Philosophy & Technology (2022) 35: 22
https://doi.org/10.1007/s13347-022-00521-7

RESEARCH ARTICLE

"Google Told Me So!" On the Bent Testimony of Search Engine Algorithms

Devesh Narayanan[1,2] · **David De Cremer**[2]

Received: 27 March 2021 / Accepted: 9 March 2022 / Published online: 26 March 2022

Abstract

Search engines are important contemporary sources of information and contribute to shaping our beliefs about the world. Each time they are consulted, various algorithms filter and order content to show us relevant results for the inputted search query. Because these search engines are frequently and widely consulted, it is necessary to have a clear understanding of the distinctively epistemic role that these algorithms play in the background of our online experiences. To aid in such understanding, this paper argues that search engine algorithms are providers of "bent testimony"—that, within certain contexts of interactions, users act as if these algorithms provide us with testimony—and acquire or alter beliefs on that basis. Specifically, we treat search engine algorithms as if they were asserting as true the content ordered at the top of a search results page—which has interesting parallels with how we might treat an ordinary testifier. As such, existing discussions in the philosophy of testimony can help us better understand and, in turn, improve our interactions with search engines. By explicating the mechanisms by which we come to accept this "bent testimony," our paper discusses methods to help us control our epistemic reliance on search engine algorithms and clarifies the normative expectations one ought to place on the search engines that deploy these algorithms.

Keywords Applied epistemology · Testimony · Search Engines · Algorithmic curation · Trust

This article is part of the Topical Collection on *Information in Interactions between Humans and Machines*

✉ Devesh Narayanan
devesh@u.nus.edu

[1] Department of Philosophy, National University of Singapore, 3 Arts Link Block AS3, #05-22, Singapore 117570, Singapore

[2] Department of Management & Organisation, National University of Singapore, Singapore, Singapore

1 Introduction

Ashish is enthusiastic about climate activism. Over the years, he has collected numerous articles and references about the threat of climate change and the actions we ought to take in response. Betty, a good friend of Ashish's, knows about this list and trusts him as an authority on the topic. One day, she asked him, "Is climate change real?" Ashish goes through his list of references and orders it—placing on top those articles that are most relevant to Betty's question and those that he thinks she might find most useful. He then hands this set of ordered articles to Betty, with no other exposition. Betty trusts that Ashish has done a good job of ordering the articles. So, she only reads the first few articles from the top. In so doing, she acquires new beliefs about climate change.

What has Ashish done here? He doesn't seem to be directly making any assertions about climate change. He is simply ordering content for Betty to engage with. Despite this, it seems like Betty's new beliefs are, in some way, influenced by Ashish's ordering. How should we think about this influence?

We raise this case because it is similar in important ways to another, much more common case. Suppose Betty had Googled the same question. This time, an algorithm would do what Ashish did—ranking and ordering a set of articles based on their relevance to Betty's inputted search string and her user profile. Assuming Betty only clicks on the first few links from her Google search results, the new beliefs she might acquire about climate change would be similarly influenced by Google's algorithms.

Algorithms play an important role in bringing us information each time we interact with online search engines. Different search engines use different algorithms—varying in terms of their design, the imperatives being optimized for (such as content novelty, user 'stickiness', or advertising revenue), and the degree of personalization, among other things. Nevertheless, all search engine algorithms,[1] in essence, perform the same basic functions. That is, based on their assessment of what is best for us and our search query, they bring to fore certain pieces of content and push other pieces to the back of the queue. In so doing, these algorithms play a directive role in structuring the information we are exposed to online and, in turn, contribute to shaping our beliefs about the world. For this reason, it is important to have a clear account of the distinctively epistemic role played by the content-ordering that these algorithms perform. Our paper attempts to provide one such account.

Specifically, this paper argues that users of search engines act as if search engine algorithms are providers of testimony, and acquire or alter beliefs on the basis of this apparent testimony. Making this connection to "testimony"—a well-explored notion in epistemology—can allow us to better understand the epistemic relation between

[1] The terms "recommendation systems" or "content-filtering algorithms" are sometimes used in the computer science literature, as well as in popular discourse, to refer to the algorithms that search engines use. To stay clear of any potential terminological inconsistencies that might arise with the use of these more precise terms, we use the more generic term "search engine algorithms" throughout.
For a discussion on the ubiquity and use of such algorithms, see Chaney et al. (2018).

the beliefs that we acquire online and the content-ordering process carried out by search engine algorithms. In Sects. 2 and 3, we clarify the connections between the epistemology of testimony and our interactions with search engine algorithms. In Sect. 4, we sketch some of the implications of our account—focusing on how users can control the influence that these algorithms have on our beliefs and how we ought to treat the search engines that deploy them.

