

THE NONRANDOM WALK OF KNOWLEDGE

BY JANE R. BAMBAUER, SAURA MASCONALE AND SIMONE M. SEPE

Abstract: A person's epistemic goals sometimes clash with pragmatic ones. At times, rational agents will degrade the quality of their epistemic process in order to satisfy a goal that is knowledge-independent (for example, to gain status or at least keep the peace with friends.) This is particularly so when the epistemic quest concerns an abstract political or economic theory, where evidence is likely to be softer and open to interpretation. Before wide-scale adoption of the Internet, people sought out or stumbled upon evidence related to a proposition in a more random way. And it was difficult to aggregate the evidence of friends and other similar people to the exclusion of others, even if one had wanted to. Today, by contrast, the searchable Internet allows people to simultaneously pursue social and epistemic goals.

This essay shows that the selection effect caused by a merging of social and epistemic activities will cause both polarization in beliefs and devaluation of expert testimony. This will occur even if agents are rational Bayesians and have moderate credences before talking to their peers. What appears to be rampant dogmatism could be just as well explained by the nonrandom walk in evidence-gathering. This explanation better matches the empirical evidence on how people behave on social media platforms. It also helps clarify why media outlets (not just the Internet platforms) might have their own pragmatic reasons to compromise their epistemic goals in today's competitive and polarized information market. Yet, it also makes policy intervention much more difficult, since we are unlikely to neatly separate individuals' epistemic goals from their social ones.

KEY WORDS: Bayesian updating, collective epistemology, Internet, polarization, pragmatic reasons, selection effects

The Internet vastly expanded our ability to access and provide information. If everything worked as expected, people should converge on the truth faster than ever. Yet this is not what we see, particularly in the context of political beliefs and economic theories. Instead, the Internet, with its myriad echo chambers and “filter bubbles,” labors under the heavy criticism that it has caused people to become more polarized in their beliefs. Even putting aside fabricated evidence and blatant lies, Internet users can wind up entrenched in opposing camps of belief because they are exposed to different facts in their targeted news feeds as well as different content recommendations. Later, any common facts they receive are interpreted with starkly different prior plausibilities based on the facts that each person received, or sought out, earlier in the process.

Before the Internet, people received evidence in a more haphazard way. Their epistemic process was more separate from their social lives such that social, political, and economic information were an imperfectly random walk. While the flow of evidence was surely influenced by friends and close

doi:10.1017/S0265052521000145

© 2021 Social Philosophy & Policy Foundation. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution, and reproduction in any medium, provided the original work is properly cited. Printed in the USA.

colleagues, we had limited ability to grow our social circles, and the costs of doing so were significant.¹ Today, by contrast, social and epistemic goals are pursued at the same time on social media. The abundance of information in the digital era has reduced the opportunity costs of looking for additional evidence, but at the same time, it has also reduced the costs of social selection. It has become so easy to simultaneously strengthen bonds of friendship and receive information that is consistent with our priors, that randomness in evidence is now the more costly course. The result is exposure to less, rather than more, evidence, and increasingly polarized beliefs instead of convergence toward accuracy.

While the empirical evidence for Internet polarization continues to mount, the conventional explanations for it are flawed, and therefore lend themselves to flawed policies. Most treat individual Internet users as passive agents who are cognitively limited in their pursuit of truth. By contrast, we hypothesize that the emergent camps of entrenched beliefs are not (necessarily) caused by any failing in human rationality or by the schemes of a manipulative corporation; they are caused by the influence our individual pursuits of friendship and camaraderie have over the course of gathering evidence. Our thesis is most similar to one of the hypotheses offered by Cass Sunstein, who has long suggested that extremism can result from the reasonable inferences drawn from selective evidence.² But while Sunstein sees groups and niche communities as the key sources of imbalanced information,³ we show that each individual, with their unique set of friends and personal priorities, will experience epistemic distortion even if they avoid groupthink situations.

We begin by briefly modeling the behavior of rational agents who are motivated to understand the truth. These agents observe the world directly, and they also learn from each other's testimony. By hearing about what others have observed, the agents can update their own beliefs, to a greater or lesser degree, without having to directly observe the relevant facts themselves. The risks of being intentionally or unintentionally deceived by others' assertions are real, but outweighed by the saved costs of personally investigating every important fact.

The Internet has exponentially increased our ability to access testimony from a wide range of sources. But it has also increased the likelihood that selection effects will taint the epistemic quest for truth. The Internet's capacity to connect people who share beliefs, and to do so at virtually no cost,

¹ R. I. M. Dunbar, "Coevolution of Neocortical Size, Group Size and Language in Humans," *Behavioral and Brain Sciences* 16, no. 4 (1993): 681–735; W.-X. Zhou et al., "Discrete Hierarchical Organization of Social Group Sizes," *Proceedings: Biological Sciences* 272 (2005): 439–44 (producing evidence that our ability to process and synthesize information on social relationships is limited by cognitive constraints).

² Cass R. Sunstein, *Going to Extremes: How Like Minds Unite and Divide* (Oxford: Oxford University Press, 2009), 22–24.

