

An information-based explanation for partisan media sorting

Journal of Theoretical Politics

1–28

© The Author(s) 2022

Article reuse guidelines:

sagepub.com/journals-permissions

DOI: 10.1177/09516298221122094

journals.sagepub.com/home/jtp**Anthony Fowler** 

Harris School of Public Policy, University of Chicago, Chicago, Illinois, USA

Kisoo Kim

Harris School of Public Policy, University of Chicago, Chicago, Illinois, USA

Abstract

Partisan voters tend to seek political news from media sources that match their predispositions. Scholars and pundits often attribute this partisan media sorting to psychological biases, and they typically assume that it leads voters to make worse decisions at the ballot box. To reinterpret this evidence and provide an alternative explanation, we develop two formal models of media choice—one in which voters only want to hear good news about their party and another in which voters only care about making good electoral decisions. Both models predict partisan media sorting, so sorting does not constitute evidence that voters are poorly informed or that they are driven by psychological biases. However, the models do produce competing predictions about when voters will consume more or less news and about whether signals from the news should influence vote choices. Reassessing the empirical literature, we find some support for both explanations.

Keywords

Media, voting, partisanship, information, confirmation bias, reinterpreting, distinguishing

American voters are purportedly stuck in partisan echo chambers. We're told that Democrats read *The New York Times*, watch MSNBC, listen to NPR, follow Alexandria Ocasio-Cortez on Twitter, and eat avocado toast in urban cafes, while

Corresponding author:

Anthony Fowler, Harris School of Public Policy, University of Chicago, Chicago, Illinois, USA.

Email: anthony.fowler@uchicago.edu

Republicans read *The Wall Street Journal*, watch Fox News, listen to AM talk radio, follow Donald Trump on Truth Social, and eat hamburgers in roadside diners. Although this caricature surely overstates the extent of partisan sorting, Americans are more likely than not to obtain political information from a source that is aligned with their predispositions.

To the extent that partisans do sort into different media outlets, scholars and pundits tend to assume that this is bad for democracy and the result of psychological biases of voters. Studies often refer to *confirmation bias* or *cognitive dissonance*, implying that partisans sort because they've already made up their minds on political matters and want to avoid encountering news stories that conflict with their predispositions. For example, in an academic study of online political searches, Knobloch-Westerwick et al. (2015) write that "Such confirmation bias has been considered problematic from a perspective of normative democratic theory: Ideally, citizens in a democracy seek diverse information and exposure to opposing views" (p. 172). In a journalistic report for *The Guardian*, Grimes (2017) writes that "Echo chambers are dangerous ... Across the political spectrum we must all work harder to analyze our sources of information and our biases."

Is partisan media sorting necessarily a sign that voters are doing a bad job? If a voter only has the time, attention, and resources to read one newspaper or watch one cable news channel, what kind of source should she select if her goal is to determine which political candidate is best for her? In this paper, we show theoretically that even if a voter is not subject to confirmation bias and cares only about identifying and selecting the best candidate, she will often prefer ideologically aligned news sources.

Consider a Democratic voter who is trying to decide whether to reelect her incumbent Democratic governor. The voter cares about both ideology and valence. She knows that the governor matches her ideological preferences better than the Republican challenger, but she's less sure about the general competence of the governor, and she's hoping to learn about that by reading the news. If she gets no news, she's inclined to reelect the incumbent, but if the governor is doing a bad job, she would prefer the challenger. There are two newspapers—one with a liberal slant and one with a conservative slant, but she only has the time to read one. Although both newspapers convey meaningful information, she knows there is a relative bias. The liberal paper is more likely to say the Democratic governor is doing a good job, and the conservative paper is more likely to say the governor is doing a bad job, regardless of whether the governor is good or bad.

Which newspaper should the voter read? If she reads the conservative paper, she is not likely to receive much information that will be relevant for her vote. The conservative paper is likely to say that the Democratic governor is doing a bad job, in which case, she'll lower her beliefs about the quality of the governor, but only by a little. If the conservative paper says the Democratic governor is doing a good job, she'll increase her beliefs about the quality of the governor by a lot. But the voter was already inclined to support the Democrat anyway, so this change in beliefs is not consequential for her vote. The liberal paper is likely to say the Democratic governor is doing a good job, in which case, the voter will increase her beliefs, but only a little. Most importantly, if the liberal paper says the governor is doing a bad job, this is a very informative signal

that is highly relevant for the voter's decision. In this case, she shifts her beliefs about the governor downward by a lot, and she's glad she read the liberal paper because it paradoxically helped her realize that the Republican candidate was the better choice.

In this situation, even if the voter has no psychological biases and doesn't care per se about the slant of news articles she reads, she has an instrumental incentive to read the ideologically aligned newspaper. On average, reading the newspaper biased in her direction helps her to cast votes that are better for her. Just as it took Nixon to go to China to convince Republicans to ease Sino-American relations, it takes bad news about a Republican candidate (or good news about a Democratic candidate) from a trusted conservative source to convince Republicans to support a Democrat. And in these unlikely circumstances, the Republicans are glad they read that paper, which motivates them to generally read the conservative paper over the liberal one.¹

In this paper, we formalize the logic above and show the conditions under which partisans will have an information-based preference for an ideologically aligned news source. We compare and contrast this model with another model of news consumption in which people always vote in line with their partisan predisposition and they simply enjoy reading good news about their party and bad news about the other party. Both models predict that voters will typically prefer the ideologically aligned news source, so partisan media sorting alone is not evidence that media consumption is driven by psychological biases or that voters are bad at seeking political information and making electoral decisions. In other words, by writing down one model that we believe represents the conventional thinking in this literature and then writing another model that relies on a different mechanism but produces the same empirical pattern, we are doing what Ashworth et al. (2021) refer to as "reinterpreting."

Although both models predict partisan media sorting, they do make different predictions about when voters will consume more or less news and about whether signals from the news should influence vote choices, providing empirical opportunities to adjudicate between these two models of news consumption. Specifically, the model of confirmation bias predicts that people will consume more news when they expect their aligned candidate to be good, while the model of informed voting predicts that people will typically consume more news when they expect their aligned candidate to be bad. Discussing previous evidence and presenting some new evidence, we find some support for both models. To the extent that partisans sort into different media, some of it appears to result from the desire of voters to cast more informed votes. By producing new predictions that diverge between the two models and then testing those predictions with data, we are doing what Ashworth et al. (2021) refer to as "distinguishing."

We also pursue several extensions of the information-based model that speak to various public and academic debates. First, we discuss the normative implications of biased media, which are theoretically ambiguous. Next, we discuss the theoretically predicted effects of media choices on vote choices. Receiving political information from an aligned vs. misaligned source will sometimes make a voter more likely to support her party, but in other cases, it can make her *less* likely to support her party. This could explain some of the heterogeneous results in the literature on the effects of media slant on voting behavior. Then, we discuss the effects of media choice on beliefs, and we explain why Democrats and Republicans can have different beliefs about political facts

even if they started with the same prior beliefs and have no biases in the ways they process information. Lastly, although our models assumed that media polarization was exogenous, we use our model to discuss how media polarization could arise endogenously from either consumer demand or from a motivation of news outlets to influence elections.

Related literature

There is a large literature on the nature, extent, and origins of media slant. Researchers have utilized a variety of approaches to quantify the ideological slant of different news outlets (e.g., Groseclose and Milyo, 2005; Kayser and Peress, 2021; Puglisi and Snyder, 2015) and they have also investigated whether media slant is driven more by consumer demand or the ideological leanings of owners and journalists (e.g., Gentzkow and Shapiro, 2010; Larcinese et al., 2011). While there is debate about whether the media as a whole is biased in one direction or another, there is general agreement that some media outlets are more liberal or conservative than others. These ideological slants manifest themselves in different ways, including each outlet's choices of what stories to cover, what sources to cite, and what candidates and positions to endorse in opinion pieces.

Another literature has focused on the extent to which partisans sort into different media outlets. Although there is disagreement about the extent of sorting, partisans are more likely than not to receive political news from ideologically aligned sources (e.g., see Flaxman et al., 2016; Guess, 2021; Iyengar and Hahn, 2009; Stroud, 2008; Tyler et al., 2022). As media choices have proliferated over time (e.g., Prior, 2007) and since partisan sorting is greater for online news consumption (Gentzkow and Shapiro, 2011), there are concerns that partisan media sorting could become worse, although many popular accounts likely overstate the extent of media sorting.

Further studies have attempted to assess the effects of media polarization and media sorting. Prior (2007) argues that polarized media create polarized elections, primarily by increasing apathy among moderates. Levendusky (2013) argues that polarized media polarizes public opinion, making partisans more sure of their initial predispositions. The predominant view is that, for one reason or another, partisan media sorting is bad for democracy. Wolton (2019) presents a formal model in which a biased media outlet can be good for voter welfare because it alters the electoral incentives of incumbents and can induce them to exert more effort, but even in this model, biased media are bad for voter information and electoral selection.