2 Algorithmic Content-Ordering as Testimony

By way of motivating some of the arguments in this paper, let us begin by looking at a previous attempt to characterize googling[2] as related to testimony, by Gunn and Lynch (2018). To quote them in full:

> Consulting Google (the actual web-browser) is like asking someone about whom to ask for an answer. And any act of googling (finding information online) by its very nature is dependent upon the beliefs and actions of other people. In that way, googling is more like testimony. (p. 43)

According to Gunn and Lynch, googling is similar to testimony because, in both cases, we are reliant on the beliefs and actions of others when acquiring or altering beliefs. To build their case, they first draw a distinction between "consulting Google (the actual web-browser)" and "the act of googling (finding information online)." The former seems to refer simply to the act of typing up a search-string and retrieving Google's list of search results. "Consulting Google," thus, simply tells you who to seek testimony from, but not what to believe. "Googling," however, refers to engaging with the content we select from a Google results page and (possibly) acquiring or altering our beliefs in the process. Since this content is produced by others, we necessarily and directly rely on others to form new beliefs while "googling." Herein lies their proposed connection to testimony. When we consult Google, we are put in contact with a wide range of potential testifiers—those who produce blogs, news websites, forums, videos, and other such forms of online content—and, in turn, adjust our beliefs on the basis of the assertions made by these testifiers. As such, googling resembles testimony because of our epistemic reliance on those who produce the content that gets presented to us when we perform a Google search.[3]

[2] Since Google is currently by far the most popular search engine (cf. StatCounter, n.d.), throughout this paper, we interchangeably use the phrases "Google's algorithms" to refer to search engine algorithms and "googling" to refer to online searching.

[3] Although Gunn and Lynch's account focuses primarily on our epistemic reliance on those who produce the content that we engage with online, one might also consider how we are reliant on fellow Google users—whose data and search habits improve the quality of the recommendations Google provides to us—as well as Google's engineers who designed these algorithms. On a broader reading of their claim that googling resembles testimony because of it is "dependent on the beliefs and actions of others," one might reasonably extend their account to include these other groups and thereby build a more comprehensive view of how googling resembles testimony. Gunn and Lynch, however, do not explicitly mention these other groups, and our own suggested improvement for their view (i.e., that we should better consider our epistemic reliance on search engine algorithms themselves) still holds even on this extended

Gunn and Lynch's account can help us begin to understand the epistemic role played by Google when we acquire or alter beliefs online. However, it is not clear whether they sufficiently take into account the work that Google's algorithms are doing to filter and order content when we consult the search engine. On one reading of their account,[4] the *only* relevant consideration is the epistemic reliance we might have on the people who produce the content we engage with. Such an account would certainly be incomplete. Suppose, instead, we had a website (say, "Xoogle") that, each time we visit it, randomly directs us to a new piece of content. Clearly, while xoogling, we would be similarly reliant on the people who produce the content we end up engaging with, even if we happened to find this content randomly. But it would seem strange to suggest that xoogling is a kind of testimony in precisely the same way that googling is. Surely, any account of how we form beliefs when we engage with content online must consider the role that search engine algorithms play in bringing that content to us. The need to develop such an account is exactly what motivates this paper.

In this paper, we argue that users treat search engine algorithms as if these algorithms provide us with testimony. The implied connection to "testimony," however, already raises an obvious worry. A common view in the epistemology of testimony is that the primary speech-act of testimony is a speaker's saying, telling, or asserting something (Adler, 2017; Searle, 1969). Suppose, in our previous case, Ashish had instead directly made an assertion, say, "Because of climate change, ice sheets have shrunk by more than 4000 gigatons in the last two decades." Since, as we know, Betty trusts Ashish to be an authority on climate change, she believes his assertion. In this case, it would not be a problem to say that Betty acquires a belief based on testimony provided by Ashish. When Ashish makes an assertion, he puts forward a proposition that he represents as true. Betty adds this proposition to her set of beliefs when she accepts his testimony. It is this connection between Ashish's assertion and Betty's belief that underpins traditional accounts in the epistemology of testimony. But no such connection seems to exist when Ashish simply handed her an ordered list of articles, or indeed, in the case of search engine algorithms. Can something count as testimony even in the absence of assertion?

