



SERRC
Social Epistemology
Review & Reply Collective

<http://social-epistemology.com>
ISSN: 2471-9560

Algorithm-Based Illusions of Understanding

Jeroen de Ridder, Vrije Universiteit Amsterdam, g.j.de.ridder@vu.nl

de Ridder, Jeroen. 2019. "Algorithm-Based Illusions of Understanding." *Social Epistemology Review and Reply Collective* 8 (10): 53-64. <https://wp.me/p1Bfg0-4ws>.

Abstract

Understanding is a demanding epistemic state. It involves not just knowledge that things are thus and so, but grasping the reasons why and seeing how things hang together. Gaining understanding, then, requires some amount of inquiry. Much of our inquiries are conducted online nowadays, with the help of search engines and social media. I explore the idea that online inquiry easily leads to what I will call *algorithm-based illusions of understanding*. Both the structure of online information presentation (with hyperlinks, shares, retweets, likes, etc.) and the operation of recommender systems make it easy for people using them to form the impression that they are conducting inquiry responsibly, whereas they are in fact led astray by irrelevant cues, spurious links between information or, even worse, various forms of misinformation.

1. Introduction

Much of our information-gathering is carried out online. Anything from reading the news and background stories to researching products or services we aim to buy, looking for travel advice, finding contact details, educating yourself about a new topic, learning new skills, and so much more is done through the use of search engines, by visiting dedicated websites, or by spending time on social media or on discussion forums. Doing so is seamlessly integrated into our lives. We've reached a point where *not* using online resources for information-gathering is becoming very hard.

As Mark Alfano and Colin Klein (2019) note in their introduction to the issue, unlimited information within easy reach might seem like an epistemic paradise. And in many ways, this is true. Judicious use of search engines and social media will find high-quality information on almost any topic, no matter how outlandish or specialized. Often, you can even find folks who share your interests and who are willing to exchange insights or simply chat.

But not all is well in paradise: we're allegedly stuck in filter bubbles and echo chambers (Miller and Record 2013; Sunstein 2017; Nguyen 2018); the internet is brimming with fake news, conspiracy theories, and other forms of misinformation, which tends to spread quickly and widely (Rini 2017; Levy 2017; Benkler et al. 2018; Faulkner 2018; Gelfert 2018; O'Connor and Weatherall 2019); and online information-gathering and sharing may be eroding our epistemic agency (Gunn and Lynch forthcoming) and promoting epistemic vices (Meyer 2019).

I want to make a case that online inquiry is prone to generate *illusions* of knowledge and understanding, rather than the real goods. This is a paradoxical result: although we have a world of knowledge at our fingertips, it's easy to deceive ourselves into thinking we have knowledge and understanding when we don't.

2. Inquiry and Understanding

The point of inquiry is “to find things out, to extend our knowledge by carrying out investigations directed at answering questions, and to refine our knowledge by considering questions about things we currently hold to be true” (Hookway 1994, 211).¹ Inquiry can have different goals, ranging from mere justified belief to demanding epistemic states such as understanding or even wisdom. I’m most interested in *understanding* as an aim of inquiry here, since the claim I will defend is the online inquiry is prone to generate illusions of understanding.

Understanding differs from knowledge in at least two ways. First, it isn’t limited to isolated propositions, but involves knowledge of a number of related propositions. Second, it requires grasping ‘how things hang together’ (Grimm 2011; Baumberger et al. 2017; Gordon n.d.). What sort of ‘hanging together’ is relevant? Explanatory relations or dependency relations: causal relations, part-whole relations, constitution relations, logical relations, and conceptual relations can all be explanatory, depending on the question at hand. I’ll take a pluralistic line here: insight into all of these relations can be conducive to understanding.

Conducting inquiry well requires monitoring and guidance. We can distinguish a number of meta-cognitive tasks relevant to responsible inquiry (cf. Hookway 2003, 199–200).²

a. *Posing good questions or identifying good problems.*

Does the question not have any false or misguided presuppositions? Is it understandable? Does it make sense given background knowledge? Is it relevant, ‘on topic’, neither too narrow nor too broad, timely? Does it admit of an answer (Hookway 2008; Watson 2018)?

b. *Identifying good strategies for carrying out inquiries.*

Given a certain question, what are suitable methods for answering it? What methods are available for solving the problem at hand? Which of these methods can I pursue with a reasonable chance of success? How long will they take? How reliable and accurate will they be?

c. *Recognizing when we possess an answer to our question or a solution to our problem.*

An answer can be right in front of us, but we must appreciate it as such. What counts as an answer or solution? When is it good enough?

d. *Assessing how good our evidence for some proposition is.*

Inquiry can produce lots of information, but we need to be able to assess its quality. Are claims epistemically justified? How strongly? How reliable

¹ Cf. also Baehr (2011); Friedman (2019).

