1]) is the reported confidence in the prior belief. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 8. Belief updating patterns: consistent with typical Bayesian models (full sample vs 1st round only)

Subsample:	Random provision		1st round, random provision		1st round, random provision (excluding grace period)	
	(1)	(2)	(3)	(4)	(5)	(6)
	Posterior belief	Posterior belief	Posterior belief	Posterior belief	Posterior belief	Posterior belief
Belief $_{pre}$	0.616***	0.297***	0.577***	0.227***	0.570***	0.231**
	(0.0151)	(0.0380)	(0.0165)	(0.0494)	(0.0349)	(0.104)
Belief $_{pre}$ × Prior confidence		0.403***		0.478***		0.465***
		(0.0501)		(0.0644)		(0.133)
Article position	0.0760***	0.231***	0.117***	0.192***	0.106***	0.295**
	(0.00972)	(0.0444)	(0.0169)	(0.0608)	(0.0282)	(0.132)
Article position × Prior confidence		-0.205***		-0.109		-0.251
		(0.0526)		(0.0780)		(0.161)
Constant	0.133***	0.205***	0.145***	-0.00859	0.146***	-1.824***
	(0.0140)	(0.0263)	(0.0185)	(0.175)	(0.0358)	(0.566)
Issue FE	Y	Y	Y	Y	Y	Y
Control for prior confidence	N	Y	N	Y	N	Y
Round FE	Y	Y	N/A	N/A	N/A	N/A
Article Pool FE	Y	Y	Y	Y	Y	Y
Article Pool FE × Prior Confidence	N	Y	N	Y	N	Y
User controls	Y	Y	Y	Y	Y	Y
# of users	1,420	1,420	1,420	1,420	207	207
Observations (user×issue×round)	7,794	7,794	3,901	3,901	1,067	1,067

Note: Each column in this table originates from a separate OLS regression of belief after reading an article on respective regressors. “User Controls” include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. Included are the observations where users were provided randomly selected articles: that is, I include all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. Article pool from which an article is randomly drawn is changing over time – article pool FE is added to the regression. Article position is the average position of the article reported by all the users after reading it. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 9. Belief updating patterns: news source fixed effect

Subsample:	Random provision		
	(1)	(2)	(3)
		Belief $_{post}$	
Belief $_{pre}$	0.616***	0.616***	0.579***
	(0.0151)	(0.0150)	(0.0151)
Article position	0.0759***	0.0786***	0.0820***
	(0.00972)	(0.0153)	(0.0155)
Constant	0.133***	0.0831**	0.415***
	(0.0140)	(0.0332)	(0.0618)
Round FE, article pool FE, user controls	Y	Y	Y
News source FE	N	Y	Y
News source FE × Most trusted party FE	N	N	Y
# of users		1,420	
Observations (user×issue×round)		7,794	

Note: Each column in this table originates from a separate OLS regression of the belief after reading an article on respective regressors. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. Included are the observations where users were provided randomly selected articles: that is, all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. The article pool from which an article is randomly drawn changes over time—article pool FE is added to the regression. Most trusted party FE is the dummy variable indicating the most trusted party of the user at the baseline. "Article position" is the average position of the article reported by all the users after reading it. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 10. Readers select their partisan media (and not prior-confirming ones): subsample analysis

	Source-Name Group (source selection allowed; with source names shown)			
Subsample:	Name *and* average position (G1)	Name only (G2)	Distance > p(50)	Distance > p(75)
	(1)	(2)	(3)	(4)
	Position of the article read			
Prior belief	-0.0520	-0.0230	-0.0327	-0.0748
	(0.0514)	(0.0407)	(0.0305)	(0.0483)
Party position of the reader's most-trusted party on the issue	0.325***	0.366***	0.347***	0.520***
	(0.114)	(0.0885)	(0.0850)	(0.104)
Constant	0.364***	0.370***	0.478***	0.485***
	(0.0445)	(0.0707)	(0.0445)	(0.0689)
Article Pool FE (includes Issue FE)	Y	Y	Y	Y
Round FE	Y	Y	Y	Y
User Controls	Y	Y	Y	Y
# of users	42	42	77	68
Observations (user×issue×round)	743	746	711	337