Interestingly, Regina Rini offers one account of testimony in the seeming absence of assertion (2017). She discusses the case where someone you trust, to whatever degree, shares a piece of fake news that, in time, appears on your social media feed. In this case, it is not immediately clear whether or not this person, by sharing or retweeting a fake news story, is actually asserting the content of that story. Rini

Footnote 3 (continued)
version of their account. As a result, we did not discuss this extension in our main text, but instead raise it as footnote for interested readers.

[4] Gunn and Lynch do acknowledge that googling is a "preference-dependent" mode of inquiry—Google tells us who to consult based on its assessment of what we might like and what links we will click (p. 43). Preference dependence is, of course, one important way in which google differs from "xoogle." However, they do not specify how exactly this feature of preference dependence is relevant to their account of how googling resembles testimony. As such, a reading that the only relevant consideration for them is our dependence on the actions and beliefs of others is, while perhaps strict, not uncharacteristic.

suggests that this ambiguity arises from "disputed norms of communication." Social media is a relatively new mode of communication, and we have not yet come to a consensus about whether, for instance, a retweet counts as an endorsement of the content being retweeted. Because of these disputed norms, we do not seem to be able to say for certain that an assertion is being made. Even so, "we treat social media sharing as if it were ordinary testimony, at least until something goes wrong" (p. 48).Rini calls this "bent testimony." In general, we might say that someone (or something) provides us with "bent testimony" if we treat them *as if* they had provided us with ordinary testimony, even though disputed norms of communication make it difficult to ascertain whether this person (or thing) is *actually* testifying.

A closer look at how we usually interact with Google suggests that search engine algorithms might be similarly providing us with bent testimony. Notice, for a start, that a Google search does not return a single answer to a query.[5] Rather, Google's algorithm compiles and orders all the articles it can find that are related to an inputted search query. This could mean tens of pages of search results and hundreds of links to relevant articles. These articles are likely to vary considerably in their assertions. So, if the average user followed enough links, it is likely that they would be exposed to a variety of different perspectives that they would then need to weigh against each other before deciding what to believe (Boutin, 2011; Weisberg, 2011). But most of us rarely do this.

Rather, as various empirical studies have suggested, many Google users rarely click more than a few links from the top. Users typically do not look further than the first few pages of search results (Jansen & Spink, 2006) and are strongly biased towards links that are higher up in the search results even when they are obviously less relevant to their search query (Pan et al., 2007). In a comprehensive study by Joachims et al. (2007), it was found that people spend the longest amount of time viewing—and are most likely to click on —the first link of a search results page, with sharp drops in viewing-time and click-through rates from the second link onwards. Various explanations have been offered for why this is the case. Oulasvirta et al. (2009) found that users are overwhelmed by the number of choices available to them in a long list of search results, and tend to look for proxies (such as order-rankings) to make their choice easier. Moreover, users are typically unaware of the influence that order-rankings have on their choice and do not view their decision to select the results at the top of their list as irrational or problematic (Epstein & Robertson, 2015). Expertise and familiarity with the search engine seem to play a role too: it has been shown that users who type more complex search strings with advanced Boolean variables tend to also look at more pieces of content further down their list of results (Jansen et al., 2000). Studies have also examined how general beliefs about the reputation and competence of the search engine—including the

[5] More recently, Google has implemented a feature called "Quick Answers" that highlights at the top of the first page of search results a single answer to certain types of simple search queries—like "Who is the Prime Minister of Singapore?", or "What's two plus six?"—where single answers are possible. But even so, below this highlighted "quick answer," you would still find the set of ordered links as you would in an ordinary Google search.

credibility of Google's brand, the design and layout of its website, and the reputability of its algorithms and engineers—can influence our trust in the relevance and truthfulness of the results at the top of a search results page (De Cremer et al., 2019; Jansen et al., 2009; Westerwick, 2013). Thus, as these various studies show, the average user seems to treat Google with "a default, *prima facie* trust" (Gunn & Lynch, 2018). Moreover, and important to the present line of reasoning, it is often the case that the first few links in a search results page usually point to articles that are quite similar to each other.[6] It is then also safe to say that the average user in reality only engages with one or two unique assertions per search.