Few studies have specifically studied *why* partisans tend to sort into like-minded media sources, with most implicitly assuming that this is attributable to confirmation bias or related psychological limitations of voters. Knobloch-Westerwick and Kleinman (2012) attempt to study whether selective exposure is explained more by confirmation bias or informational utility, but they simply test whether selective exposure varies depending on which party is expected to win the upcoming election, and they provide little theoretical rationale for their empirical tests. Gentzkow and Shapiro (2006) study a theoretical model in which news consumers simultaneously draw inferences about the quality of a media source and the state of the world. Partisans sort into ideologically aligned sources that they rationally view as more credible. Our model abstracts away

from this mechanism and shows that partisans will have an incentive to sort into like-minded media sources even if everyone knows that all sources are equally credible.

Our model of information-based media sorting is most closely related to theoretical models showing that advice from like-minded experts is more valuable for decision making than advice from a neutral or misaligned expert (Calvert, 1985; Crawford and Sobel, 1982; Patty, 2009; Suen, 2004). The theoretical mechanism we highlight that drives like-minded news to be more beneficial for vote choices is not a novel one. We have attempted to apply this insight from models of information to this particular setting in the simplest possible way that captures the importance of both ideology and valence in electoral selection. Our models also improve upon this literature by providing theoretical predictions that could allow us to distinguish between explanations based on either confirmation bias or information seeking.

There is also a large literature on the effects of media choices, media signals, and media slant on vote choices (e.g., DellaVigna and Kaplan, 2007; Gerber et al., 2009; Green et al., 2018; Hopkins and Ladd, 2014; Ladd and Lenz, 2009). As we later discuss, some of this evidence helps us to distinguish between different explanations for media sorting. And in other cases, we believe our information-based model provides a useful framework for thinking about media effects, and it might help us to understand some of the heterogeneous results in this literature.

Two models of news consumption

This section presents two different models of media choice. The first is a simple formalization of what we believe to be the conventional thinking in the political behavior literature on partisan media sorting. In this model, partisans enjoy consuming good news about their party and bad news about the other party. The second model assumes that voters care only about electing the better candidate. By showing that these two models that rely on very different kinds of voter preferences and motivations both produce partisan media sorting, we hope to theoretically reinterpret this empirical literature (see Ashworth et al., 2021). Although both models predict partisan media sorting, they do produce different predictions about when partisans will consume more or less news. To help us empirically distinguish between these two explanations for partisan media sorting, this section elucidates those divergent predictions, and the next section presents and discusses empirical evidence related to those predictions.

In both models, there is an upcoming election between an incumbent and challenger. The two candidates have different ideologies, and without loss of generality, assume the incumbent has ideology L and the challenger has ideology R . Each candidate has a valence V which is either *High* or *Low*. There is a single voter with ideology L who decides whether to reelect the incumbent or replace her with the challenger. The voter knows the candidates' ideologies but does not know their valence and believes the incumbent's valence is *High* with probability ρ while the challenger's valence is *High* with probability ρ' . The voter also believes that the valences of the incumbent and challenger are independent of one another, such that information about the incumbent's valence has no effect on beliefs about the challenger's valence.

There are two media outlets, L and R , that receive a common noisy signal $Y \in [0, 1]$ about the valence of the incumbent. Y is informative about the valence of the incumbent because when the incumbent's valence is *High*, Y follows a distribution with differentiable probability density function f_H and cumulative distribution function F_H , and when the incumbent's valence is *Low*, Y follows the functions f_L and F_L , and the two distributions satisfy the monotone likelihood ratio property, i.e., $\frac{f_H}{f_L}$ is increasing in Y . This implies that a higher value of a signal realization is more suggestive of the incumbent having *High* valence.

Each news outlet coarsens the signal into a binary report by reporting positively about the incumbent when Y exceeds its threshold and reporting negatively otherwise, but the two outlets have different threshold levels t_L , $t_R \in [0, 1]$ with $t_L < t_R$.² Because outlet L has a lower threshold than R , it is ex-ante more likely to report positively. Furthermore, because the outlets receive a common signal, whenever the two outlets report differently (when $t_L < Y < t_R$), L will report positively and R will report negatively.

To ease our analysis, we assume that the two thresholds t_L and t_R are symmetric (i.e., $t_L = 1 - t_R$) and the two density functions f_H and f_L are symmetric around $\frac{1}{2}$ (i.e., $f_H(y) = f_L(1 - y)$).³ Then, we can denote π_C to be $F_L(t_L) = 1 - F_H(t_R)$, or the probability that both media outlets correctly report on the quality of the incumbent (i.e., both send positive signals when the incumbent's valence is *High* and both send negative signals when the incumbent's valence is *Low*). And we denote π_W to be $F_H(t_L) = 1 - F_L(t_R)$, or the probability that both outlets report incorrectly, and π_D be $F_H(t_R) - F_H(t_L) = F_L(t_R) - F_L(t_L)$, or the probability that the two media receive an intermediate signal and send different signals, with only outlet L reporting positively. Then it follows that $\pi_C > \pi_W$ from the monotone likelihood property, meaning that the media signals are informative about the incumbent's valence. Table 1 shows how the probabilities of different media reports depend on the state of the world.

Knowing the reporting structure of the two outlets, the voter chooses whether to read L , R , or neither. In both models below, we assume the voter chooses the option that maximizes her expected utility (although her preferences will differ between models), and if she is indifferent between consuming news or not, she chooses not to consume news.⁴

Confirmation-Bias model

Suppose the voter already knows that she prefers the ideologically aligned incumbent, and she will vote for that candidate regardless of any new information that arises. Suppose further that she only tunes into political news for entertainment purposes. She

Table 1. How news reports depend on the quality of the incumbent.

	If $V = High$	If $V = Low$
Both outlets report that the incumbent is good	π_C	π_W
Only L reports that the incumbent is good	π_D	π_D
Both outlets report that the incumbent is bad	π_W	π_C

derives utility $\alpha > 0$ when she reads a news report that is good for the incumbent. She derives a disutility of -1 when she reads a news report that is bad for the incumbent.⁵ That is, her expected utility of news consumption is

$$\begin{aligned} EU(\text{Not Read}) &= 0, \\ EU(\text{Read } M) &= \alpha \Pr(Y > t_M) - \Pr(Y \leq t_M), \end{aligned}$$

for a news outlet $M \in \{L, R\}$.

Proposition 1. A news consumer in the confirmation-bias model always prefers her aligned outlet over the other outlet.

Because the probability of good news is greater and the probability of bad news is lower for the aligned outlet, the voter prefers the aligned outlet over the misaligned outlet. The voter's expected utility from reading the aligned outlet is

$$EU(\text{Read } L) = \alpha\{\rho(\pi_C + \pi_D) + (1 - \rho)(\pi_W + \pi_D)\} - \{\rho\pi_W + (1 - \rho)\pi_C\},$$

and her expected utility from reading the misaligned outlet is

$$EU(\text{Read } R) = \alpha\{\rho\pi_C + (1 - \rho)\pi_W\} - \{\rho(\pi_D + \pi_W) + (1 - \rho)(\pi_D + \pi_C)\},$$

and the difference between the two is

$$EU(\text{Read } L) - EU(\text{Read } R) = \pi_D(1 + \alpha),$$

which is always positive.

Proposition 2. The consumer in the confirmation-bias model reads the aligned outlet (as opposed to nothing) if and only if

$$\alpha > \frac{\rho\pi_W + (1 - \rho)\pi_C}{\rho(1 - \pi_W) + (1 - \rho)(1 - \pi_C)}.$$

The above expression shows that, all else equal, the voter is more likely to read L as opposed to nothing when ρ is higher or α is higher. If $\rho > \frac{1}{2}$, she is more like to read when π_W is lower or π_C is higher, and the opposite is the case when $\rho < \frac{1}{2}$. In other words, when the incumbent is expected to be high quality, the expected value of consuming news is greater when the news is more reliable, and when the incumbent is expected to be low quality, the expected value of consuming news is greater when the news is less reliable.

The comparative-static prediction about ρ may be particularly useful for generating empirical predictions. One source of variation in ρ comes from news reports themselves. Suppose the voter has the opportunity to consume multiple news reports. The confirmation-bias model predicts that hearing good news about the incumbent should, all else equal, make her more likely to consume more news from her preferred outlet, and hearing bad news should, all else equal, make her less likely to consume more news.

Information-Based model

Alternatively, suppose the voter derives no utility from reading good or bad news reports per se. Instead, she cares only about making the right electoral decision.

The voter receives a payoff of $\gamma > 0$ when she votes for an ideologically aligned candidate and a payoff of 1 when she votes for a candidate with high valence. Let V_S and I_S represent the valence and ideology of whichever candidate is ultimately selected by the voter. Then, her expected utility is given by

$$EU = \gamma 1\{I_S = L\} + \Pr(V_S = High).$$

Again, we assume she selects whichever option maximizes expected utility, and when she is indifferent between consuming news or not, she consumes no news.