Note: Each column in this table originates from a separate OLS regression of position of the chosen article on the latest position users held before making the choice (prior) and the party position of the reader's most-trusted party on the issue. The first column contains a subsample of users who were assigned to the subgroup of the Source-Name group where both the sources' names and their average positions were available. The second column shows the subsample where only the sources' names were available. The third column shows the subsample where the prior belief and the party position are farther than the median distance (0.132). The fourth is similar, except the cutoff is at the 75th percentile (0.232). "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "Expected source position" is the average position on an issue of the articles written by the sources that share the same favorite party as the source at hand. The favorite party of the source is determined by the party that has the smallest squared position distance across issues with the source. Article pool changes over time—article pool FE is added to the regression. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 11. Source selection: comparison between treatment groups (finest subgroups)

Group:	*Source-Name Group*		*Source-Position Group*	*No-Choice Group*			
	G1 (selection allowed; with names + positions shown)	G2 (selection allowed; with only names shown)	G3 (selection allowed; with only positions shown)	G4 (random provision; name + position notified)	G5 (random provision; only name notified)	G6 (random provision; only position notified)	G7 (random provision; no info notified)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Position of the article read						
Prior belief	-0.0520	-0.0230	-0.0461	-0.0228	-0.0333	0.0516	-0.0160
	(0.0514)	(0.0407)	(0.0471)	(0.0642)	(0.0383)	(0.0327)	(0.0363)
Party position of the reader's most-trusted party on the issue	0.325***	0.366***	-0.00159	0.103	0.175	0.0618	0.0619
	(0.114)	(0.0885)	(0.127)	(0.121)	(0.107)	(0.104)	(0.149)
Constant	0.364***	0.370***	0.470***	0.464***	0.487***	0.380***	0.445***
	(0.0445)	(0.0707)	(0.0582)	(0.0541)	(0.0395)	(0.0496)	(0.0431)
Pool, Round FE, User Controls	Y	Y	Y	Y	Y	Y	Y
# of users	42	42	38	48	48	53	45
Observations	743	746	680	514	606	1,018	976

Note: Each column in this table originates from a separate OLS regression of position of the chosen article on the latest position users held before making the choice (prior) and the party position of the reader's most-trusted party on the issue. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added when applicable.). "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "Expected source position" is the average position on an issue of the articles written by the sources that share the same favorite party as the source at hand. The favorite party of the source is determined by the party that has smallest squared position distance across issue with the source. The article pool from which an article is randomly drawn changes over time—article pool FE is added to the regression. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 12. Learning: robust to drop any issue

Drop:	Issue 1	Issue 2	Issue 3	Issue 4	Issue 5	Issue 6	Issue 7	Issue 8
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$\lvert Belief_{post} - Belief_{pre} \rvert$							
Source-Name Group	0.0189*	0.0288**	0.0246**	0.0180*	0.0224**	0.0224**	0.0248**	0.0175*
(selection allowed; with source names shown)	(0.0111)	(0.0118)	(0.0111)	(0.0104)	(0.0110)	(0.0104)	(0.0107)	(0.0105)
Source-Position Group	-0.00580	-0.00633	0.00248	-0.00332	-0.00704	0.000430	0.00213	-0.000554
(selection allowed; with only positions shown)	(0.0166)	(0.0168)	(0.0177)	(0.0163)	(0.0155)	(0.0171)	(0.0175)	(0.0171)
Constant	0.179***	0.143***	0.139***	0.139***	0.140***	0.139***	0.146***	0.142***
(Omitted: *No-Choice Group* – random provision of articles)	(0.0171)	(0.0200)	(0.0201)	(0.0197)	(0.0192)	(0.0197)	(0.0197)	(0.0197)
p-value: *Source-Name* = *Source-Position*	0.170	0.0635	0.253	0.237	0.0922	0.238	0.237	0.333
User controls, Issue FE	Y	Y	Y	Y	Y	Y	Y	Y
# of users	330	336	330	328	332	330	333	333
Observations (user×issue)	1,561	1,560	1,562	1,558	1,574	1,576	1,561	1,578

Note: Each column in this table originates from a separate OLS regression of the absolute value of change in belief on the treatment group dummies. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. $Belief_{pre}$ ([0,1]) is the reported belief at the baseline. $Belief_{post}$ ([0,1]) is the reported belief after reading an article in the experimental period (not grace period). Extreme opinion is the belief that equals to 0 or 1. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 13. Moderation of extreme belief: robust to drop any issue