With this reasoning in mind, let us return to our original example. Suppose Betty, after a Google search, says, "Google told me that temperature fluctuations are common in nature, and man-made carbon emissions have little to do with it." It is likely to be the case that the top choices for Betty's search results point to articles that make this particular assertion. Articles that feature later on in her search results might have made a different or conflicting assertion (say, "There is an overwhelming scientific consensus that man-made carbon emissions contribute directly to rising temperatures"), but, Betty is likely not to read those. Here, Betty's statement—that Google told her some fact about climate change—seems like a reasonable thing to say, because she is treating Google's algorithms as if they are *recommending* the content featured at the top of a search result.

As such, because of the way we ordinarily use Google, the assertions that are brought to the top of our search results by its algorithm are often the only ones we will engage with (and potentially accept). We trust that the algorithm will place content documenting true assertions at the top of search results. This trust is what allows us to reduce the many, possibly diverse points of view represented in the results for any search query to a single assertion that fits the statement: "Google told me X." In other words, this trust is what reduces content-curation to a recommendation. These algorithmic recommendations, in practice, seem comparable to the assertions of an ordinary testifier.

However, according to Rini's account, testimony is bent when there are disputed norms of communication that make it difficult to decide whether a seeming testimony includes an assertion or not. For example, in her discussion of the bent testimony provided by retweeters of fake news, the relevant disputed norm is whether the retweeting of an article counts as an endorsement of its claims. Such disputed norms are an important part of what makes bent testimony "bent." Are there comparable disputes about googling?

To locate the relevant disputed norms about googling, it is helpful to once again compare our interactions with retweeters and with search engine algorithms. In Rini's example, disputes seem to arise because, when a retweeter recommends a

[6] This is often, but not necessarily the case. The articles at the top of a search result have high 'similarity scores', are hosted on websites of comparable repute, and are closely related to the search input we type in. If, for instance, two sources deemed reputable by the algorithm (say, *The New York Times* and *The Washington Post*) make dramatically different assertions about a topic, these might appear together at the top of a search results page. It's hard to say how often this happens.
For details on the design of Google's 'PageRank' algorithm, see (Page, 2006).

piece of content, it seems reasonable to think that they are also endorsing its claims. But they might not intend such an endorsement, and it is difficult to ascertain the retweeter's intent only based on the contents of their tweet. So, absent other information, we "implicitly assume that our social media interlocutors … believe what they share" (Rini, 2017, p. 48). There are, of course, some differences in the work that retweeters and search engine algorithms put in to bring a piece of content to the end user. However, there is at least one important similarity: in both cases, it appears as if a piece of content is being recommended—whether by the retweeter or by the algorithm. It is in this shared practice of content recommendation that we might locate the disputed norms relevant to our interactions with search engine algorithms.

Consider the following example: suppose Ashish writes to Betty one day, "I came across this video lecture describing how the oil industry withheld information about climate change from us for several years. I think you'll find it very interesting!". Can Betty ascertain whether Ashish is endorsing the claims in the video? It is hard to say. Ashish doesn't say whether or not he approves of the video, only that he thinks Betty will. It could be that Ashish completely disagrees with the claims made in the video, but thinks that it would provide some insights that could benefit Betty's ongoing exploration of climate change. But Betty cannot be faulted for thinking that this video recommendation—especially from someone she perceives to be an authority on the topic—reflects Ashish's endorsement for its claims. Absent contextual information about why and how a piece of content was recommended, it is thus difficult to ascertain whether the recommendation counts as an assertion of the content's claims. The relevant disputed norms when it comes to content recommendations, therefore, seems to arise from slippage between the statement of recommendation, "I think that this P-asserting piece of content is good for you" and the assertion, "I think that P."

Earlier in this paper, we suggested that most of us treat search engines as if they were recommending the content ordered at the top of a results page. Usually, algorithms order content based on relevance, novelty, personalized fit, or other similar criteria. However, when we selectively choose to engage only with content ordered at the top, it is unclear whether this choice is based on our assessment that this content is, say, more relevant to our inputted search string, or whether it is more likely to contain true claims. At least, both possible bases for our choice seem reasonable. As such, there seems to be a similar disputed norm about whether a search engine algorithm merely recommends some P-asserting piece of content based on truth-independent criteria (such as relevance to one's inputted search string), or whether it endorses P as true.