Proposition 3. Let $\bar{\rho}_M$ be the expected valence after reading a positive report from media outlet M , and $\underline{\rho}_M$ be the expected valence after reading a negative report from outlet M . These posterior beliefs are functions of ρ , π_C , π_D , and π_W . It follows from our previous assumptions that

$$\underline{\rho}_L < \underline{\rho}_R < \rho < \bar{\rho}_L < \bar{\rho}_R.$$

Specifically, those posteriors are characterized by Bayes' Rule as

$$\begin{aligned} \underline{\rho}_L &= \frac{\rho\pi_W}{\rho\pi_W + (1-\rho)\pi_C}, \\ \underline{\rho}_R &= \frac{\rho(\pi_W + \pi_D)}{\rho(\pi_W + \pi_D) + (1-\rho)(\pi_C + \pi_D)}, \\ \bar{\rho}_L &= \frac{\rho(\pi_D + \pi_C)}{\rho(\pi_D + \pi_C) + (1-\rho)(\pi_D + \pi_W)}, \text{ and} \\ \bar{\rho}_R &= \frac{\rho\pi_C}{\rho\pi_C + (1-\rho)\pi_W}. \end{aligned}$$

Figure 1 helps to visualize Proposition 3 and the way media choice affects electoral choice and therefore voter utility. The horizontal axis represents the voter's beliefs about the valence of the incumbent, which start at ρ and shift upward or downward after consuming news. The vertical axis represents the voter's expected utility associated with different electoral choices. If the voter reelects the incumbent, her expected utility is shown by the dotted line. If she replaces the incumbent with the challenger, her expected utility is the flat, dashed line. And if she follows the optimal voting rule by only reelecting the incumbent when the expected utility of reelecting exceeds the expected utility of replacing, her expected utility is described by the solid upper contour.

The red and blue arrows and dots represent the potential changes in the voter's beliefs and subsequent voting decisions after reading R and L respectively. The thicker arrows and larger dots show the more likely scenario under each outlet, and the thinner

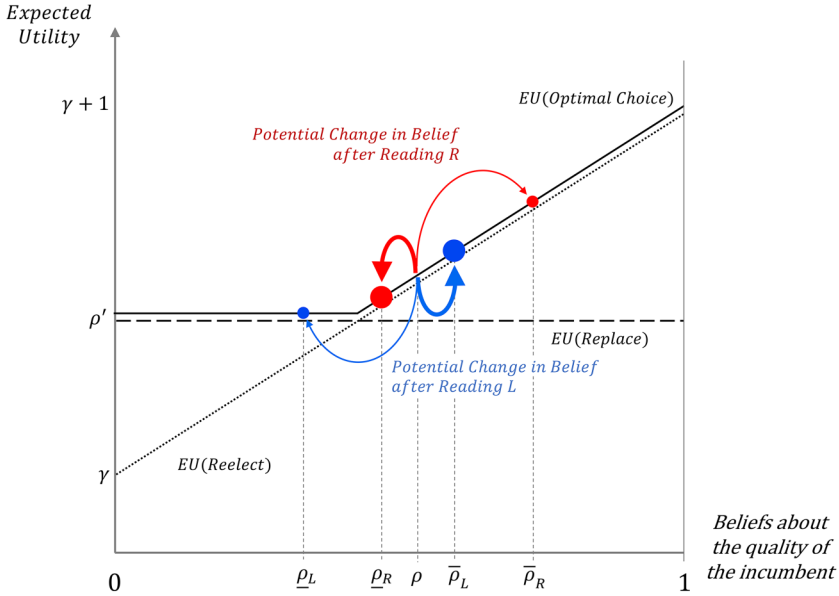


Figure 1. Effects of media signals on beliefs and vote choice.

arrows and smaller dots show the less likely scenario. In the figure, the effect of a positive versus negative signal from L on beliefs has the same magnitude as that of R . However, the signal from L is relevant for the voting decision, while the signal from R is not. In this scenario, the expected utility from consuming L is greater because in the unlikely scenario that L provides a negative signal about the incumbent, the voter is convinced to support the challenger and she is glad that she has obtained this information that enables her to make a better decision.

Let ρ^* be the level of expected incumbent valence at which a voter is indifferent between reelecting and replacing the incumbent in the absence of any new information, or $\rho^* = \rho' - \gamma$. We focus on the most interesting case in which $\rho^* \in (0, 1)$. This implies that $\gamma < \rho'$, which suggests that for sufficiently low beliefs about the incumbent’s valence, the voter prefers to replace the incumbent, and for sufficiently high beliefs, the voter prefers to reelect the incumbent. Figure 2 helps to visualize our subsequent theoretical results. The vertical lines in Figure 2 partition 6 different possible cases to consider. As the figure illustrates, we can think about moving from case to case by varying ρ , the prior belief about the valence of the incumbent, while holding all else equal. Of course, there is nothing special about ρ relative to other parameters in the model, and we could have moved across cases by varying those. For example, consider Case 6 where the voter will support the incumbent regardless of any potential signal she receives from the news. This case can arise because ρ is high, ρ' is low, π_C is small, π_W is large, γ is large, or some combination of the above.

Proposition 4. If $\bar{\rho}_R \leq \rho^*$ or $\rho^* \leq \underline{\rho}_L$ (cases 1 and 6 in Figure 2), the voter will consume no news.




Figure 2. Illustration of theoretical cases.

If no possible signal realization from the news would lead the voter to change her vote from what she was already inclined to do in the absence of new information, she has no incentive to consume news.

Proposition 5. If $\underline{\rho}_L < \rho^* < \rho$ (cases 4 and 5), the voter consumes news from her aligned news outlet.

If the voter is inclined to support the aligned incumbent in the absence of new information, but she could be persuaded to support the challenger, then the voter will want to consume news, and she will strictly prefer the aligned outlet over the misaligned outlet.

To prove Proposition 5, we break it into two cases, the first of which is easy to prove with no further math (as with Proposition 4). Suppose $\underline{\rho}_L < \rho^* < \underline{\rho}_R < \rho$ (case 5). In this case, a negative signal from the misaligned outlet would not be enough to persuade the voter to support the challenger. Only a negative signal from the aligned outlet would be pivotal for her voting decision. It's easy to see that the voter derives no benefit from reading the misaligned outlet, so she would never prefer that over reading no news, but she does derive utility from reading the aligned outlet because sometimes, it informs her that the aligned incumbent is so low quality that she is better off electing the challenger.

The harder case is when $\underline{\rho}_L < \underline{\rho}_R < \rho^* < \rho$ (case 4). In this scenario, a negative signal from either outlet would be enough to persuade her to support the challenger. So both

news outlets provide an instrumental benefit to the voter, and we have to assess which one provides more of a benefit. The voter's expected utility from each outlet can be understood by thinking through the likelihood and expected benefit associated with each of the possible cases—the incumbent can be high- or low-valence and the media outlet can report positively or negatively. If the incumbent is high-valence and the voter receives a correct, positive signal and reelects, then she receives a payoff of $1 + \gamma$. If the incumbent is low-valence and the voter receives an incorrect, positive signal and reelects, she receives a payoff of γ . And if the voter receives a negative signal (correct or incorrect) she will replace the incumbent and receive an expected payoff of ρ' .

In this case, the voter's expected utility from choosing the aligned news outlet is

$$\rho(\pi_D + \pi_C)(1 + \gamma) + \rho\pi_W\rho' + (1 - \rho)(\pi_D + \pi_W)\gamma + (1 - \rho)\pi_C\rho'.$$

Her expected utility from choosing the misaligned outlet is

$$\rho\pi_C(1 + \gamma) + \rho(\pi_D + \pi_W)\rho' + (1 - \rho)\pi_W\gamma + (1 - \rho)(\pi_D + \pi_C)\rho'.$$

The difference between these two expected utilities is

$$\pi_D(\rho + \gamma - \rho').$$

This means the voter will prefer the aligned media outlet when

$$\rho > \rho' - \gamma, \text{ or } \rho > \rho^*,$$

which we assumed above in the statement of Proposition 5.

Proposition 6. If $\rho < \rho^* < \bar{\rho}_R$ (cases 2 and 3), the voter consumes news from her misaligned outlet.

The proof of Proposition 6 is nearly identical to the proof of Proposition 5. If $\rho < \bar{\rho}_L < \rho^* < \bar{\rho}_R$ (case 2), the misaligned outlet is the only one that could potentially persuade her to switch from supporting the challenger to supporting the aligned incumbent, so this is the only outlet that provides an expected benefit. If $\rho < \rho^* < \bar{\rho}_L < \bar{\rho}_R$ (case 3), she will prefer the misaligned outlet over the aligned one so long as $\pi_D(\rho + \gamma - \rho') < 0$ or $\rho < \rho^*$, which is an assumption of Proposition 6.