³ Cass R. Sunstein, "The Law of Group Polarization," *Journal of Political Philosophy* 175 (2002): 177; Cass R. Sunstein and Reid Hastie, *Wiser: Getting Beyond Groupthink to Make Groups Smarter* (Cambridge, MA: Harvard University Press, 2014), 44–45.

distorts the selection of testimony that we encounter on the searchable web. Abundant speech wreaks havoc on the epistemic pursuit. In this sense, our essay is in line with the observations of Richard Sorabji about the dangers of social media as a source of information, and one may expect us to share his disappointment with Facebook's conduct and fecklessness.

Yet there are some needling details that suggest Facebook and other social media companies are receiving undue credit and blame for the state of modern discourse. First, our theory of human behavior assumes that Internet users permit a selection effect in the information that is presented to them. They are either ignorant about the selection effect—an implausible proposition at this point—or they do not care enough about their epistemic goals to avoid or correct for selection. There must be pragmatic goals—for example, maintaining social ties—that can rationally interfere with the quest to improve knowledge. Second, competition among content creators drives *all* media outlets (not just the Internet giants) to compromise their epistemic goals in order to keep the interest of an increasingly polarized audience. Each of these complicates the standard story about Internet echo chambers and shrinks the set of efficacious policy responses.

I. TESTIMONY AND KNOWLEDGE BEFORE FACEBOOK

Humans learn useful things by communicating with one another without having to directly experiment and learn from the world. When a speaker makes an assertion with the intent that it be accepted as true, listeners will use that testimony to adjust their understanding of the world. Although the speaker could be lying, the speaker's reputation will suffer if listeners discover the deceit. Unintentional mistruths can also occur, but the speaker's reputation has some disciplining effect on this too. The reputational sanctions are particularly relevant for assertions based on "harder" information—that is, factual claims that can be independently verified by listeners without a good deal of interpretation.

Thus, even though a listener may continue to have doubt and uncertainty about a proposition, her probabilistic predictions about its truth are refined and improved with the help of testimony. This ability to learn from the testimonial reports of others is one of our species' superpowers as it relieves us from having to start every epistemic quest with only the raw materials of our firsthand observations.

In principle, having access to more testimony should bring epistemic benefits by increasing the evidence available to individuals and hence producing more accurate beliefs. In fact, we show that advances in communications technology that increase our access to testimony also increase the likelihood of selection effects, causing beliefs to become more radicalized and less true. To understand why this is so, let's first explore the relationship between testimony and knowledge in the pre-Internet era using an illustration.

Suppose Simone is a rational alien from planet Bayesianus.⁴ He, like everyone on his planet, is an epistemic agent. He is motivated exclusively by a desire to ascertain the truth and he is able to think in an ordered, logical way. His beliefs about the truth of a proposition at any given time can be expressed as a credence between 0 (false) and 1 (true). His credences are updated according to the axioms of probability as he acquires new information over time.

Simone is not an expert of American politics, but has become focused on the question φ : “Is the president of the United States a bad one, who makes worse-than-average policy decisions?” Being completely ignorant at the beginning, Simone starts with a prior of $cr(\varphi = T) = \frac{1}{2}$.⁵ He then teleports to Earth to investigate the answer.

Shortly after arrival, let’s assume that Simone receives a piece of evidence about the president’s performance. For clarity, we provide a highly stylized and constrained example, though the ideas generalize. Suppose Simone can receive one of three possible signals, a, b , or c , by chatting with somebody on the street (a), reading an op-ed in the first newspaper he finds (b), or casually listening to talk radio (c). Each of these pieces of evidence is a signal that Simone can use to update his credence, where $s \in \{a, b, c\}$.

However, Simone also knows that the evidence in a, b, c is merely suggestive on its own and cannot alone give him confidence. This is because the likelihood distribution functions of receiving these signals under conditions where φ is true or φ is false are not very informative. Simone knows, for example, that an expression of disappointment about the president from a single nonexpert’s opinion on talk radio is quite likely to occur whether the president is doing a good job or not. Moreover, the information contained in these signals is likely to be softer in nature—the product of someone else’s interpretation rather than hard data that Simone can independently verify. Further, even if Simone could independently verify the harder facts that testifiers used as the basis for their conclusions, he knows he lacks the expertise to independently process and interpret this evidence.

Nevertheless, Simone is at least able to update his initial credence to $cr(\varphi = T|s)$, where this credence will depend on the likelihood distribution functions of receiving each signal under conditions where φ is true or φ is false and in accordance with the Bayes rule.⁶ Suppose that, depending on the signal he receives, Simone will update his prior as follow:

⁴ The Bayesian approach endorses strict rationality conditions, requiring individuals to form and update their credences in accordance with objective probability rules that are specified by the model. Roger Myerson, *Game Theory: Analysis of Conflict* (Cambridge, MA: Harvard University Press, 1991).

⁵ This means that Simone starts his inquiry about φ with a cognitive clean slate.

⁶ The Bayes equation here reads: $\frac{cr(s|\varphi=T)}{cr(s|\varphi=T)+cr(s|\varphi=F)}$. Note, however, that this formula is simplified, as the complete version of the Bayes rule would require multiplication of the numerator and each term of the denominator by the common prior (i.e., here, $\frac{1}{2}$). However, the prior can be omitted, as it cancels out.