² I’ve added (a) below to Hookway’s list. The formulation of tasks (b) through (e) is Hookway’s, but their explication is mine.

is our own perception, intuition, memory, reasoning? Who are relevant experts and good sources? How should the evidence pro and con be weighed?

e. Judging when we have taken account of all or most relevant lines of investigation. Doing inquiry well also involves knowing when to stop. When have we done enough to close inquiry (at least for the time being)? What methods or lines of inquiry are essential for answering the question at hand and which ones are optional or even redundant?

I'll now go on to argue that various features of the online environment can interact in detrimental ways with the tasks on this list.

3. Online Inquiry and the Illusion of Understanding

Put in very general terms, the worry is this: Online inquiry makes it very easy for us, sometimes to the point of inevitable, to outsource the meta-cognition required for good inquiry to features of the online environment. Since much of the internet isn't designed with the purpose of facilitating responsible inquiry and good meta-cognition, online inquiry can easily go astray, leaving us with illusions of understanding instead of the real thing. In what follows, I'll unpack this idea in various more specific ways.

3.1 Understanding-Seeking Online Inquiry

Understanding involves grasping 'how things hang together'. Hence, when the goal of online inquiry is understanding, we need accurate information about understanding-generating dependence relations. As we'll find again and again below, the internet is both good and bad for this.

Let's start with an example of how the online environment can be conducive to understanding-seeking inquiry. The internet connects information in all sorts of ways: through the use of hyperlinks, websites, and search engines. At least some of these connections represent dependence relations relevant to understanding. Think of Wikipedia: entries contain lots of hyperlinks to sources and connected concepts, and some articles contain schematic overviews of related information. Such hyperlinks are a great way of showing dependence relations, particularly conceptual relations, coherence relations, and epistemic support relations, but also explanatory relations and causal connections.

However, hyperlinks can impede inquiry just as easily. They can make spurious connections between bits of information and the ease of clicking them can derail inquiry by tempting you towards information that's irrelevant to your inquiry. Several other features also create trouble. First, consider the ordering of search results. From an epistemic perspective, it would be best if the top results were mostly or exclusively links to truthful and accessible information. Unfortunately, however, nothing guarantees this. Although its exact workings are proprietary information, Google's PageRank algorithm (Page et al. 1999), which remains

a key part of its search engine, determines the importance of webpages by looking at the amount of links to it. Pages with more incoming links are deemed more important and thus show up higher in the results.³

There might be a story about how, in a better world than ours, this measure of importance could track the epistemic quality of webpages through the wisdom of crowds, but it's no great leap to imagine how PageRank-importance can come apart from truthfulness and accessibility. Recent simulation models confirm this suspicion (Masterton and Olsson 2018). This means that top search results do not necessarily include those webpages that would be most conducive to the acquisition of knowledge and understanding. Clearly, the same will be true for sponsored links appearing near the top search results. You can verify this for yourself by typing in keywords connected to hot-button issues: just try 'age of the earth', 'Obama Muslim', or 'vaccination'. Since few people make it past the first page of results, online searches can easily leave you with poor quality information.⁴

Similar problems occur when you look for information on other platforms. Twitter search results are ranked by their popularity (i.e., amount of engagements), which lacks any straightforward connection to the quality of information. While Twitter discussions can be a good way to learn about the different sides of an issue, they often also contain misinformation, polarizing rhetoric, and other unhelpful material. The same goes for posts united by hashtags. YouTube's recommendations depend on relevance, but also on predicted watch time. The longer users are predicted to keep watching a video, the higher it will show up in the list of recommended videos. The effect of this is that YouTube's recommendations tend towards the extreme and sensational, rather than the truthful (Lewis 2018; Weill 2018; Chaslot 2019; Roose 2019).

Abstracting away from the particulars of individual platforms, the general point is this: Online information is connected in a multitude of ways, both static and dynamic (hyperlinks, responses to search queries, recommendations, hashtags, etc.). While some of those connections represent understanding-generating dependence relations, many others are irrelevant or even undermine and obstruct understanding. Online inquiry thus puts you at constant risk of chasing down spurious relations.