In cases 1 and 6, the voter's beliefs are so extreme that no signal could persuade her to change her vote, so she consumes no news (Proposition 4). In cases 4 and 5, the voter is inclined to support the aligned incumbent but could be convinced to change her vote if she hears bad news from either outlet (case 4) or the aligned outlet (case 5). In both cases, she derives a higher expected utility by choosing the aligned outlet (Proposition 5). And in cases 2 and 3, the voter is inclined to support the challenger in the absence of new information but could be persuaded, and in these cases, she prefers the misaligned news outlet (Proposition 6).

The colored arrows in Figure 1 help to visualize some of our findings. In case 5, if the voter selects the misaligned outlet, she will never receive a signal that could change her vote, so her expected utility is the same as if she received no new information (red arrows). But if she selects the aligned outlet, most of the time she will hear good news, which will not change her vote, but some of the time, she will hear bad news, which will change her vote (blue arrows). Her expected utility from selecting the

aligned news outlet is higher than if she got no information because in the case where there is bad news about her incumbent, she changes her vote and increases her expected utility. In cases 1 and 6, none of the red or blue arrows reach a region where she would change her vote, so consuming news has no value.

To summarize the results of the information-based model, if the voter is inclined to support the ideologically aligned candidate in the absence of any new information—which, in practice, must be the case more often than not, she will prefer the ideologically aligned media outlet over the misaligned outlet. This is the same as the prediction from the confirmation-bias model. Furthermore, she will prefer to consume news from the aligned outlet so long as $\rho_L < \rho^*$. Since ρ_L is an increasing function of ρ and since ρ^* does not depend upon ρ , this means that, in the most likely range of possible beliefs, the voter is less likely to consume news as ρ increases. Notice that this is exactly opposite the prediction of the confirmation-bias model, providing an opportunity to empirically adjudicate between these two potential explanations of partisan media sorting.

Although we focused on the case of a left-leaning voter evaluating an ideologically aligned incumbent for expository purposes, notice that this is without loss of generality. We could have discussed a voter of either ideology and an incumbent of either ideology, and we would be left with the same predictions about media choice and news consumption. Whenever the voter is naturally inclined to support her party in the absence of new information, she will prefer the ideologically aligned media source, and she will be less likely to consume news when her beliefs are especially in favor of her party.

Although we have been discussing the case in which the voter learns only about the incumbent, this is also without loss of generality. The model could be modified to allow the voter to learn about only the challenger or both the incumbent and challenger, and the results would be unchanged. In a low-salience election, it might make sense to think about the voter learning about the quality of the incumbent by consuming coarse signals from the news and learning little about the challenger. But in a high-salience election, it might make more sense to think about the voter already knowing the quality of the incumbent and then gleaning information about the quality of the challenger by consuming coarse signals from the news. The same logic and results discussed above would apply in this model as well.

The existing model closest in spirit to our information-based model is Suen (2004). Our model applies the same logic and mechanism to the context of media choice and voting, and we set it up in a way that enables us to compare and contrast the predictions of the information-based and confirmation-bias models. Other related models have additional features or assumptions that are unnecessary or irrelevant in this particular setting. For example, in Calvert (1985), an advisor provides binary signals about two competing options, while in our information-based model, a media outlet provides one binary signal about the relative utility of two options, which more closely mirrors the case of a newspaper endorsing a political candidate. This means that the mechanisms driving the results differ between the two models. In Calvert's model, biased advisors often convey more information than unbiased ones, but in our information-based model, the aligned and misaligned news outlets convey the same amount of information, but one is more likely to convey decision-relevant information.

Empirical evidence distinguishing between confirmation bias and informed voting

Neither model presented above is intended to be an accurate and complete description of media choice. Surely, many individuals select their news outlets partly because they want to cast more informed votes, partly because they enjoy consuming news that confirms their predispositions, and for many other reasons outside the scope of this paper. If we combine the two models and allow voters to put weight on casting informed votes and on consuming desirable news, it would likely be a more realistic description of political news consumption. However, by developing two separate models and exploring their competing predictions, we hope to distinguish between these two different motivations for consuming political news and provide a framework for assessing the circumstances under which each is relevant for understanding media choice (see Ashworth et al., 2021 for a general discussion of how models can help us distinguish between competing explanations).

Although the contribution of this paper is primarily theoretical, this section briefly discusses empirical evidence in the literature and presents some new evidence to better assess the extent to which partisan media sorting is attributable to confirmation bias and/or the desire to cast more informed votes. No one piece of evidence is presented as providing a definitive answer, and we hope this discussion will motivate future empirical studies. Table 2 summarizes the empirical evidence discussed and presented in this section.

One straightforward prediction of our information-based model but not the confirmation-bias model is that unexpected news stories from an aligned source should influence vote choices. In the confirmation-bias model, voters already know which candidate they support and they simply consume news for its entertainment value. But in the information-based model, voters consume news to inform their decisions at the ballot box. In this model, realized signals from the news should influence voting behavior, and unexpected signals should be particularly influential.

Table 2. Summary of empirical evidence.

Finding	Source	Consistent with
U.S. newspaper endorsements affect vote choice, and unexpected endorsements have especially large effects.	Chiang and Knight (2011)	Informed voting
1997 Labour endorsements shifted voting behavior of UK citizens who read the affected newspapers.	Ladd and Lenz (2009) and Table 3	Informed voting
Labour supporters exposed to the unexpected endorsement reduced their news consumption.	Table 4	Informed voting
Republicans (Democrats) consumed more (less) news after Trump-Ukraine scandal broke.	Figure 3	Informed voting
Republicans increased news consumption after Comey letter.	Tyler, Grimmer, and Iyengar (2022)	Confirmation bias
Democrats increased news consumption after release of Access Hollywood tape.	Tyler, Grimmer, and Iyengar (2022)	Confirmation bias

Newspaper endorsements provide a nice opportunity to test this hypothesis because by endorsing one candidate over another, the news outlet is sending a binary signal to voters about which candidate it believes is better. Many previous studies have already investigated the effects of newspaper endorsements, and they generally find that these signals affect the decisions of many voters. In one particularly thorough study, Chiang and Knight (2011) collect data on the endorsements of hundreds of newspapers during the 2000 and 2004 U.S. presidential election campaigns, and they merge it with rolling surveys throughout the campaigns. Comparing readers of a particular paper before and after an endorsement, Chiang and Knight find that voters typically shift in the direction of the endorsed candidate, and this effect is particularly pronounced when the endorsement is unexpected—i.e., a conservative-leaning paper endorses a Democrat or a liberal-leaning paper endorses a Republican. This result is exactly in line with the predictions of the information-based model and could not be explained if partisans simply sort into like-minded papers because of confirmation bias.⁶

One concern with many of the studies in this literature is that newspaper readership is endogenous. In fact, the confirmation-bias model would predict that many partisans would stop reading the news altogether if their ideologically aligned newspaper endorses the other party's candidate. Ladd and Lenz (2009) address this concern by exploiting the British Election Panel Study (BEPS) in which the same respondents were interviewed over time. During the 1997 UK general election campaign, several conservative-leaning newspapers unexpectedly endorsed the Labour Party. Implementing a differences-in-differences design, Ladd and Lenz observe which newspapers respondents read in 1996, and they compare changes in vote choice for people who read the paper that would later make an unexpected endorsement to people who read a different paper. In Table 3, we replicate Ladd and Lenz's empirical strategy with some slight modifications explained below.

We limit our sample to respondents who reported reading a conservative-leaning newspaper in 1996.⁷ To avoid dropping anyone from the sample because of changes in their political views, we code variables for Labour support which take the value of 1 if a respondent reports supporting the Labour Party, 0 if they report supporting the Conservative Party, and 0.5 otherwise. We regress Labour support in 1997 on an indicator for whether a respondent's 1996 paper later endorsed the Labour Party, and we also include our measures of Labour support from 1996 and 1992. Column 1 of Table 3 shows the results of this regression, and we estimate that the unexpected newspaper endorsements shifted support for the Labour Party by .09 points on this 0-1 scale. At least 9% of individuals who read a conservative-leaning paper shifted their voting behavior in the direction of the Labour party as a result of their paper unexpectedly endorsing the Labour Party (more than 9 percent given partial shifting, e.g., going from Conservative support to abstaining).

Columns 2-4 of Table 3 show the results of the same analysis but separately restricting the sample to people who expressed support for the Labour Party, Conservative Party, or neither in 1996. We detect large effects for people who sided with the Conservative Party or neither party in 1996, but as expected, for people who were already planning to support the Labour Party, the unexpected good news about the Labour Party was largely inconsequential for their votes.

Table 3. Effects of unexpected newspaper endorsement on vote choice.