Importantly, this disputed norm is not a mere eccentricity that only exists in the interactions between users and search engine algorithms. Google itself routinely describes its own search ranking process by appealing to both truth-related and truth-independent criteria interchangeably. When describing their methods for detecting spam, for analyzing a search query's keywords to locate relevant websites, and for analyzing the accessibility and user-friendliness of website designs, they typically refer to truth-independent criteria such as "usefulness" or "relevance to the search query" (Google Search, n.d.-a). However, when describing their methods for assessing the expertise, trustworthiness and authoritativeness of content, and their

"search quality" guidelines, they regularly suggest that their highly-ranked content should contain true claims (Google Search, n.d.-b; Google Search, n.d.-c). Some of Google's newer features—such as "Quick Answers," "Featured Snippets," and "Knowledge Graph"—which present prominent information in a card or image separately from the list of search results—are described as providing convenience and ease-of-use to users (Google Search, n.d.-d), even though these cards are arguably as close as it gets to a direct assertion from Google about what it considers to be the correct answer to our inputted search query. Popular writing about search engines also frequently make similar appeals to truth-independent criteria (such as relevance, convenience, transparency) while concurrently suggesting that search engines are crucial for helping us find true and accurate information (cf. Hillis et al., 2012; McKinsey, 2011; Roesler, 2021; Rothman, 2018). These divergent accounts suggest persistent and widespread confusion about whether search engine algorithms merely recommend pieces of content based on truth-independent criteria or whether they are helping us find true information. This is, we suggest, the relevant disputed norm at play when we treat search engine algorithms as providers of bent testimony.

To summarize our argument so far: we have suggested that Rini's account of "bent testimony" provides a useful basis for understanding the influence that search engine algorithms have over our beliefs. The testimony of these algorithms is bent because:

(1) We trust that search engine algorithms reliably order relevant and useful content at the top of a search results page. Consequently, we only engage with a few pieces of content, and therefore only encounter one or two unique claims per search. In so doing, we treat search engine algorithms as if they are "recommending" the content featured at the top.
(2) Disputed norms exist about whether the recommendation of a piece of content counts as an assertion of its claims.

As such, statements of the form "Google told me X!" are more than just a turn of phrase. Because of the way in which we usually interact with search engine algorithms, it indeed seems as if these algorithms are making assertions and providing testimony—much like any other testifier. When we form beliefs based on the content we engage with online, we do so at least partly on the basis of such testimony. However, importantly, this testimony is "bent"—search engine algorithms seem to provide us with testimony only when certain disputed norms are present in our interactions with them.

3 Bent Testimony Versus 'Real' Testimony

Throughout this paper, following Rini (2017), we have added the modifier "bent" to the well-explored notion of testimony. At this point, it is important to briefly clarify this modifier's purpose. In line with Rini's use of the term, one might say that X provides them with bent testimony when:

(1) Within certain interactional contexts, and in particular, when there are disputed norms of communication, it is not clear whether X is providing testimony or not.
(2) Despite this, one acts *as if* X was providing us testimony, and acquires or alters beliefs on that basis.

Importantly, this does not entail that X is *actually* providing testimony, but simply that, within certain contexts, we treat X *as if* it was a testifier. The primary overlap between bent testimony and "real" testimony is that in both cases, we acquire or alter beliefs on the basis of the actions or assertions of another. There may be other conditions necessary for something to count as "real" testimony—and, in fact, there are several competing accounts in the epistemology of testimony about what these conditions might be. However, it is not necessary for the purposes of this paper to engage with these debates. This paper does not argue that search engine algorithms *in fact* provide us with testimony, only that in our usual interactions with search engines, we act as if they do.[7]

Treating search engine algorithms as providers of bent testimony has far-reaching implications for how we ought to interact with search engines while acquiring beliefs online. Before discussing these implications, however, we will consider two possible objections about algorithmic bent testimony.