$cr(\varphi = T|a) = 0.391$, $cr(\varphi = T|b) = 0.483$, and $cr(\varphi = T|c) = 0.605$.⁷ The precision of these calculations and the likelihood distribution functions that produce them, like the hypothetical as a whole, is far-fetched. But if our illustration works under the demanding and constrained assumptions of objective Bayesianism (which assumes that likelihood distribution functions are objective and known), it will also work *a fortiori* under more relaxed assumptions. This includes assuming that (i) agents are Bayesians but have only a subjective representation of likelihood distribution functions (that is, can conceive of, but do not know for sure, certain states of the world, together with their likelihood, that might apply to a certain object of investigation),⁸ and (ii) individuals are unable to engage in formal complex computations but can efficiently use heuristics in its place (for example, one can know how to throw a ball without necessarily knowing the laws of physics governing this action).

Now, imagine Simone receives signal c , by listening to the radio, so that he updates his credence to $cr(\varphi = T|c) = 0.605$ —a moderate belief that the president is bad. Because the Internet has not yet become popularized, Simone has limited options for what to do next. Media markets have high entry barriers because printing presses and journalism staff are expensive, and competition from television networks has captured the demand for smaller, local papers.⁹ And access to broadcast media is restricted to some extent by limited usable frequency bandwidths and resulting licensing schemes. Like most other people, Simone can only turn to a handful of “prestige” national newspapers and network televisions, and maybe one local newspaper, to look for further evidence.¹⁰

While pondering his options, Simone turns on the television in his hotel room and watches an interview with Dave, a respected political philosophy professor. Dave can help Simone learn whether φ is true because Dave, too, is an immigrant from planet Bayesianus who initially came to Earth with a clean slate of knowledge (namely, $cr(\varphi = T) = 1/2$). But unlike Simone, Dave has spent years studying American politics. Hence, the evidence available to Dave is more informative than the evidence independently available to Simone, both because this evidence will tend to be harder in nature and because Dave has the skills to verify the veracity of the information he receives. Let’s then assume that the signals Dave can receive about φ , the quality of the president’s policy decisions, are qualified signals, $q \in \{d, e, f\}$, where the likelihood distribution functions of these signals is much more informative than those of the set of signal $s \in \{a, b, c\}$ about φ being true

⁷ These credences obtain under the Bayes rule from the following likelihood distribution functions conditional on φ being true, $cr(a|\varphi = T) = 0.25$, $cr(b|\varphi = T) = 0.29$, $cr(c|\varphi = T) = 0.46$, and conditional on φ being false, $cr(a|\varphi = F) = 0.39$, $cr(b|\varphi = F) = 0.31$, $cr(c|\varphi = F) = 0.30$.

⁸ Leonard J. Savage, *The Foundations of Statistics* (New York: Dover Publications, 1954), 8–9.

⁹ David Halberstam, *The Powers That Be* (Urbana, IL: University of Illinois Press, 2000).

¹⁰ Lee McIntyre, *Post Truth*, MIT Press Essential Knowledge Series (Cambridge, MA: MIT Press, 2018), 63–64.

or false.¹¹ Before doing the interview, Dave has received d . His updated credence on φ under the Bayes rule is $cr(\varphi = T|d) = 0.068$, strongly convincing Dave that the president's policy is in fact a good one. If Dave had received one of the other signals, his updated credences would have been $cr(\varphi = T|e) = 0.524$, and $cr(\varphi = T|f) = 0.912$.

What will happen to Simone's moderate belief that the president is bad after he listens to Dave, the expert with a high credence that the president is good?¹² By applying the Bayes rule, Simone's updated credence will be 0.101,¹³ meaning that Simone will also come to be convinced that the president is, in fact, good.

Simone wanders the streets a little while longer, picking up more signals, but because none have the same authority and epistemic influence as Dave's signal, Simone's updating never again veers very far from 0.101. He returns to his planet fairly confident that the president is making good policy decisions. His beliefs are informed by the well-honed testimony of Dave.

II. TESTIMONY AND KNOWLEDGE UNDER THE INFLUENCE OF FACEBOOK

Now, suppose we move Simone forward in time by a couple decades so that he visits Earth during the modern Internet era. His first testimony is the same as before—signal c that updates his prior to $cr(\varphi = T|c) = 0.605$.

This time, however, it is going to be much easier for Simone to access evidence about φ . The Internet and social networks can provide virtually infinite sources of testimony, while many of the entry barriers to the media and broadcast markets have been removed to promote increased competition. In this changed environment, the opportunity costs to look for further evidence have considerably declined. One would expect this to translate into access to more total evidence and hence better beliefs. Instead, it will translate into polarization. Where do things go wrong?

Let's go back to our illustration to try to answer this question. Imagine that, like a majority of adults in the United States these days,¹⁴ Simone decides to use social networks to get more evidence on φ . After his initial exposure in the real world to the signal c and before watching the interview

¹¹ The likelihood distribution functions conditional on φ being true are $cr(d|\varphi = T) = 0.06$, $cr(e|\varphi = T) = 0.11$, $cr(f|\varphi = T) = 0.83$, and conditional on φ being false, $cr(d|\varphi = F) = 0.82$, $cr(e|\varphi = F) = 0.10$, $cr(f|\varphi = F) = 0.08$.