	(1)	(2)	(3)	(4)
	DV = Labour Support 1997			
	All Respondents	Conservative '96	Neither '96	Labour '96
Unexpected Endorsement	.087 (.025)	.094 (.062)	.111 (.036)	.023 (.044)
Labour Support 1996	.450 (.044)			
Labour Support 1992	.374 (.039)	.406 (.205)	.461 (.049)	.179 (.066)
Constant	.076 (.019)	.099 (.021)	.238 (.026)	.718 (.067)
Observations	603	138	314	151

Robust standard errors in parentheses

We can also use this setting and data to go beyond Ladd and Lenz's analyses and test additional implications of our models. The confirmation-bias model predicts that as beliefs about the quality of a party's candidate increase, news consumption should increase, while the same effect is ambiguous in the information-based model depending on where one's beliefs start—although we expect that the effect should typically be negative.⁸ Using the same BEPS data, we code the number of days per week that a respondent reports reading the newspaper. Focusing again on respondents who read a conservative-leaning paper in 1996, we regress our measure of newspaper readership in 1997 on an indicator for whether a respondent's paper unexpectedly endorsed the Labour Party in 1997, and we also include our measure of newspaper readership in 1996, 1995, 1994, and 1992 (all the years for which this question was asked).

Results are shown in Column 1 of Table 4. On average, we estimate that an unexpected endorsement of the Labour Party decreased newspaper readership by half a day per week. This might initially seem to be consistent with the confirmation-bias model since we'd expect conservatives who didn't want to hear bad news about their party to stop reading the paper as soon as they see the unexpected Labour endorsement. However, many people who already supported Labour in 1996 read these conservative-leaning papers. Perhaps they liked the non-political reporting of those papers—an important feature of real-world news consumption that is not present in either model.⁹ To more closely test the predictions of the models, we separately analyze people who supported the Conservative Party, the Labour Party, or neither in 1996.

Columns 2-4 of Table 3 show that to the extent that the unexpected Labour endorsements decreased newspaper readership, it was largely coming from people who already supported Labour in 1996. If these voters are motivated by confirmation bias, they should increase their news consumption after their paper endorses their preferred party. But instead, they go in the opposite direction, consistent with the information-based model, and this negative effect is statistically significant ($p = .016$). Once these voters see that their conservative-leaning paper endorses the Labour Party, this is a strong signal in

Table 4. Effects of unexpected newspaper endorsement on newspaper readership.

	(1)	(2)	(3)	(4)
	DV = Days Per Week 1997			
	All Respondents	Conservative '96	Neither '96	Labour '96
Unexpected Endorsement	-.511 (.201)	-.395 (.484)	-.432 (.302)	-.829 (.340)
Days Per Week 1996	.430 (.081)	.521 (.174)	.255 (.115)	.617 (.148)
Days Per Week 1995	.127 (.058)	-.055 (.089)	.233 (.084)	.042 (.108)
Days Per Week 1994	.113 (.050)	.166 (.118)	.081 (.066)	.187 (.094)
Days Per Week 1992	.136 (.051)	.255 (.111)	.124 (.075)	.088 (.086)
Constant	.749 (.401)	.335 (.786)	1.231 (.562)	.529 (.837)
Observations	603	138	314	151

Robust standard errors in parentheses

favor of their initial inclination, and because they are now more confident in their vote choice, there is less benefit to reading more political news. We also estimate that respondents who supported the Conservative Party in 1996 decreased their news consumption, potentially consistent with the confirmation-bias model, but this estimate is notably smaller than that for the Labour supporters, and it is not statistically distinguishable from zero. When we pool Conservative and Labour supporters and run an interactive regression, we cannot statistically reject the possibility that the effects are the same for the two groups.

Another opportunity to test whether beliefs about the relative quality of one's party's candidate influences news consumption arose with the Trump-Ukraine scandal. On September 26, 2019, news broke that a government official had filed a whistleblower complaint alleging that President Trump had attempted to arrange an illegal quid pro quo with the government of Ukraine. The event was significant enough that even Americans who do not regularly consume political news would have likely heard about it through friends and non-news outlets like late-night television (see Baum 2002 for a discussion of this kind of "soft news"). Furthermore, this event likely lowered the beliefs of many American voters about the quality, honesty, and leadership of President Trump and his administration.

This provides a nice opportunity to test the predictions of our models since the confirmation-bias model would predict that Democrats should consume more news (relative to Republicans) after this event, and the information-based model makes the opposite prediction. If partisans just care about reading good news about their party, Republicans should tune out once the Trump-Ukraine scandal broke, and Democrats should tune in. But if partisans are trying to think about casting the most informed votes, Republicans

should have had an extra incentive to tune in to figure out if they should continue to support Trump and the Republicans in 2020.

To assess the news consumption of partisans over this period, we utilize data from Nationscape, a large-scale survey of approximately 900 Americans per day during this period (Tausanovitch and Vavreck, 2020). Polls were conducted online through Lucid. One of the questions asked respondents whether they had read news about politics from a national newspaper such as the *New York Times*, *Wall Street Journal*, *USA Today*, or *Washington Post* in the past week. Although this is surely a coarse and imperfect measure, many Americans would plausibly turn to their national newspaper of choice if they want to learn more about the Trump-Ukraine scandal and decide for themselves whether the Trump administration behaved inappropriately or not.

Figure 3 shows the proportion of self-identified Democrats and Republicans reporting that they read a national newspaper over time starting 30 days before the news broke about the scandal through two weeks afterward. We have data on the exact time that each survey was taken, and we plot smoothed averages resulting from a kernel regression (Epanechnikov kernel with a bandwidth of 2 days). One limitation of this analysis is that we do not have panel data, so the pool of respondents is changing over time, and respondents may have changed their reported party identification as a result of current events (Montagnes et al., 2019). However, the proportions of respondents identifying as Democrats or Republicans do not change meaningfully over this period.

We see in Figure 3 that Democrats are more likely to report reading a national newspaper than Republicans, but there was not a meaningful change in the level of newspaper




Figure 3. Changes in news consumption after reports of whistleblower.

consumption or the difference between Democrats and Republicans after news broke about the Trump-Ukraine scandal. If anything, Democrats slightly decrease their reported consumption, and Republicans slightly increase their consumption, consistent with the information-based model, although the changes are substantively small.

Tyler et al. (2022) conduct additional tests of how the news consumption of partisans responds to good or bad news. Specifically, they examine the release of the Comey letter (good news for Republicans) and the release of the Access Hollywood tape (good news for Democrats) during the 2016 presidential election. They find, consistent with the confirmation-bias model, that partisans typically increase news consumption in response to good news, but inconsistent with confirmation bias, members of the other party do not decrease news consumption in response to bad news. One advantage of this study over our analysis of the Trump-Ukraine scandal is that their measure of news consumption is not self reported but is inferred from web-browsing activity.

We have attempted to test several predictions of our models using survey experiments fielded online through Amazon Mechanical Turk. Specifically, we randomly exposed self-reported partisans to receive positive or negative signals about their governor, and we randomly assigned whether these signals came from a conservative- or liberal-leaning outlet. Consistent with the information-based model and inconsistent with the confirmation-bias model, respondents do shift their beliefs about their governors and their future vote intentions in response to these media signals. However, other complications prevent us from more thoroughly testing the competing predictions of our models. For example, many respondents residing in New York did not perceive the *New York Times* to be relatively more liberal than the *Wall Street Journal*, and few respondents in any treatment condition were interested in seeking information about their governors, perhaps because Mechanical Turk workers often aim to complete their tasks as quickly as possible. In the interest of transparency, and to potentially benefit future researchers who might implement similar experiments in the future, we describe the experiments and results in detail in the Appendix.

Examining the evidence as a whole, we conclude that there is some evidence to support both models of partisan media sorting. The confirmation-bias model cannot explain why voters change their partisan support in response to unexpected newspaper endorsements, and it cannot explain why Labour supporters decreased their news consumption in response to good news about their party. But the information-based model does not explain why Democrats increased news consumption in response to the Access Hollywood tape or Republicans increased consumption in response to the Comey letter. To the extent that partisans sort into different media outlets, some of this is likely attributable to confirmation bias and a desire to hear good news about one's party, but some of this is also likely attributable to voters trying to obtain more relevant information to make better electoral decisions.

Extensions

Here, we discuss several extensions of our information-based model that may be relevant to various public and academic debates. First, what are the normative implications of biased and polarized media? To be sure, too much media polarization is bad for voter information and decision making. In our model, if $\pi_D = 1$ and therefore the media

outlets always disagree, the voter learns nothing from the media and cannot improve her vote choices. But if there is no media polarization, many voters obtain a signal from the media that is irrelevant for their electoral decisions. As has been demonstrated in other work, an individual can often make better decisions through an ideologically aligned advisor (e.g., Suen 2004), so biased media can allow individual voters to make electoral decisions that are better for them.