3.1 Retweeters Could Have Testified, but Algorithms Strictly Cannot

Consider, once again, Rini's original account of the bent testimony of social media sharing. One might be confused about whether the specific act of sharing content counts as testimony, but it does not seem at all controversial to suggest that social media posts *could* contain testimony. Rini's puzzle is motivated precisely by this tension between the ordinary testimony of social media posts that make direct assertions, and the bent testimony of social media sharing. It might not be immediately clear that such a tension exists when we talk about the bent testimony of search engine algorithms—since it might not be obvious that algorithms *could have* testified. This might then be an important disanalogy between retweeting and googling. If we entirely rejected the notion that it is even in principle possible for algorithms to make assertions and provide testimony, it would be difficult to conceive how their testimony could be, in the case of content-ordering, "bent."

Testimony has often been considered to be strictly originating from human beings (Tollefsen, 2009). This is because, at least:

[7] Some might prefer to avoid using a well-explored term like "testimony"—with modifier or without—in favour of terms like "evidence" or "influence." But, for this paper, it does not matter so much what we call it. We would only have to engage in these terminological discussions if we were arguing that search engine algorithms were, in fact, providing testimony (and, perhaps, not just "evidence"). Hopefully even those who prefer a more restricted use of "testimony" would agree that, in some cases, we might act as if someone (or something) was giving us testimony and acquire/alter beliefs on this basis, even if they, in fact, were not testifying. This concession is all we need for the arguments in the paper to proceed.

(1) Testimony requires the testifier to sincerely *believe* the claims they are testifying about
(2) Testifiers *intend* to present evidence or make an assertion about a matter in dispute (Coady, 1992; Lackey, 2008; Pagin, 2016).

By comparison, even metaphorically, it seems inaccurate to suggest that search engine algorithms "believe" what they recommend and "intend" to provide testimony. It seems like humans alone can possess beliefs and intentions, and, as such, it is incorrect to talk about algorithms in comparable terms.

Certainly, this claim is not without controversy. Some have argued that animals can have beliefs (Bermúdez, 2007; Rescorla, 2009), as can groups such as companies or governments (Tuomela, 1992). Similarly, a full discussion about whether algorithms can have beliefs and intentions, then, needs more careful arguments to be mounted on either side—but such a discussion may not be necessary for this paper. We suggest that neither belief nor intention is strictly necessary for testifiers to provide testimony. Consider the following brief examples, both inspired by Jennifer Lackey[8]:

(1) A creationist teacher discusses the theory of evolution with her class, despite not believing the theory. Her students, unaware of what their teacher actually believes, come to believe the theory of evolution.
(2) Ashish is talking with some friends about climate change, and asserts,"The recent wildfires and hurricanes are all related to climate change – and it's all our fault!". Betty happens to be passing by and overhears this conversation. She comes to believe this claim—even though Ashish did not intend for her to hear his assertion.

In both cases, beliefs are being acquired on the basis of testimony—even if, in one case, the testifier didn't believe the claim being testified, and in the other, they didn't intend to provide testimony. It is then unclear why we should deny the possibility of algorithmic testimony solely on the basis of them not having beliefs and intentions.

We have shown earlier how, from the point of view of users, the testimony from search engine algorithms looks and feels like ordinary testimony. One would need compelling and obvious reasons to demarcate this seeming testimony from "real" testimony. However, it is not clear that such reasons exist. As such, claims about the impossibility of algorithmic testimony appear to be overstated.

Moreover, several scholars have further developed the notion that there does not seem to be a principled reason to deny the possibility of algorithmic testimony. Many have gone further than we have here, to build positive accounts of how algorithms can be testifiers. Some, like Bruno Latour, wholly reject the idea that there

[8] Despite opening up the possibility for algorithmic testimony for this essay, Lackey seems to be among those scholars who believe only humans can testify (2008, p. 189). See Freiman and Miller (2020) for a more comprehensive engagement with her concerns about non-human testimony.

is any meaningful distinction at all to be drawn between the testimony of humans and non-human artefacts (Latour, 2005, 2012). Other accounts preserve a distinction between human and algorithmic testimony but treat both forms as valid. Freiman and Miller (2020) provide a thorough overview of the various positions available to those seeking to defend algorithmic testimony. Nickel (2013) and Wheeler (2020), for instance, suggest that machines can assert only if they are capable of insincerity—of violating the epistemic norm that assertions should represent the truth. They discuss how machines can be designed to violate epistemic norms that we would normally apply to testifiers, and if so, we cannot but think of such machines as testifiers. As another example, Freiman and Miller (2020) develop an account of "quasi-testimony,"[9] which features a similar notion of conformity to ordinary norms of testimony. On this account, machines can be said to be quasi-testifiers if they satisfy two conditions: first, that they are designed to produce outputs that resemble testimony phenomenologically, and second, that their outputs conform either with epistemic norms that they are parasitic on (as Nickel and Wheeler suggest), or to what would have been an epistemic norm of testimony in a similar context.[10]