There is a another relevant effect. The mere activity of browsing through search results, clicking various hyperlinks, or following recommendations can give you a false *feeling* of understanding. When you're browsing through digital connections between bits of information, it's easy to think that you're adding to your understanding. After all, it will seem as if you're grasping more and more dependence relations between the issues you're exploring. If you made it past the second page of search results, say, or watched several recommended videos, you'll feel like you've put in the work and earned your epistemic

³ This is a crude simplification. In fact, several other factors and processes influence the ranking of search results, most importantly the commercial interests of advertisers. Since this will typically only detract further from the epistemic quality of search results, I'll ignore these complications here.

⁴ See the statistics on Click-Through-Rates here: <https://www.advancedwebranking.com/ctrstudy/>.

credits. But, as we noted before, it may well be that you're in fact tracking down various spurious connections and misleading information. Hence, you can come to *feel* very intellectually confident for all the wrong reasons—an illusion of understanding.

3.2 Posing Good Questions or Identifying Good Problems

The first step of conducting inquiry well is to figure out the question one is trying to answer. As Lani Watson points out, asking good questions is an intellectual skill: “A good questioner acts competently in order to elicit worthwhile information” (Watson 2018, 358). This involves making contextually appropriate judgements about *when*, *where*, *how*, and *from whom* to elicit information, as well as appropriate judgements about *what* information to elicit. There are many ways in which a question can be bad or misguided: it can fail to make sense, it can be ambiguous, based on false or misleading presuppositions, irrelevant, too broad or too narrow, or fail to admit of answer.

Using the internet can help to improve questions: entering search terms in a search engine, Wikipedia, Twitter, or Reddit and getting unexpected or unhelpful results can help you to recognize ambiguity, false presuppositions, or irrelevance in your earlier thinking. Seeing search results can also help to narrow down or broaden a question, or abandon it altogether. Note, however, that all of these things require active meta-cognition on your part: you need to be alert to the possibilities that there is something wrong with your question and prepared to improve it. In other words, you should not just go along with whatever your search happens to throw up.

Several features of online tools can also cause inquiry to take a bad turn right off the bat. Consider autocomplete functionality first. To borrow an example from Alfano et al. (2018, 310), suppose you're interested in alternative sources of energy and start typing 'alt' into your search window. Google might suggest 'alt right' as a way of completing your query, guiding you to a very different set of results. While it's somewhat unrealistic to suppose that someone who is genuinely interested in alternative sources of energy will suddenly be drawn towards alt-right ideology, autocomplete can work in more subtle ways too. You might be interested in why academics have tenure, so you start typing 'why are professors', at which point Google proposes '... allowed to double dip'.⁵ If you go along with this suggestion your inquiry takes a turn for the worse, focusing on a question with a dubious presupposition and questionable relevance for your original query.

Another worry about search and autocomplete functionality is that it simply goes along with whatever you enter. If questions or search terms are tendentious or have false presuppositions, Google will point you to websites sharing the presuppositions of your query without any indication that they may be false. For instance, roughly half of the results on the first page when you enter 'earth 6000 years old' link to creationist websites defending the claim that the earth is very young. The formulation of search terms can also have implicit presuppositions. Safiya Noble (2018, 111–16) discusses the example of searching for

⁵ This is actually what showed up in my search window (in July 2019).

information about ‘black on white crime’. Rather than directing you to websites with reliable statistics showing that this is a relatively insignificant category in the overall numbers, the top results include white supremacist and nationalist websites with inflammatory and racist content, which reinforce the misguided notion that white people are under threat.

The above examples presume that you start out with a relatively clear sense of what you’re looking for. But sometimes, search queries are triggered by nothing more than a vague sense of a topic. You put in some keywords and see what shows up. In effect, this is letting Google (or another search engine) and its underlying algorithms determine what your specific question is, by coming up with autocompletes and search results. Again, this can sometimes help you to get clearer on an issue, but it might as well saddle you with irrelevant or false information, confusions, bad presuppositions, and other unhelpful things. For instance, Noble (2018) documents extensively how Google reinforces stereotypes by including lots of misrepresentations of oppressed and marginalized groups of people among its top results.