¹² For simplicity, but without loss of generality, we assume that each individual knows the evidence distribution of the other individuals.

¹³ Here, the formula of the Bayes rule is: $\frac{cr(s|\varphi=T) \times cr(q|\varphi=T)}{[cr(s|\varphi=T) \times cr(q|\varphi=T)] + [cr(s|\varphi=F) \times cr(q|\varphi=F)]}$. Therefore, the conditional probability of φ being true, conditional on all the possible combinations of signals that Simone and Dave can receive, are: $cr(\varphi = T|a,d) = 0.045$, $cr(\varphi = T|a,e) = 0.414$, $cr(\varphi = T|a,f) = 0.869$, $cr(\varphi = T|b,d) = 0.064$, $cr(\varphi = T|b,e) = 0.507$, $cr(\varphi = T|b,f) = 0.907$, $cr(\varphi = T|c,d) = 0.101$, $cr(\varphi = T|c,e) = 0.628$, $cr(\varphi = T|c,f) = 0.941$.

¹⁴ Jeffrey Gottfried and Elisa Shearer, *News Use Across Social Media Platforms* (Pew Research Center, 2016) (reporting that 62 percent of adults in the United States get their news from social media, and 71 percent of that is from Facebook).

with Dave, Simone joins a social networking site and asks, “How has the president been doing? I hear c .”¹⁵ Simone’s post is viewed by Jane and Saura, who are interested in the comment, or maybe even searched for something like it, because they, too, each independently observed c under similar likelihood distribution functions as Simone.

In this case, one could think that Simone’s update credence after talking with Jane and Saura—who share his same moderate belief that the president is bad—is unlikely to change much. After all, unlike Dave, Jane and Saura are not expert with epistemic authority and do share the same belief as Simone. Yet, when the three of them confer and verify that their observations of c were independent, each will wind up with a stronger posterior credence that the president is bad— $cr(\varphi = T|c, c, c) = 0.782$.¹⁶ And if the next search or post from either of the three results in another, seemingly independent, discovery of c , and then another, and another (which could happen, for example, if other people who independently received the same signal c join the conversation), it would take only ten independent opinions for Simone’s, Jane’s, and Saura’s priors to be updated to a radical extreme, i.e., $cr(\varphi = T|c \times 10) = 0.986$.¹⁷

To borrow from Jason Brennan, Simone and his social network friends have turned from “Hobbits” to “Hooligans,” the rabid sports fans of politics, simply by interacting in good faith to learn from each other.¹⁸ Note that this is not a case of “post-truth,” although the effects are similar. The concept of “post-truth” has come to denote circumstances in which objective facts are less influential—if not altogether irrelevant—in shaping public opinion than appeals to emotion and personal beliefs.¹⁹ Likewise, for Simone, the exposure to evidence that radicalizes him is likely to eclipse the value of harder evidence. But this does not happen because Simone will be biased and value how “he feels” more than hard evidence. Instead, Simone will rationally discount the value of that additional evidence based on the other evidence he already possesses.

To better see this, suppose that by the time Simone watches the interview with Dave after he has already engaged with his newfound social network friends. While Dave’s testimony will still have some influence on Simone’s belief, this influence will be much attenuated, $cr(\varphi = T|c \times 10, d) = 0.840$,²⁰ in spite of Dave’s greater expertise and access to harder evidence. Simone

¹⁵ Another version that analogizes to a Google search rather than Facebook interactions might be that Simone googles c and views related search results.

¹⁶ This credence is obtained through the formula: $\frac{cr(c|\varphi=T)^n}{cr(c|\varphi=T)^n + cr(c|\varphi=F)^n} = \frac{0.46^n}{0.46^n + 0.30^n}$. For $n = 3$ (as there are three individuals who received the same signal c) the credence is 0.782.

¹⁷ This figure is obtained through the Bayes formula in fn. 16 for $n = 10$.

¹⁸ Jason Brennan, *Against Democracy* (Princeton, NJ: Princeton University Press, 2016).

¹⁹ McIntyre, *Post Truth*, 13.

²⁰ This credence is obtained through the formula: $\frac{[cr(c|\varphi=T)^{10}] \times [cr(d|\varphi=T)]}{[cr(c|\varphi=T)^{10}] \times [cr(d|\varphi=T)] + [cr(c|\varphi=F)^{10}] \times [cr(d|\varphi=F)]} = \frac{0.46^{10} \times 0.06}{[0.46^{10} \times 0.06] + [0.30^{10} \times 0.82]}$. Please note that this computation is sensitive to decimal digits, so a new computation could give slightly different results although the qualitative result remains unchanged.

would have to get respite from the information circulating on his social network and encounter several random signals in order to reverse the effects of his social networks' conversations.