However, what matters for society is aggregate electoral selection, not whether each individual was able to cast votes more in line with their interests. Better voting decisions at the individual level need not imply better aggregate electoral outcomes, and vice versa (Ashworth and Fowler, 2020). If we extended our model to include many different voters with different preferences, what would matter most for voter welfare is the slant of the media outlet that is selected, in equilibrium, by the median voter. If the median voter is someone who is almost indifferent between the two parties on ideology and cares most about valence, then in the context of our model, they will prefer a media outlet that is not strongly biased in favor of either party, so polarized media choices could be bad for voter welfare in this scenario. But if there are many different media outlets with different slants, this could be good for voter welfare—allowing each voter to sort into whichever outlet provides information that is most relevant for their vote choice, leading the aggregate electoral outcomes to be similar to those in a world in which everyone is perfectly informed.

Next, many scholars are interested in the effects of biased media on vote choice, and several empirical studies have investigated the effects of media choices on vote choice. For example, DellaVigna and Kaplan (2007) examine the introduction of Fox News across cable markets and find that Fox News exposure benefited George W. Bush in the 2000 presidential election. Hopkins and Ladd (2014) examine heterogeneity across individuals and find that the effect appears for Republicans but not Democrats. Gerber et al. (2009) conduct a field experiment in which they randomly assigned households to receive a left-leaning newspaper, a right-leaning newspaper, or neither in the run-up to a gubernatorial election. Perhaps surprisingly, they find that both papers increase support for the Democratic candidate.

In most of these studies, the theoretical expectation is that consuming left-leaning news should make a voter more likely to support left-leaning candidates, and vice versa. But in our information-based model, the effects are subtler. Instead, the expected effect of media choice on vote choice varies across the different cases shown in Figure 2. In cases 1 and 6, there should be no effect. In cases 3 and 4, the effects should go in the intuitive direction. In these cases, the voter votes entirely according to the signal they receive from their chosen outlet, and the left-leaning outlet is more likely to provide good news about the left-leaning candidate. However, in cases 2 and 5, the effect of media choice on vote choice, on average, goes in the *counterintuitive* direction. For example, in case 5, the voter is inclined to support the aligned incumbent in the absence of any additional information. If she consumes news from the misaligned outlet, this will have no effect on her vote choice, and she will continue to support the aligned incumbent. But if she consumes news from the aligned outlet, she will sometimes hear bad news about the aligned incumbent that persuades her to vote for the other party. So in this case, reading the aligned outlet paradoxically makes the voter less likely to vote

in line with her typical party. She still supports her party most of the time, but she is more likely to vote for the other party when she is exposed to her like-minded outlet.

This means that the effects of media choice on vote choice are theoretically ambiguous, and this may help us to understand some of the null or surprising findings in this literature. For example, how did a conservative newspaper increase Democratic support in Gerber, Karlan, and Bergan's experiment? One possibility is that there were many voters initially leaning toward the Republican but open to being persuaded, and in this particular case, both newspapers provided good news about the Democratic candidate relative to the Republican. And if enough voters were initially in cases 2 or 5 as opposed to cases 3 or 4, we might expect that the right-leaning paper would, on average, benefit the Democratic candidate. Future studies on the effects of media choice may benefit from thinking about this subtlety, and future researchers may be able to use these insights to explicitly test for heterogeneity according to the initial beliefs or predispositions of the voters being studied.

Third, what does our model have to say about the effect of media signals on beliefs? Many political scientists argue that voters form their beliefs poorly and in biased ways. For example, scholars often point to the fact that partisans have different beliefs and don't react similarly to the same political events as evidence of partisan biases in belief formation (e.g., Bartels 2002). Of course, one reason that partisans can have different beliefs and respond to the same event in different ways is that they may have different values and prior beliefs. But our information-based model also shows that even if partisans have the same values and prior beliefs, they can still respond differently to the same political event because they have rationally sorted into different media outlets. In our model, if two voters start with the same prior belief but consume media from outlets with different thresholds, they will necessarily have different beliefs after consuming news.

For example, suppose the media outlets are providing signals about a left-leaning incumbent. If both outlets report good news for the incumbent, both voters will increase their beliefs, but the voter consuming right-leaning news will shift her beliefs more. If both outlets report bad news, both voters will decrease their beliefs, but the voter consuming left-leaning news will shift more. And if the outlets report different signals, the voters will respond to the same events in the world by shifting their beliefs in opposite directions. None of this is normatively troubling for democracy or voter competence. Voters have a rational incentive to sort into different media outlets according to their ideologies, and this will mean that those with different ideologies will likely have different beliefs and respond differently to the same events.

Although we assume in our models that media bias and polarization are exogenous, we discuss here how media polarization can arise endogenously in several different ways. As we've discussed, there is an optimal level of bias for each voter, so there is demand for biased media among the consumers. Furthermore, if the ideological orientation of a news outlet is far enough from a voter's orientation, she derives no benefit from that outlet because that news will never be consequential for her vote. Given this, it would likely be difficult to sustain an equilibrium without some media polarization. Suppose there is a continuum of political ideologies among the electorate and suppose that all media outlets are converged at the preferred threshold of the median voter. An outlet would likely have an incentive to deviate for two reasons. First, instead of receiving a portion

of the market share among voters around the median, an outlet could receive all of the market share among a group of consumers to the left (right) by shifting their threshold to the left (right). Second, by shifting to one direction, an outlet will induce new people to consume news that previously opted out. For these reasons, consumer demand for media bias would induce profit-maximizing outlets to endogenously polarize (see Chan and Suen, 2008 for an explicit model of this phenomenon).

Media polarization could also arise in our framework from a desire of the owners of media outlets to influence elections. If the owner of a media outlet has their own ideology, they may prefer to set their threshold at whatever point maximizes their own expected utility by sending positive signals about the incumbent whenever their expected valence is high enough such that the owner would like them to be reelected and negative signals otherwise. Alternatively, if the owner is a partisan who simply wants to maximize the probability that her preferred party is elected, it's not obvious what threshold they should select. As discussed above, if an outlet's threshold is too extreme, few people will choose that outlet because it doesn't provide enough information. Furthermore, left-leaning outlets can benefit right-leaning candidates, and vice versa, because there will be voters in cases 2 and 5 of Figure 2. Nevertheless, media polarization can arise endogenously in our framework from either the demand for biased news among consumers or from the desire of owners to influence elections.

A final, natural extension of our information-based model would allow the voter to consume news from multiple outlets. However, within our framework, if consuming news is costly, the voter would never select more than one outlet. This is because once they have already consumed news from their optimal outlet, news from other outlets will never be pivotal for their voting decision. This result would not necessarily hold in another model where the media outlets do not receive a common signal about the incumbent's valence. The fact that different media outlets may obtain different information may provide a rationale for politically interested consumers to seek multiple outlets. But if different outlets are reporting on the same information but with different slants, the voter has an incentive to focus on the one, most ideologically aligned outlet.

Conclusion

We have provided a theoretical explanation for partisan media sorting based on the informational value of news from an aligned vs. misaligned source. We have compared and contrasted the predictions of our explanation with the most common explanation that voters enjoy hearing good news about their party (and bad news about the other party). Our theoretical exercise demonstrates that partisan media sorting, in and of itself, is not evidence of confirmation bias or poor decision making on the part of voters because we would expect to observe partisan media sorting even if voters are perfectly rational and only care about making the best electoral decisions.

However, our information-based model of partisan media sorting does make different predictions than our confirmation-bias model about whether vote choices will be influenced by media signals and when partisans will consume more or less news. These competing predictions allow us to partly adjudicate between these two explanations and ask which mechanism is a more relevant explanation of partisan media sorting in the real

world. After reviewing the empirical evidence and contributing some new empirical tests, we find some evidence to support both the information-based and the confirmation-bias models. Although some voters surely do enjoy consuming like-minded news for noninstrumental reasons, much of the observed partisan media sorting is also likely attributable to the desire of voters to obtain news that will be most relevant for their voting decisions.

We have also discussed various extensions of our information-based model that are relevant to the normative implications of slanted media, conflicting results on the effects of media choice on vote choice, the extent to which partisans have different beliefs about political facts, and explanations for media bias. When partisan voters have limited attention or time but wish to cast informed votes, they will often choose to consume slanted news. Some of the implications are counterintuitive. For example, consuming aligned news can (in some cases) make voters more likely to vote for the other party, and rational and informed citizens will often respond quite differently to the same news event. Perhaps most importantly, slanted media are not necessarily bad for voter welfare, and partisan media sorting could be a sign that voters are selecting the news outlets that provide the most relevant information for their voting decisions.