What’s more, the effects of this can extend beyond just one failed search. The first bit of information people encounter has outsized influence on their subsequent thinking about an issue; a cognitive bias known as the anchoring effect (Kahneman 2011, chapter 11). This means that the first couple of autocomplete suggestions or the first few search results you see after typing in some poorly thought-through keywords can continue to pull your subsequent thinking in certain directions, even if you discard them and redo your search.

Similar things can happen if you turn to YouTube, Twitter, or Reddit to find out more about a topic. You can get distracted by YouTube recommended videos, Twitter’s trending topics, or Reddit’s popular posts. YouTube and Twitter also employ autocomplete in their search window. The order of search results is in part a function of their personalized relevance (based on your browsing history, search history, and possibly other information about your online behavior), but also of their popularity (measured by, e.g., interactions). It’s unclear to what extent accuracy or truthfulness plays a part.⁶

3.3 Identifying Good Strategies for Carrying Out Inquiries

Inquiring well requires good inquisitive strategies and identifying those needs judgment. The omni-availability of smartphones, tablets, and computers with internet connections makes Google searches and a suite of default go-to websites for various topics the default strategy for any inquiry whatsoever, regardless of its topic or goal. Again, this isn’t necessarily a bad thing. Google discloses an unbelievable amount of information and is incredibly efficient, typically getting you relevant results in seconds. And plenty of designated websites, forums, scholarly repositories, MOOCs, or YouTube channels can be incredibly helpful in accessing high-quality information.

⁶ This effect is especially troubling for YouTube, in light of its documented history of directing users to radicalizing content.

However, online inquiry often does narrow down our options, especially when we're not actively looking for a more diverse range of information. This matters because the default online options aren't always the epistemically best choice. We've noted above how the order of search results is influenced significantly by other considerations than the quality of information, so Google isn't always your best bet for high-quality information, unless you curate its results yourself by clicking only those results that you already know are from reliable sources. Always only consulting one's favorite news site is very easy, but not conducive to gaining broad understanding by learning about multiple perspectives, especially with so many partisan sites around. Letting YouTube's recommender system be your guide can be epistemically risky, as we noted before.

Amazon's book recommendations can similarly guide you towards misinformation: Diresta (2019) describes how anti-vaxxers and 'natural health' proponents game Amazon's algorithms by leaving lots and lots of five-star reviews to promote materials recommending unsupported medical treatments and making wild claims about the risks of vaccines. Consulting TripAdvisor for travel information will get you mostly information by and for international tourists (and potentially fake or sponsored information), rather than local insider knowledge, which (some) traditional printed travel guides tended to offer.⁷ Similar biases and distortions can be at work with information and reviews found on Yelp, Google Maps, etc.

In all these cases, the problems are similar. Many sectors of the digital economy have developed in such a way that one or a few big tech companies monopolize certain kinds of online information or services (cf. Hindman 2018). These companies are driven by commercial interests and have to use big data and algorithms to streamline their services, so the quality of the information they offer can easily suffer by being biased or manipulable in undesirable ways. Since they provide the default inquisitive strategies for certain kinds of information, inquirers who outsource the task of identifying good strategies for inquiry to the online environment are at risk of becoming poorly informed.

3.4 Recognizing When We Possess an Answer to our Question or a Solution to our Problem

Recognizing answers is easy when it comes to straightforward questions about facts, but when we're looking for understanding, it can be much harder to tell when we have it. Once more, there is both good and bad to be found online. Some moderated web forums do a fine job of pointing users to good answers. They let them vote answers up or down, keep track of how often users report that an answer worked for them, or close topics if they have been addressed to the satisfaction of the original poster. Similarly, websites of officially designated authorities or curated websites with expert information also make it easy to get good answers.

⁷ In 2017 a prank showed how TripAdvisor's algorithms could be gamed so as to make a non-existent fake restaurant the top choice in London by creating a nice-looking website, posting lots of fake reviews, adding photos of fake dishes, and creating an illusion of popularity success by accepting only reservations for months ahead (Butler 2017).

In several other ways, however, the online environment makes trouble for the task of reliably recognizing good answers and solutions. In §3.1, we noted how the links and other connections between online information can reflect dependence relations conducive to understanding, but can just as well reflect spurious or irrelevant relations that leave you epistemically worse off when you take them to be understanding-generating relations. In light of the meta-cognitive task we're considering in this subsection, we can put the point differently, too. The fact that two pieces of information are digitally connected might make you think you've found a way in which things 'hang together', but there's no guarantee that this relation promotes genuine understanding.