Drop:	Issue 1	Issue 2	Issue 3	Issue 4	Issue 5	Issue 6	Issue 7	Issue 8
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	$Extreme_{post}$							
Source-Name Group	-0.0593*	-0.0773**	-0.0682**	-0.0701**	-0.0782***	-0.0713**	-0.0616**	-0.0645**
(selection allowed; with source names shown)	(0.0302)	(0.0320)	(0.0282)	(0.0296)	(0.0281)	(0.0300)	(0.0290)	(0.0294)
Source-Position Group	-0.0169	0.00397	-0.0135	-0.00338	-0.0232	-0.00669	-0.00435	-0.0179
(selection allowed; with only positions shown)	(0.0447)	(0.0528)	(0.0453)	(0.0472)	(0.0460)	(0.0493)	(0.0497)	(0.0462)
$Extreme_{pre}$	0.374***	0.366***	0.367***	0.366***	0.371***	0.376***	0.367***	0.368***
	(0.0289)	(0.0282)	(0.0273)	(0.0281)	(0.0276)	(0.0279)	(0.0287)	(0.0283)
Constant	-0.0115	0.249***	0.250***	0.241***	0.243***	0.233***	0.236***	0.242***
(Omitted: *No-Choice Group* – random provision of articles)	(0.0426)	(0.0540)	(0.0506)	(0.0506)	(0.0511)	(0.0517)	(0.0505)	(0.0497)
p-value: *Source-Name* = *Source-Position*	0.372	0.138	0.239	0.171	0.247	0.210	0.265	0.336
User controls, Issue FE	Y	Y	Y	Y	Y	Y	Y	Y
# of users	330	336	330	328	332	330	333	333
Observations (user×issue)	1,561	1,560	1,562	1,558	1,574	1,576	1,561	1,578

Note: Each column in this table originates from a separate OLS regression. $Extreme_{post}$ is a dummy variable indicating whether the belief ([0,1]) equals 0 or 1 after reading an article in the experimental period (not the grace period). $Extreme_{pre}$ is an analogous dummy variable for the baseline belief. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 14. Moderation of extreme beliefs (continuous measure of extremism)

Model:	Basic	First Difference (FD)	Lagged Dependent Variable (LDV)	Baseline balance
	(1)	(2)	(3)	(4)
	$(\text{Belief}_{post}-0.5)^2$	$(\text{Belief}_{post}-0.5)^2$- $(\text{Belief}_{pre}-0.5)^2$	$(\text{Belief}_{post}-0.5)^2$	$(\text{Belief}_{pre}-0.5)^2$
Sample Mean	0.120	-0.0130	0.120	0.133
Source-Name Group (selection allowed; with source names shown)	-0.00650 (0.00752)	-0.0140** (0.00604)	-0.0101* (0.00551)	0.00750 (0.00806)
Source-Position Group (selection allowed; with only positions shown)	0.00579 (0.0124)	-0.0260*** (0.00995)	-0.00946 (0.00930)	0.0318** (0.0128)
$(\text{Belief}_{pre}-0.5)^2$			0.480*** (0.0249)	
Constant (Omitted: *No-Choice Group* – random provision of articles)	0.124*** (0.0121)	-0.0108 (0.0112)	0.0591*** (0.00987)	0.134*** (0.0144)
p-value: *Source-Name* = *Source-Position*	0.344	0.245	0.947	0.0725
User controls, Issue FE	Y	Y	Y	Y
# of users	336	336	336	336
Observations (user×issue)	1,790	1,790	1,790	1,790

Note: Each column in this table originates from a separate OLS regression. Belief_{pre} ([0,1]) is the reported belief at the baseline. Belief_{post} ([0,1]) is the reported belief after reading an article in the experimental period (not the grace period). "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 15. Ideological distance: a good proxy for unfamiliarity

Subsample:	Random provision	
	(1)	(2)
	Evaluation error = | User evaluation of article position – Article position |	
"Source's ideology distance to the reader"	0.18***	0.10***
	(0.026)	(0.029)
Constant	0.11***	0.10***
	(0.018)	(0.018)
Issue, round FE, user controls	Y	Y
Control for |Article position - Belief $_{pre}$|	N	Y
# of users	796	796
Observations (user×issue×round)	2,300	2,300

Note: Each column in this table originates from a separate OLS regression. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue-round level observations are pooled together, and issue FE and round FE are added. Included are the observations where users were provided randomly selected articles: that is, I include all the observations of the group with random provision (No-Choice Group) and the observations collected during the grace period for the other groups. User evaluation of article position is the user-reported article's position evaluated right after reading the article. Article position is the average position of the article reported by all the users after reading it. "Source's ideology distance to the reader" ($\in [\mathbf{0}, \mathbf{1}]$) is the mean absolute distance between the baseline user belief and the source's average article position. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 16. Evolution of beliefs by finest subgroup