Thus, while "post-truth" is often associated with a form of ideological purity that discounts the value of evidence *tout court*,²¹ we claim that Simone's transformation is more akin to a form of "rational dogmatism." Unlike pure dogmatists, Simone does not *a priori* exclude the relevance of Dave's testimony.²² Yet, after he has interacted with his social network friends, Simone will come to be convinced—based on the evidence he has gathered in that interaction—that Dave's evidence does not matter that much given the weight and consistency of his other evidence.²³

Note that polarization would still have been the outcome if Simone had seen the interview with Dave as soon as he got to Earth. If this had occurred, Simone's social networking post would have shared different content (namely, "I learned *d*—it does seem like the president is doing a good job.") Then, instead of hearing from Jane and Saura, he would have heard from Jean and Sara, who also, independently, received a similar signal and have a similar starting credence. Jean, Sara, and Simone would then push each other rapidly to the far opposite extreme. Similarly, things would not have changed much had Simone received another signal, namely, *a*. In this case, he would have had a credence that φ is true, $cr(\varphi = T|a) = 0.391$ ²⁴ before interacting on social media. If his search or his post led him then to meet another couple of individuals who also independently received *a*, his credence would have become 0.208,²⁵ and if he had communicated with ten more individuals, 0.012,²⁶ another radical credence.

Moreover, even the expert, Dave, is not immune from these effects. If he had conferred with friends—other experts like him—on social media about *d*, his own updated beliefs after meeting others with the same signal,

²¹ McIntyre, *Post Truth*, 13.

²² This explanation is also distinguishable from, though not inconsistent with, Thomas Kelly's hypothesis that individuals have incentives to look for flaws in the evidence or arguments of signals that conflict with their priors, thereby giving more attention (but also more skeptical attention) instead of less to signals that run against their initial hypotheses. See Thomas Kelly, "Disagreement, Dogmatism, and Belief Polarization," *Journal of Philosophy* 105 (2008): 611–33.

²³ Relatedly, Allen Buchanan argues that there is a kind of tribalistic mentality involved in polarization, which plays a distinctive role in hindering communication and reducing one's evidence pool. Under this mentality, people are labeled ("liberals" or "conservatives") in a stereotyping, homogenizing way (denying significant differences among members of the groups), where the stereotyping in effect holds that everyone in group X is an unreliable source of testimony—because "those people" are all either stupid and hopelessly confused or utterly insincere (in either case, not worth listening to). See Allen Buchanan, *Our Moral Fate: Evolution And the Escape from Tribalism* (Cambridge, MA: MIT Press, 2020).

²⁴ This credence is obtained under the Bayes rule in note 6 and the likelihood distributions functions in note 7, i.e., $\frac{cr(a|\varphi=T)=0.25}{cr(a|\varphi=T)=0.25+cr(a|\varphi=F)=0.39}$.

²⁵ The updated credence here is calculated analogously to the case of Jane and Saura in note 16.

²⁶ The updated credence here is calculated analogously to the case in note 17.

independently derived, would lead to an even more extreme credence that would not be easily dislodged, even if he later received the signal f from a different, respected expert. Concededly, there is a qualitative difference between expert and laymen polarization, as long as we have reason to expect experts to behave like scientists and to value countervailing evidence as some of the most informative. But if academics have mixed motives that interfere with their quest for truth—if, for example, professional reputations are less often advanced by disproving the consensus opinion than they once were—then this difference might become increasingly slim. In the digital era, we cannot presume experts to be immune from what Lee McIntyre has dubbed “the dark side of interactive group effects.”²⁷

The abundance of available signals has reduced the opportunity costs of looking for additional evidence, but the search and customization functions have also removed the role of randomness and constraint in epistemic journeys. Counterintuitively, the result is exposure to less, rather than more, total evidence. Easy access to a wide range of sources of testimony has made the subtle art of selection a determinative step in any unwary traveler’s quest for truth.²⁸

With an abundance of testimony, epistemic discovery is better served when search functions are tempered by aggregation across the largest possible body of evidence. A further modification of our example helps illustrate. Assume that the society Simone visits on Earth is only composed of twenty-one people: ten individuals who independently received signal c , ten individuals who independently received signal a , and Dave. Now further assume that an individual who received c and an individual who received a fall in love and then exchange their signals (*omnia vincit amor*, after all). These individuals’ update credence on φ would then be 0.496.²⁹ Assume now that our two individuals decide to get married and invite to their wedding the remaining population of polarized agents, so that all ten individuals who received c and the ten who received a get to communicate their evidence with each other. Our little society’s updated credence on φ would be 0.457.

Dave is also invited to the wedding, but he is late. By the time he arrives, the guests have already exchanged their evidence and updated their credence. After talking to Dave and hearing d , they would further adjust their credence on φ to 0.058. Yet, if Dave had received f , rather than d , our society would end up with a credence of 0.897. Therefore, neutralizing (or, more realistically, mitigating) selection effects also neutralizes post-truth and

²⁷ McIntyre, *Post Truth*, 60.

²⁸ Note, however, that our theory does not violate the commutativity principle, requiring that the order in which information is acquired does not change the probability of belief. If Simone first observed Dave, and then interacted with ten people who reported “ c ,” the result would be the same.