There are almost always more search results (websites, tweets, posts, etc.) than you have time to survey and usually indications of their quality or relevance are lacking (besides their place in the overall order of results, which is a highly fallible proxy for quality and relevance). Hence, you might take yourself to have found good enough answers once you've clicked a few of them, but it's often unclear that this is really so.

The vast amount of connections between online information also makes it difficult to assess whether we have (enough) understanding. More links suggest there are more relations that could add to your understanding of an issue, so when have you done enough? To appreciate the point, compare the online situation to a traditional encyclopedia entry or a handbook article. They are supposed to give you a more or less self-contained and complete overview of a topic (with references to further reading and background sources that can deepen your understanding). In contrast, countless further clickable connections make it very unclear whether and when you've done enough to reach some desired level of understanding. At the same time, it's tempting to think that you must have acquired understanding once you have explored a good number of links. In other words, it's easy to deceive yourself with an illusion of understanding.

3.5 Assessing How Good our Evidence for Some Proposition Is

Even offline, assessing and weighing evidence was never an easy task. One often needs prior (expert) knowledge and well-honed intellectual skills and virtues and we all need to deal with individual and social cognitive biases. The online environment can be helpful, but also adds risks and complications. You can double-check claims and arguments by comparing sources, consulting certified authorities, or inspecting the raw data; find out more about the background and reputation of sources and purported experts; trace down references; or check with other people who are trying to figure things out.

Now for the downsides. The worry is that everything from the highest-quality evidence to complete nonsense and misinformation can be put online just as easily and—even more worrying—might become just as popular and frequently visited.⁸ Trying to double-check

⁸ A widely cited study of Twitter dynamics showed that false information spreads faster and farther than true information (Vosoughi et al. 2018). A smaller-scale analysis on BuzzFeed already showed how several fake news stories were more widely shared on Facebook than real news before the 2016 US elections (Silverman 2016).

information, to look for counterevidence, to find and vet experts, and to consult others can thus easily put you in touch with misinformation, misleading evidence, fake authorities, spurious dependence relations, or otherwise epistemically bad inputs, resulting in illusions of knowledge and understanding.

These effects are compounded by the fact that many popular websites and social media offer proxy indicators that have a tenuous relation or no relation at all with epistemic quality. We already discussed Google's PageRank above. Similar things can be said about 'likes', 'shares', 'retweets', and other common measures of popularity. Results that come out of recommender systems on YouTube, Amazon, and other websites suffer from the same problem. The fact that 'people like you' have also watched or bought X is hardly a reason for taking X to be true and trustworthy. The risk is that we rely unduly on indicators which are bad proxies for epistemic quality, thus cultivating illusions of knowledge and understanding.

3.6 Judging When We Have Taken Account of all or Most Relevant Lines of Investigation

The final meta-cognitive task involved in inquiry is knowing when to stop. Again, the internet is Janus-faced when it comes to this task. For straightforward inquiries into everyday factoids, there are plenty of reliable one-stop websites. If you have prior knowledge of which sources, experts, or organizations are reliable authorities on a topic of interest, it's easy to consult their websites and find the information you're looking for. Also, discovering that multiple, independent online sources agree about certain issues or converge on an answer can be an indication that you can stop inquiry.

Things get trickier when you're seeking understanding of complex issues and don't have a prior sense of where to look for reliable information. In addition to the issues already discussed in §3.4, there are further difficulties. The sheer amount of search results and never-ending recommendations make it hard to determine when you've done enough. Endless scrolling on platforms like Facebook and Twitter has the same effect. Because indicators for the epistemic quality of online information are often lacking, you can end up spending lots of time chasing down poor quality information, while not knowing it.

Since we are inclined to take the time invested in an inquiry as a proxy for its thoroughness and quality, we might stop prematurely, because we lull ourselves into thinking we've invested enough effort. The availability of misinformation, especially of the kind that shows up high in search results or recommendations, can also lead us to prolong inquiry when it ought to stop, because you might take it as a new relevant lead. Misinformation similarly creates problems for relying on consensus or convergence as an indicator of true information. When spurious dissent shows up online, you can be led astray by purely manufactured controversy, all the while taking yourself to take into account multiple perspectives responsibly.