Acknowledgements

We thank Scott Ashworth, Pamela Ban, Peter Buisseret, Dan Butler, Ethan Bueno de Mesquita, Wiola Dziuda, Arda Gitmez, Justin Grimmer, Tim Groseclose, Seth Hill, Will Howell, Federica Izzo, Gary Jacobson, Thad Kousser, Daniel Moskowitz, Konstantin Sonin, Chris Tausanovitch, Lynn Vavreck, Stephane Wolton, Adam Zelizer, and seminar participants at the University of Chicago and UCSD for helpful comments.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iD

Anthony Fowler  <https://orcid.org/0000-0001-6566-2556>

Notes

1. As we discuss when we present our formal models, this logic could also go in the opposite direction. If a voter is naturally inclined to support the other party in the absence of any new information, she will have an incentive to seek out the ideologically misaligned media source. However, if we think these situations are rare, then our information-based model will typically predict that partisans favor aligned media sources.
2. The assumption that media outlets coarsen the signal can be justified in a few different ways. First, media outlets have limited space and time to convey the news, so they cannot relay all the information they have. Second, media outlets often constrain themselves to convey binary or coarse signals by, for example, endorsing specific political candidates. Third, the consumers of news have limited time and attention, which further coarsens the information media outlets can convey. For instance, if consumers only have the time and interest to glance at

- headlines, it will be as if they are receiving coarse signals about the quality of their incumbent politicians.
3. If the media outlets or the signal distributions are asymmetrically biased, this complicates our models without yielding new insights. If one outlet is notably more biased than the other, it will provide little informational value to its consumers, so the more interesting case for studying partisan media sorting is one in which the two outlets are comparably biased.
 4. This assumption is equivalent to having a very small cost of consuming news. We could explicitly include a cost in our models and none of the subsequent results would change.
 5. The dictionary of the American Psychological Association defines confirmation bias as “the tendency to gather evidence that confirms preexisting expectations, . . .” (see dictionary.apa.org/confirmation-bias). Our model attempts to capture this notion by assuming that individuals enjoy confirmatory evidence and wish to avoid conflicting evidence. Some uses of the term *confirmation bias* also include the possibility that individuals dismiss conflicting evidence when they are confronted with it. This latter notion is not explicitly built into our model, but a model with this kind of motivated reasoning would also imply that consuming news has little effect on beliefs or vote choices.
 6. Chiang and Knight (2011) motivate this analysis with a model of media effects which is similar, in spirit, to our information-based model. However, they do not endogenize the voter’s choice of media outlet or decision to consume media; they take the voter’s newspaper and decision to consume news as given.
 7. Because the treatment of interest arises for people who read a conservative-leaning paper that unexpectedly endorsed the Labour Party, we believe we can make a more compelling apples-to-apples comparison by focusing only on people who initially read a conservative-leaning paper.
 8. One way to think about these empirical tests is that we have already played the game once where consumers decided whether to consume news and from which outlet to consume it. After people have received new information, we can play the game again and ask whether these individuals would like to consume more news now that their beliefs about the quality of the incumbent have potentially shifted. The posterior beliefs from the previous game become the prior beliefs in the new game. We have empirical predictions about how news consumption should vary with the prior beliefs, and the differential signals from the previous game provide exogenous variation in those prior beliefs.
 9. As expected, there is partisan media sorting in this sample, with Labour supporters more likely to read a left-leaning paper and Conservative supporters more likely to read a right-leaning paper in 1996. But because there are so many more Labour supporters in the sample in 1996, there are roughly as many Labour supporters as Conservative supporters reading right-leaning papers.

References

- Ashworth S, Berry CR and de Mesquita EB (2021) *Theory and Credibility: Integrating Theoretical and Empirical Social Science*. Princeton: Princeton University Press.
- Ashworth S and Fowler A (2020) Electorates vs. Voters. *Journal of Political Institutions and Political Economy* 1(3): 477–505.
- Bartels LM (2002) Beyond the running tally: Partisan bias in political perceptions. *Political Behavior* 24: 117–150.
- Baum MA (2002) Sex, lies, and war: How soft news brings foreign policy to the inattentive public. *American Political Science Review* 96(1): 91–109.
- Berinsky AJ, Huber GA and Lenz GS (2012) Evaluating online labor markets for experimental research: Amazon.com’s Mechanical Turk. *Political Analysis* 20(3): 351–368.

- Berinsky AJ, Margolis MF and Sances MW (2014) Separating the shirkers from the workers? Making sure respondents pay attention on self-administered surveys. *American Journal of Political Science* 58(3): 739–753.
- Calvert RL (1985) The value of biased information: A rational choice model of political advice. *Journal of Politics* 47(2): 530–555.
- Chan J and Suen W (2008) A spatial theory of news consumption and electoral competition. *Review of Economic Studies* 75(3): 699–728.
- Chiang C-F and Knight B (2011) Media bias and influence: Evidence from newspaper endorsements. *Review of Economic Studies* 78(3): 795–820.
- Crawford VP and Sobel J (1982) Strategic information transmission. *Econometrica* 50(6): 1431–1451.
- DellaVigna S and Kaplan E (2007) The fox news effect: Media bias and voting. *Quarterly Journal of Economics* 122(3): 1187–1234.
- Flaxman S, Goel S and Rao JM (2016) Filter bubbles, echo chambers, and online news consumption. *Public Opinion Quarterly* 80(1): 298–320.
- Gentzkow M and Shapiro JM (2006) Media bias and reputation. *Journal of Political Economy* 114(2): 280–316.
- Gentzkow M and Shapiro JM (2010) What drives media slant? Evidence from U.S. Daily newspapers. *Econometrica* 78(1): 35–71.
- Gentzkow M and Shapiro JM (2011) Ideological segregation online and offline. *Quarterly Journal of Economics* 126(1): 1799–1839.
- Gerber AS, Karlan D and Bergan D (2009) Does the media matter? A field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics* 1(2): 35–52.
- Green DP, Zelizer A and Kirby D (2018) Publicizing scandal: Results from five field experiments. *Quarterly Journal of Political Science* 13(3): 237–261.
- Grimes DR (2017) Echo Chambers Are Dangerous—We Must Try to Break Free of Our Online Bubbles. *The Guardian*. December 4: 2017. <bit.ly/2ySQ0kQ>.
- Groseclose T and Milyo J (2005) A measure of media bias. *Quarterly Journal of Economics* 120(4): 1191–1237.
- Guess AM (2021) (Almost) everything in moderation: New evidence on Americans' online media diets. *American Journal of Political Science*.
- Hopkins DJ and Ladd JM (2014) The consequences of broader media choice: Evidence from the expansion of fox news. *Quarterly Journal of Political Science* 9(1): 115–135.
- Iyengar S and Hahn KS (2009) Red media, blue media: Evidence of ideological selectivity in media use. *Journal of Communication* 59(1): 19–39.
- Kayser MA and Peress M (2021) Does the media cover the economy accurately? An analysis of sixteen developed democracies. *Quarterly Journal of Political Science* 16(1): 1–33.
- Kennedy R, Clifford S, Burleigh T, et al. (2020) The shape of and solutions to the MTurk quality crisis. *Political Science Research Methods* 8(4): 614–629.
- Knobloch-Westerwick S, Johnson BK and Westerwick A (2015) Confirmation bias in online searches: Impacts of selective exposure before an election on political attitude strength and shifts. *Journal of Computer-Mediated Communication* 20(2): 171–187.
- Knobloch-Westerwick S and Kleinman SB (2012) Preelection selective exposure: Confirmation bias versus informational utility. *Communication Research* 39(2): 170–193.
- Ladd JM and Lenz GS (2009) Exploiting a rare communication shift to document the persuasive power of the news media. *American Journal of Political Science* 53(2): 394–410.
- Larcinese V, Puglisi R and Snyder JM Jr. (2011). Partisan bias in economic news: Evidence on the agenda-setting behavior of U.S. Newspapers. *Journal of Public Economics* 95(9–10):1178–1189.

- Levendusky M (2013) *How Partisan Media Polarize America*. Chicago: University of Chicago Press.
- Montagnes BP, Peskowitz Z and McCrain J (2019) Bounding partisan approval rates under endogenous partisanship: Why high presidential partisan approval may not be what it seems. *Journal of Politics* 81(1):321–326.
- Patty JW (2009) The politics of biased information. *Journal of Politics* 71(2): 385–397.
- Prior M (2007) *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. Cambridge: Cambridge University Press.
- Puglisi R and Snyder MJr (2015) The balanced U.S. Press. *Journal of the European Economic Association* 13(2):240–264.
- Stroud NJ (2008) Media use and political predispositions: Revisiting the concept of selective exposure. *Political Behavior* 30(3): 341–366.
- Suen W (2004) The self-perpetuation of biased beliefs. *The Economic Journal* 114(495): 377–396.
- Tausanovitch C and Vavreck L (2020) *Democracy Fund + UCLA Nationscape Project (Version 0) [Codebook and Dataset accessed from PIs on October 19, 2019]*. Los Angeles, CA: Democracy Fund and UCLA.
- Tyler M, Grimmer J and Iyengar S (2022) Partisan enclaves and information bazaars: Mapping selective exposure to online news. *Journal of Politics* 84(2): 1057–1073.
- Wolton S (2019) Are biased media bad for democracy? *American Journal of Political Science* 63(3): 548–562.