	(1)	(2)	(3)
	$\mid Belief_{post}$ - $Belief_{pre} \mid$	$Extreme_{post}$	$\mid Belief_{post}$ - Party Position$\mid$
Sample Mean	0.140	0.284	0.166
G2	-0.00935	0.0777*	0.00175
(selection allowed; with only names shown)	(0.0183)	(0.0454)	(0.0175)
G3	-0.0285	0.0978*	0.00298
(selection allowed; with only positions shown)	(0.0197)	(0.0550)	(0.0149)
G4	-0.00381	0.0700	0.00774
(random provision; name+position notified)	(0.0157)	(0.0486)	(0.0137)
G5	-0.0317**	0.113**	0.00449
(random provision; only name notified)	(0.0156)	(0.0465)	(0.0131)
G6	-0.0389**	0.129**	-0.00699
(random provision; only position notified)	(0.0165)	(0.0523)	(0.0131)
G7	-0.0266	0.0986**	0.0225
(random provision; no info notified)	(0.0170)	(0.0453)	(0.0140)
$Extreme_{pre}$		0.368***	
		(0.0259)	
$\mid Belief_{pre}$ - Party Position$\mid$			0.360***
			(0.0343)
Constant	0.159***	0.146**	0.0934***
(Omitted: *G1* –	(0.0227)	(0.0675)	(0.0203)
selection allowed; with names+positions shown)			
User controls, Issue FE	Y	Y	Y
# of users	336	336	316
Observations (user×issue)	1,790	1,790	1,682

Note: Each column in this table originates from a separate OLS regression. $Belief_{pre}$ ([0,1]) is the reported belief at the baseline. $Belief_{post}$ ([0,1]) is the reported belief after reading an article in the experimental period (not the grace period). $Extreme_{post}$ is a dummy variable indicating whether the belief ([0,1]) equals to 0 or 1 after reading an article in the experimental period (not grace period). $Extreme_{pre}$ is an analogous dummy variable for the baseline belief. "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 17. Convergence to party's position – alternative definitions

	(1)	(2)	(3)
	$\|\text{Belief}_{post}$ - Party Position$\|$	$\|\text{Belief}_{post}$ - Party Position (Alt. 1)$\|$	$\|\text{Belief}_{post}$ - Party Position (Alt. 2)$\|$
Sample Mean	0.166	0.165	0.171
Source-Name Group	-0.00487	-0.00702	-0.00866
(selection allowed; with source names shown)	(0.00956)	(0.00966)	(0.00909)
Source-Position Group	-0.00323	-0.00385	-0.00771
(selection allowed; with only positions shown)	(0.0115)	(0.0113)	(0.0108)
$\|\text{Belief}_{pre}$ - Party Position (relevant alternative)$\|$	0.364***	0.364***	0.329***
	(0.0340)	(0.0337)	(0.0327)
Constant	0.102***	0.0939***	0.0757***
(Omitted: *No-Choice Group* – random provision of articles)	(0.0184)	(0.0183)	(0.0172)
p-value: *Source-Name* = *Source-Position*	0.903	0.814	0.941
User controls, Issue FE	Y	Y	Y
# of users	316	316	316
Observations (user×issue)	1,682	1,682	1,682

Note: Each column in this table originates from a separate OLS regression. Belief_{pre} ([0,1]) is the reported belief at the baseline. Belief_{post} ([0,1]) is the reported belief after reading an article in the experimental period (not grace period). "Party Position" is the average baseline belief of the users who share the same most-trusted party at the baseline and have baseline confidence ([0,1]) of 0.8 or more. "Party Position (Alt. 1)" is the average baseline belief of the users who share the same most-trusted party at the baseline and whose trust in the most-trusted party is higher than 0.5 (≈ median trust conditional on most-trusted). "Party Position (Alt. 2)" is the average baseline belief of all the users who share the same most-trusted party at the baseline. For each column, the relevant measure of party position is used to control for the initial distance. "User Controls" include baseline characteristics of the users: age, gender, political knowledge score, sources of news consumption, and time spent on news consumption. User-issue level observations are pooled together, and issue FE is added. Standard errors (in parentheses below coefficients) are clustered by user. Stars are based on p-values: *** p<0.01, ** p<0.05, * p<0.1.