²⁹ Here, as well as in the cases that follow, the procedure to calculate the update credence is the same as in the case of Simone and Dave.

rational dogmatism effects. Once two groups of moderates talk to each other, Hobbits remain Hobbits, avoiding that intensively selected aggregation of evidence that takes place on the Internet. This is because when evidence is aggregated widely across individuals, what matters in the epistemic process is the kind of evidence one has. The aggregation of softer information no longer produces radical changes in one's belief when that information comes from many heterogeneous sources of testimony. At the same time, the aggregation across a representative sample of people would increase the value of harder evidence when those facts are easily accessible from many, nonpolarized individuals, (that is, Dave's well-honed testimony would reach a much larger number of people) for a net epistemic gain. Under the current state of affairs, however, the dark-side effect of Internet interactions based on soft information and selection mechanisms dominates the bright-side effect of increased access to harder information, producing an overall informational loss.

The question then is how one should proceed toward correcting selection effects. At first blush, the automated filtering of content by Facebook, Google, and other tech giants would seem the natural place to start studying and treating the disorder. Perhaps these platforms can devise new algorithms to expose their users to more representative (or, at least, less nonrandom) sets of messages. But we suspect that the information ecosystem as a whole, and the ways people use it, will be unusually resistant to policy intervention. The next two sections explore why this is so.

III. WHAT HAS KNOWLEDGE DONE FOR ME LATELY?

There is a question we have not yet asked—the proverbial elephant in the room: Why don't people recognize the epistemic impairment that arises from selection effects? And why don't we actively correct for their adverse impact on truth-seeking?

Simone, being Vulcan-like in his single-minded quest for logic and truth, and his knowledge about the likelihood distributions for different sorts of signals, should have worked out that the probability of receiving another signal c after reporting his own testimony would have been great regardless of whether the proposition ϕ were true or false. Simone should proactively search for evidence in a way that eliminates path dependency—that is independent from the signals he has received before. In time, however, as Simone learns about the customized nature of Internet searches and news feed algorithms, an alien like him will adjust his understanding of likelihood ratios for signals. But what about the rest of us, who, presumably, have seen countless times that Internet content is highly curated and path-dependent? Why do some of us, at least, ignore the power of content selection?

If people were motivated only by epistemic goals, they would search for people who are different from them to exchange evidence, like scientists

do.³⁰ But most people do not behave like scientists in everyday life, even when they are motivated by truth-seeking. While some have hypothesized that our inherent cognitive bias makes us ripe for manipulation and exploitation by those with an agenda,³¹ we suspect selection will occur even absent a cognitive bias. For although we have epistemic reasons to engage in Internet discourse, we also have pragmatic reasons that are truth-independent. The main (but not only) pragmatic goal that can explain much of the epistemic failing in selection is the social one. We don't want to fastidiously survey a random sample of information on the Internet; we want to find our friends and engage with them. But these are precisely the people who are most likely to give us a biased stream of cues about the way the world works.

In real life, our social circles are also, to some extent, nonrandom. Our first social circle—the family—shares genetic similarities in addition to geographic and cultural similarities; and school friends and neighbors, too, tend to sort themselves in ways where social peace, continuity, and shared values can take precedence over factual truth. Even before the Internet, social sorting was becoming more intense. Education and religiosity gaps emerged between urban and rural neighborhoods, and marriages bonded people who were more and more similar to each other. Thus, the evidence and signals we exchanged with each other for epistemic purposes have always suffered from some selection effects. But the Internet has vastly increased these effects, making it virtually costless, and therefore more likely, for people who are similar in ways that matter socially (and, therefore, likely to already share similar beliefs) to find and engage with each other.

Those who love the Internet (and we confess we are among them) usually believe that it has two awesome features: you can connect with people, and you can exchange information. But these two goals are often in tension for the reasons we explained above. Worse still, they cannot be neatly separated. Friends want to talk politics, learn from each other, maintain their good standing with each other, and converge on a shared understanding of facts. As the Internet makes it easy to share signals about complex topics like policy, economics, or public health, epistemic goals are more frequently at odds with pragmatic goals, and the latter can predominate the former. And so, there are sometimes pragmatic reasons sounding in social harmony and the desire for group acceptance³² to ignore selection effects in information, even when the selection problem is obvious. The digital era has less room for John Stuart Mill's harm principle, in which truth emerges through exposure

³⁰ Jonathan Rauch, *Kindly Inquisitors: The New Attacks on Free Thought*, (Chicago and London: The University of Chicago Press, 2013), 31–56.

³¹ McIntyre, *Post Truth*.

³² Leon Festinger, *A Theory of Cognitive Dissonance* (Stanford, CA: Stanford University Press, 1957); Solomon Asch, "Opinions and Social Pressure," *Scientific American* 193, no. 5 (1955): 31–35.

to opposing views, because our exposure is largely determined by social phenomena.³³

There are also times, though, that epistemic goals are well aligned with pragmatic goals. This is most immediately observable when new knowledge will have a direct and highly consequential bearing on the learner. Suppose, for example, that Jane is a highly social agent who often prefers to agree with her friends even at the expense of knowing the truth. If Jane's friend Saura told her to drink bleach in order to avoid catching a deadly virus, Jane is likely to stop and think. She will see that her pragmatic interest in social bonding may be diverging from her pragmatic interest in staying safe and healthy. Jane's behavior on the Internet and elsewhere will look much more like the purely epistemic agent since the stakes of being wrong, in either direction, are so great.