4. Conclusion

The internet isn't the epistemic paradise it's sometimes cracked up to be. When we're conducting inquiry online in order to acquire knowledge or understanding—as we often and inevitably do these days—there is a very real risk that we end up with *illusions* of knowledge and understanding rather than the real epistemic goods. Various systemic and design features of the internet in general, and more specific websites and platforms in particular interfere with the process of inquiring well. When we outsource the meta-cognitive tasks required for responsible inquiry to the online environment, we can easily be led astray and take ourselves to know and understand more than we really do. We might mistake popularity and search ranking for reliability, hyperlinks and recommendations for explanatory dependency relations; we might ask the wrong questions or buy into false presuppositions; exclude relevant lines of inquiry prematurely; fail to recognize when we've found a good rather than a bad answer or have acquired genuine rather than illusory understanding; form mistakenly optimistic judgements about the epistemic quality of information; and protract inquiry unnecessarily or stop it too soon. The world that the internet puts at our fingertips is one of misinformation, as much as it is of one information.

In closing, though, let me hint at three lessons to learn—one cautionary but two more hopeful. First, intentionally created fake news and other kinds of misinformation, which have received so much attention in the recent scholarly literature, are no doubt part of the problem, but the issues run deeper. Several of the worries I discussed above derive from systemic features of the internet: the way it orders and organizes information, its lack of epistemic quality control, its openness to all, and various more specific design features of platforms, search engines, and recommender systems. This means that online inquiry would still easily generate illusions of understanding, even if the internet could be purged of much misinformation.

Second, it is no part of my argument to claim that the risks I have identified are inevitable, or that the internet *necessarily* leads to illusions of knowledge and understanding. Whether it does so depends in large part on myriads of design choices, business decisions, and government policies. This means the epistemic potential of the internet can be improved. This is not an easy job. It needs concerted efforts from computer scientists, psychologists, communication scientists, and philosophers, in addition to better business incentives, policies, and law-making.

Third, for all I have said, the internet remains an amazing and unparalleled source of knowledge and understanding, as long as you know where and how to look for it. Skilled searches, judicious use of social media, and a curated diet of websites can improve your epistemic standing vis-à-vis almost any topic whatsoever. But the emphasis is on skill, judgment, and curation. Online inquiry requires cognitive skills and intellectual virtues. Cognition cannot be outsourced entirely to the online environment. Good inquiry is hard work, and this is as true online as it was offline.

References

- Alfano, Mark and Colin Klein. “Trust in a Social and Digital World.” *Social Epistemology Review and Reply Collective* 8 (10): 1-8.
- Alfano, Mark, J. Adam Carter, and Marc Cheong. 2018. “Technological Seduction and Self-Radicalization.” *Journal of the American Philosophical Association* 4 (3): 298–322.
- Baehr, Jason. 2017. *The Inquiring Mind*, New York: Oxford University Press.
- Baumberger, Christoph, Claus Beisbart, and Georg Brun. 2017. “What is Understanding? An Overview of Recent Debates in Epistemology and Philosophy of Science.” In *Explaining Understanding* edited by Stephen Grimm, Christoph Baumberger, and Sabine Ammon, 1–34. New York: Routledge.
- Benkler, Yochai, Robert Faris, and Hal Roberts. 2018. *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. New York: Oxford University Press.
- Butler, Oobah. 2017. “I Made My Shed the Top Rated Restaurant On TripAdvisor.” *Vice* December 6. https://www.vice.com/en_uk/article/434gqw/i-made-my-shed-the-top-rated-restaurant-on-tripadvisor.
- Chaslot, Guillaume. 2019. “The Toxic Potential of YouTube’s Feedback Loop.” *Wired* July 13. <https://www.wired.com/story/the-toxic-potential-of-youtubes-feedback-loop/>.
- De Regt, Henk. 2017. *Understanding Scientific Understanding*. New York: Oxford University Press.
- Diresta, Renée. 2019. “How Amazon’s Algorithms Curated a Dystopian Bookstore,” *Wired*. March 5. <https://www.wired.com/story/amazon-and-the-spread-of-health-misinformation/>.
- Faulkner, Paul. 2018. “Fake Barns, Fake News.” *Social Epistemology Review and Reply Collective* 7 (6): 16–21.
- Friedman, Jane. 2019. “Inquiry and Belief.” *Noûs* 53 (2): 296–315.
- Gelfert, Axel. 2018. “Fake News: a Definition.” *Informal Logic* 38 (1): 84–117.
- Gordon, Emma C. (n.d.). “Understanding in Epistemology.” In *Internet Encyclopedia of Philosophy* edited by James Fieser and Bradley Dowden. www.iep.utm.edu/understa/.
- Grimm, Stephen. 2014. “Understanding as Knowledge of Causes.” In *Virtue Epistemology Naturalized* edited by Abrol Fairweather, 329–345. New York: Springer.
- Grimm, Stephen. 2011. “Understanding.” In *The Routledge Companion to Epistemology* edited by Sven Bernecker and Duncan Pritchard, 84–94. London: Routledge.
- Gunn, Hanna Kiri and Michael Lynch (forthcoming). “The Internet and Epistemic Agency.” In *Applied Epistemology* edited by Jennifer Lackey. New York: Oxford University Press.
- Hindman, Matthew. 2018. *The Internet Trap: How the Digital Economy Builds Monopolies and Undermines Democracy*. Princeton, NJ: Princeton University Press.
- Hookway, Christopher. 2008. “Questions, Epistemologies, and Inquiries.” *Grazer Philosophische Studien* 77: 1–21.
- Hookway, Christopher. 2006. “Epistemology and Inquiry: The Primacy of Practice.” In *Epistemology Futures* edited by Stephen Hetherington, 95–110. Oxford: Oxford University Press,