Appendix

In July and August of 2020, we fielded two field experiments online through Amazon Mechanical Turk or MTurk (see Berinsky et al., 2012). In one study, we recruited respondents from the state of Georgia, and in the other, we recruited respondents from the state of New York. In both cases, respondents were at least 18 years old, they self-identified as either Democrats or Republicans (i.e., Independents and others were excluded), we utilized IP tracking to verify that they indeed resided in either Georgia or New York (Kennedy et al., 2020), and we utilized a screener question to remove inattentive respondents who didn't read the question (Berinsky et al., 2014).

In both experiments, respondents were shown real news articles that focused on their governor's handling of the COVID-19 pandemic. In each case, we selected one article that portrayed the governor in a positive light and another that portrayed the governor in a negative light. For the Georgia experiment, the positive article discussed how the reopening of Georgia had not been as bad as many critics feared, and the negative article focused on the ways in which the governor's office had manipulated data to make the pandemic seem better than it was. For the New York experiment, the positive article praised the governor's leadership, and the negative article criticized the governor for requiring nursing homes to accept patients with COVID-19. For each real article, we created two different versions that could be shown to respondents. The content of the article was identical between the two versions, but the media outlet was altered. We closely mirrored the appearance of the actual media outlets such that a respondent would perceive the article as being taken directly from a news website. In the Georgia experiment, the articles were presented to be from either Fox News or MSNBC, and in the New York experiment, the articles were presented to be from either the *Wall Street Journal* or the *New York Times*. At the end of the

experiment, respondents were debriefed and shown the original article and the correct source. Both experiments were approved by our university's Institutional Review Board.

All respondents were shown a news article, and they were randomly assigned to one of four possible treatment conditions. The article could have provided a positive or negative signal about the governor's handling of the pandemic, and the article could have been from a liberal or conservative outlet. Before reading the article, respondents were asked to state their partisan identification (and only Democrats and Republicans were able to continue the study), and they were asked to evaluate the performance of their governor on a scale from 0 to 10. After reading the article, respondents were asked if they want to revise their evaluation of the performance of their governor, they were asked about their intention to vote for their governor if he sought reelection, and they were asked whether they would like to be provided with more information about their governor at the end of the survey.

Our information-based model makes predictions about how the different treatments should affect beliefs about one's governor. For example, for the Democratic governor of New York, a positive article from the *Wall Street Journal* should have a greater positive effect than a positive article from the *New York Times*. The model also makes predictions about how these media signals should affect future news consumption. For example, exposure to negative news about a Democratic governor should make Democrats more likely to seek more information about the governor and Republicans less likely to do so.

Unfortunately, we believe that several practical limitations prevented these experiments from offering a clean test of our theoretical predictions. One challenge is that it was difficult to recruit a significant number of Republicans through MTurk. Another challenge is that MTurk workers appear to complete their tasks as quickly as possible, so few respondents requested or sought additional information about their governor. In the New York experiments, many respondents did not perceive a meaningful ideological difference between the *New York Times* and *Wall Street Journal* (this doesn't necessarily mean the readers of these outlets are not aware of the ideological slants). Unfortunately, we only asked respondents about their perceptions of the ideology of the outlets after the experiment (we didn't want to prime them or inform them about the purpose of the experiment before we measured our outcome variables), and the treatments influenced these perceptions, so we cannot condition on them. Furthermore, in the Georgia experiment, respondents were not significantly influenced by any of the treatments, perhaps because they already had strong views about their governor and were already informed about the stories.

In the interest of transparency, we show the results of our experiments below. Table A1 tests how the various treatments affected respondents' beliefs about the performance of their governor and their future vote intentions. In the table, we classify the treatments as positive or negative (with respect to the information provided about the performance of the governor) and as ideologically aligned or misaligned with the governor. In Georgia, Fox News is classified as aligned with the Republican governor, and in New York, the *New York Times* is classified as aligned with the Democratic governor. We regress each outcome variable on indicators for our treatments. Theoretically, the negative story from that aligned source should be the worst for the governor, so we make that

the omitted category, and the three other coefficients of interest can be interpreted as the effect of another treatment relative to that omitted category. To improve precision, we also control for each respondent’s party and their pre-treatment belief about the performance of their governor.

In Columns 1 and 2 of Table A1, we see that our treatments did not significantly shift the views of respondents in Georgia about their governor. This could be because respondents in Georgia already held strong views about their governor. Governor Kemp was not particularly popular among our respondents, with the average rating and vote intention (scores ranging from 0 to 10) just above 3. This could also be because respondents in Georgia were already familiar with the stories we showed them. Another relevant factor is that our respondents may not have perceived one or both of the media outlets as especially credible—a real-world feature of news consumption that is outside the scope of our model. In Column 2, we see that seeing the positive vs. negative story increases vote intentions for Governor Kemp by 0.40 points if the source was MSNBC, but only 0.13 points if the source was Fox News. So many of the respondents (most of whom are Democrats), may have not put much stock in the stories if Fox News was the source—even if Fox News presented a negative story about a Republican governor.

In Columns 3 and 4 of Table A1, we see that the treatment did have a meaningful effect on the views of New York respondents about Governor Cuomo. The positive articles increase assessments of Cuomo relative to the negative articles. However, the effect of the article doesn’t appear to depend on the source. For example, a positive article from the *New York Times* increases beliefs about Governor Cuomo just as much if not more so than a positive article from the *Wall Street Journal*. This result is inconsistent with the predictions of our information-based model, but further investigation offers one explanation.

Table A1. Effect of news treatments on support for governor.

	(1)	(2)	(3)	(4)
	Georgia		New York	
	Rating	Vote Intention	Rating	Vote Intention
Misaligned/Negative	-0.23 (0.18)	-0.12 (0.27)	-0.07 (0.18)	-0.04 (0.25)
Aligned/Positive	-0.16 (0.18)	0.13 (0.27)	1.41** (0.18)	0.97** (0.25)
Misaligned/Positive	-0.11 (0.23)	0.28 (0.27)	1.34** (0.18)	0.91** (0.25)
Pre-Treatment Rating	0.84** (0.04)	0.98** (0.04)	0.86** (0.03)	0.96** (0.04)
Democrat	-0.58** (0.22)	-1.52** (0.26)	0.25 (0.16)	1.21** (0.22)
Constant	0.94** (0.29)	0.39 (0.34)	-0.08 (0.21)	-1.24** (0.29)
Mean DV	3.57	3.19	6.19	6.10
Observations	192	192	384	384

Robust standard errors in parentheses; ** p < .01, * p < .05.

At the end of the survey, we asked respondents to rate the ideological/partisan leanings of the two outlets utilized in the experiment on a scale from 0 (most liberal) to 10 (most conservative). While 73 percent of our Georgia respondents perceived that Fox News is more conservative than MSNBC, only 56 percent of our New York respondents perceived that the *Wall Street Journal* is more conservative than the *New York Times*, meaning our sample of respondents may lack the knowledge of these outlets required for our prediction to hold. This does not mean that the typical consumer of these outlets is not aware of their ideological biases, but it does mean that MTurk may not be an appropriate sample for testing the implications of our model.

Table A2 tests the next set of predictions of our information-based model regarding which individuals should be more likely to seek additional information. We asked all respondents if they would like to be directed to more information about their governor at the end of the survey, and for those that requested the information, we measured which ones clicked the link that we provided. Unfortunately, perhaps because MTurk workers try to complete their tasks quickly and move on to their next tasks, few respondents requested more information, and even fewer clicked the link, limiting the ability of our experiment to measure effects on information seeking. Furthermore, our predictions are about the extent to which good or bad news differentially affects Democrats vs. Republicans, so the underrepresentation of Republicans in our survey further limits the precision of our estimates.

We regress each measure of information seeking on an indicator for whether each respondent received good or bad news about the governor, an indicator for whether each respondent is from the same party as the governor, and the interaction of the two. The prediction of our information-based model is that this interactive coefficient is negative. Good news about the aligned candidate should make one less likely to seek more information, while good news about the misaligned candidate should make one more likely to seek more information. None of our estimated interactive coefficients are statistically distinguishable from zero, although for reasons discussed above, the standard errors are typically large relative to the share of respondents seeking more information.

Table A2. Effect of positive news on information seeking.

	(1)	(2)	(3)	(4)
	Georgia		New York	
	Requested	Clicked	Requested	Clicked
Positive News	.044 (.059)	-.004 (.042)	-.074 (.079)	-.063* (.031)
Same Party	.093 (.070)	.015 (.049)	-.005 (.068)	.025 (.039)
Positive News*Same Party	-.046 (.104)	-.039 (.062)	.057 (.095)	.044 (.045)
Constant	.089* (.038)	.054 (.030)	.270** (.056)	.063* (.031)
Mean DV	.141	.050	.251	.066
Observations	199	199	395	395

Robust standard errors in parentheses; ** $p < .01$, * $p < .05$.