Thus, we would predict that the variable mix of pragmatic and epistemic motives will cause people to seek information and update their beliefs in different ways depending on how directly consequential the particular proposition may be.

IV. POLARIZED PEOPLE POLARIZE THE MEDIA

Individual Internet users are not the only actors with mixed motives. Media companies of every sort and, though we don't like to admit it, academics too, have a mix of pressures and goals. They want to pursue truth and epistemic authority, but the pragmatic need to attract readers (or students) and turn a reliable profit is also a concern. When media markets had significant entry barriers, media companies could simultaneously strive for epistemic truth *and* fill their coffers. Since consumers had limited choice, every media "niche" was well populated, so that any strategic motives only *partly* jeopardized their epistemic mission.³⁴

The Internet dramatically lowered the costs for new media entrants, making competition for people's attention quite fierce. At the same time, for the reasons explained above, the users that media companies are in competition for are increasingly polarized. Media companies know that they will not be successful if they provide contrarian evidence in a market of polarized readers. After all, a news agency, for example, can anticipate that readers with an unlimited supply of accessible and inexpensive testimony that is curated based on prior beliefs will eventually abandon the

³³ Thomas Schelling comes to a similar conclusion about residential segregation—small preferences of affinity, in this case living around people like oneself, lead to dramatic segregation, without assuming any racist motives. See Thomas Schelling, "Dynamics of Segregation," *Journal of Mathematical Sociology* (1971); see also Thomas Schelling, *Micromotives and Macrobehavior* (New York: W. W. Norton and Co., 1978).

³⁴ Hotelling theory might predict that the media companies would stay fairly close together to maximize profits, perhaps splitting the difference between the center of the media space and where they think the truth actually lies.

news agency as a reliable source of testimony because they believe—with good reason—that the company routinely provides “misinformation.”

Thus, even traditional media organizations whose tenure and business models predate the Internet will have to supply cherry-picked arguments and true facts (which is easy enough to do in the age of abundant information) so that it can maintain its place as a reliable source of information for an audience. Alternatively, they can resort to models where a contrarian view is always presented to foster a sense of objectivity, regardless of whether this view has any evidentiary basis and often at the expense of further inquiry.³⁵ One way or the other, epistemic motives are sidelined—it is a matter of survival.

Just as social sorting was intensified but certainly not *caused* by the Internet, the competitive pressure for media companies to compromise truth-seeking was also underway well before social networking websites became popular. In the book *Network Propaganda*, Yochai Benkler, Robert Faris, and Hal Roberts show compelling evidence that the strategy to stratify and pitch information to a self-radicalizing audience (particularly, but not exclusively, on the right wing of the ideological spectrum) was the playbook for the Fox News channel, which learned from the talk radio shows that came before it.³⁶ Rush Limbaugh was the pioneer of “community” talk shows, where followers were allowed to call in, but only after vetting them and making sure they shared Limbaugh’s views. As put by Tom Nichols in his book *The Death of Expertise*, “[d]ebate ... was not the point: the object was to create a sense of community among people who already were inclined to agree with each other.”³⁷ MSNBC and Democracy Now! radio use the same strategies for the ideological far left. Every media company is constrained, to some degree, by the risk that it will lose its base if it veers too far from wherever its base currently lands on an issue.

There is little reason to believe that Internet media companies, including Facebook and Google, are immune from the same competitive pressures. There is room to wonder whether these companies became dominant *because* they model their algorithms based on the behavior and revealed preferences of their users. If so, aggressive attempts to undo selection effects could drive users to other platforms that better serve their niches.

Consider, for example, Richard Sorabji’s diagnosis and recommendations for the treatment of filter bubbles. Sorabji suggests that a change in business model would improve epistemic outcomes without badly undermining profits. Specifically, he suggests that social networking platforms should use subscriptions or contextual advertising to fund their operations rather than tailored advertising that profiles each participant. As a practical

³⁵ McIntyre, *Post Truth*, 80–81.

³⁶ Yochai Benkler, Robert Faris, and Hal Roberts, *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics* (Oxford: Oxford University Press, 2018).

³⁷ Tom Nichols, *The Death of Expertise: The Campaign Against Established Knowledge and Why it Matters* (Oxford: Oxford University Press, 2017).

matter, Sorabji may underestimate the impact that this change would have on revenues. Since behavioral advertising has a “click-through” rate several times higher than contextual advertising,³⁸ we would expect overall revenues to plummet without the help of tailored matching. But even assuming that a shift in business model is possible, it will not help if users are actively seeking out content based on some pragmatic objective that biases the flow of testimony they encounter.