- Hookway, Christopher. 2003. "How to Be a Virtue Epistemologist." In *Intellectual Virtue: Perspectives From Ethics and Epistemology* edited by Linda Zagzebski and Michael DePaul, 183–202. New York: Oxford University Press.
- Hookway, Christopher. 1994. "Cognitive Virtues and Epistemic Evaluations." *International Journal of Philosophical Studies* 2 (2): 211–227.
- Kahneman, Daniel. 2011. *Thinking, Fast and Slow*. New York: Farrar, Strous, and Giroux.
- Levy, Neil. 2017. "The Bad News About Fake News." *Social Epistemology Review and Reply Collective* 6 (8): 20–36.
- Lewis, Paul. 2018. "'Fiction is Outperforming Reality': How YouTube's Algorithm Distorts Truth" *The Guardian* February 2. <https://www.theguardian.com/technology/2018/feb/02/how-youtubes-algorithm-distorts-truth>.
- Masterton, George and Erik Olsson. 2018. "From Impact to Importance: The Current State of the Wisdom-of-Crowds Justification of Link-Based Ranking Algorithms." *Philosophy & Technology* 31 (4): 593–609.
- Meyer, Marco. 2019. "Fake News, Conspiracy, and Intellectual Vice." *Social Epistemology Review and Reply Collective* 8 (10): 9-19.
- Miller, Boaz and Isaac Record. 2013. "Justified Belief in a Digital Age: On the Epistemic Implications of Secret Internet Technologies." *Episteme* 10 (2): 117–134.
- Nguyen, C. Thi. 2018. "Echo Chambers and Epistemic Bubbles." *Episteme* online first. doi: [10.1017/epi.2018.32](https://doi.org/10.1017/epi.2018.32).
- Noble, Safiya Umoja. 2018. *Algorithms of Oppression*, New York: NYU Press.
- O'Connor, Cailin and James Owen Weatherall. 2019. *The Misinformation Age: How False Beliefs Spread*. New Haven, CT: Yale University Press.
- Page, Lawrence, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report, Stanford InfoLab.
- Rini, Regina. 2017. "Fake News and Partisan Epistemology." *Kennedy Institute of Ethics Journal* 27 (2S): 43–64.
- Roose, Kevin. 2019. "The Making of a YouTube Radical." *New York Times* June 8. <https://www.nytimes.com/interactive/2019/06/08/technology/youtube-radical.html>.
- Silverman, Craig. 2016. "This Analysis Shows How Viral Fake Election News Stories Outperformed Real News On Facebook." *Buzzfeed* November 16. <https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook>.
- Sunstein, Cass. 2017. *#Republic: Divided Democracy in the Age of Social Media*. Princeton, NJ: Princeton University Press.
- Vosoughi, Soroush, Deb Roy, and Sinan Aral. 2018. "The Spread of True and False News Online." *Science* 359 (6380): 1146–51.
- Watson, Lani. 2018. "Educating for Good Questioning: a Tool for Intellectual Virtues Education." *Acta Analytica* 33 (3): 353–370.
- Weill, Kelly. 2018. "How YouTube Built a Radicalization Machine for the Far-Right," *The Daily Beast* December 17. <https://www.thedailybeast.com/how-youtube-pulled-these-men-down-a-vortex-of-far-right-hate>.