Two empirical observations suggest there is evidence that Sorabji’s recommendations could backfire. First, even individuals who primarily use Facebook as a source of news are exposed primarily to news content that their *friends* share, not content shared by advertisers or paying propagandists. Their individual choices in both choosing “friends” and selecting from the content that appears on the news feed contribute much more to selection effects than Facebook’s algorithms do.³⁹ Second, a restriction on behavioral advertising could reward polarizing content producers and penalize the more objective ones. In Europe, when the Data Protection Directive made targeted advertising much more difficult, advertisers were forced to rely on contextual cues based on the content of a website to decide where to place their ads. The biggest losers were general interest news websites, and the winners were websites that produced focused content that helped stratify Internet users into pigeonholes for the advertisers.⁴⁰ Thus, the policies that Sorabji recommends could backfire by sending both users and revenue streams to more fragmented social media platforms.

This unearths two counterintuitive facts that complicate prior assumptions about sound policy. First, media concentration can have a positive influence on epistemic authority and social trust. Perhaps in time, with the right cultural or political shifts, people will change their behavior to better reward media companies when they prioritize truth and open-minded inquiry. In the short-term, though, media competition will have a negative effect on knowledge, introducing incentives for distortionary information practices at the expense of epistemic goals. The conventional wisdom of antitrust that competition enhances quality is strained. It may not be that consumers are better off when media companies compete, at least in the context of political knowledge, if consumers choose among media platforms based on pragmatic preferences (such as social goals) that come in direct conflict with truth.

³⁸ Jun Yan et al., “How Much Can Behavioral Targeting Help Online Advertising?” *Proceedings of the 18th International Conference on World Wide Web* (2009): 261. The example of WhatsApp as an illustration of the validity of subscription services is also not convincing. Since the founders of WhatsApp sold the company to Facebook, it is entirely reasonable to guess that the company, like many startups, was prospecting in the hopes of an acquisition by Facebook or Google.

³⁹ Eytan Bakshy et al., “Exposure to Ideologically Diverse News and Opinion on Facebook,” *Science* 348 (2015): 1130–32.

⁴⁰ Avi Goldfarb and Catherine E. Tucker, “Privacy, Regulation, and Online Advertising,” *Management Science* 57 (2011).

Second, and relatedly, the selection effects that accelerate polarization among social participants on the Internet will be replicated in traditional media and even, to some degree, at universities and other educational institutions. Any authority that depends to some degree on the patronage of readers or listeners will have to compromise their epistemic mission to get closer to one or the other camp of a bimodal distribution. (At least, this is so for sources of information about issues like politics that have an indirect and ambiguous impact on the listener, and require interpretation to make sense of observations.) This is significant because conventional advice to get off social media and use a traditional, trusted source for news has missed the breadth of our epistemic affliction: even traditional sources of authority are catering to a social clique.

V. WHAT TO DO?

The dysfunction we have attempted to explain here is vexing and unusually hardened to policy intervention. We know this because we authors have different ideas about the best (no, the least-bad) solution. All three of us want to reduce selection effects and influence social factors so that seeking out opposing evidence is rewarded instead of punished. But there are no obvious interventions that have a high chance of success. So we will conclude with two rough sketches of policy implications that can hint at where the levers for reasonable government action might be.

A. Improve the epistemic function of the Internet

Policymakers might work to improve the epistemic value of social networks and other Internet platforms. If online media platforms and their users were prevented from having too much influence on the information that is pitched to them—if, for example, Google's search bar had to return the same results for the same search terms for every user, or if Facebook's news feed were forced to display all new posts in reverse-chronological order rather than based on an algorithm that predicts the user's interests, people would interact with a more random or representative sample of signals. Facebook and Google might even attempt to match users with messages having opposing valences to reduce the selection effect. This shares some conceptual similarities to investigations by U.S. and EU regulators over the bias in Google's search results. A solution of this sort would promote the epistemic goal of a user at the expense of the social goals. But it might backfire if users are highly motivated (for social reasons) to find people like them and abandon the large platforms for competitors that use intensive selection.⁴¹

⁴¹ Theoretically there is an alternative to trying to reduce the prevalence of fake news: namely, to raise the prevalence of fake news to a point where people would ignore all attempts

B. Foster greater integration of epistemic and pragmatic goals

An alternative option is to focus policy interventions on individuals rather than the media. If people's social standing depends on their epistemic prowess, there would be fewer clashes between the pursuit of knowledge and pursuit of friendship. If culture were shifted through education and other means to highly value critical thinking and open-mindedness, friends would enjoy each other's company *and* debate, help each other seek unbiased sources of evidence, and help each other reject flawed arguments.

These are not good options. The first may be unconstitutional. The second, leaning on education, is a legal scholarship cliché, all but conceding that the problem is intractable. For information law scholars, this is one of the thrills of studying communications and free thought: they are not easily tamed.

Law, University of Arizona, USA

*Law and Economics, Center for the Philosophy of Freedom,
University of Arizona, USA*

Law, University of Arizona, USA

at testimonial persuasion and just pursue social goals, while seeking the advice of experts for epistemic goals. This is, in some sense, what light-hearted social media platforms like TikTok are providing. But it remains to be seen whether a contaminated flow of information on Facebook will drive users to better (less selective) sources of information. It also runs the risk of depressing the weight that individuals place on epistemic pursuits.