Since, as suggested earlier, we are not arguing that search engine algorithms are in fact testifiers—it falls outside of the scope of this paper to fully engage with these various accounts of algorithmic testimony. We primarily wish to show that Rini's original puzzle does not disappear entirely when we move away from retweeting, towards googling. For this purpose, it suffices to show, as we have done, that plausible defenses of algorithmic testimony have been mounted elsewhere, and that we do not have very good reasons to insist on a clear line separating human testimony from all other kinds of testimony. As such, it would be premature to rule out the possibility of algorithms as providers of bent testimony simply based on the *prima facie* unintuitiveness of algorithmic testimony.

3.2 The Bent Testimony of Content-Algorithms Does Not Sustain the Corresponding Beliefs Acquired by the Search User

Suppose that a trusted friend hands you a sealed, opaque box and tells you that it contains a beetle. If you accept his claim, for the moment, your belief that the box contains a beetle is sustained by your friend's testimony. You then open the box, and

[9] Freiman and Miller (2020) are among those who believe that there is a meaningful distinction to be drawn between the testimony of humans and algorithms. Their use of the modifier "quasi" while discussing algorithmic testimony is mainly to emphasize this distinction. Despite this, they believe that we might similarly acquire or alter beliefs on the basis of both ordinary testimony as well as "quasi" testimony.

[10] Freiman and Miller suggest that when a machine's output, by design, resembles human testimony (e.g., an automated announcement in a natural language on a loud speaker), the machine's designers "count on its users to correctly decipher the meaning of the output and correctly assess its validity because they recognize the testimony-like epistemic norms under which the output is produced" (p. 13). In turn, when users expect this machine output to conform to epistemic norms that would be in place for similar interactions (e.g., an announcement on a loud speaker made by a human), we are treating this machine as a "quasi-testifier." See Freiman and Miller (2020, pp. 11 – 14) for a more thorough discussion of quasi-testimony.

indeed, find a beetle inside. However, once the box is opened, the evidence from your own perception greatly diminishes your epistemic dependence on your friend's testimony. Your belief that the box contains a beetle, then, is no longer sustained by testimony.

Now, suppose your friend instead gives you a sealed, transparent box and tells you that it contains a beetle. This time, you do not need to open the box to see the beetle inside. Although your friend still provides you with testimony, your belief that the box contains a beetle is, at best, only weakly—if at all—reliant on his testimony.

One might think that this latter case resembles a Google search. When we perform a Google search, we receive a considerable amount of additional information about the content being featured, in addition to their position on the list of search results. Typically, we would be able to see the source of the website hosting this content, the date it was posted, and a brief preview of the content, among other things. These other pieces of information might help us decide which links to click on and which content to engage with, thereby influencing the beliefs we acquire while searching. As such, even if we accept the notion that Google's algorithms are providers of bent testimony, it is unclear whether the beliefs we acquire while searching are reliant on this testimony.

There are certain clear cases where it is indeed the case that other factors—such as the reputation of the website hosting the content, or the perceived quality of the previewed content—are dominant factors in our decisions to select articles from a search result page. Suppose Betty was using Google to navigate to the Wikipedia page for the *Paris Agreement*—perhaps because she trusts Wikipedia to give her a quick and truthful introduction on the topic. Here, clearly, it would not matter which position Wikipedia occupies in her list of search results. In such cases, where we have strong opinions about which sources are trustworthy, the ordering of content by search algorithms might not play an important role. Moreover, another important disanalogy between googling and our vignette is that, in the former case, we are explicitly requesting for specific information to be provided to us, via our inputted search query. Thus, one potentially easy way to assess pieces of content—independent of their position on the list of search results—is to evaluate the extent to which they are relevant and useful for answering our search query. If Betty searches for the *Paris Agreement*, and the first few results happen to contain guides for travelling to Paris, she would, correctly, disregard them.

However, we often consult Google about topics we might not know enough about to have prior opinions about which sources are trustworthy. And, sometimes—especially when searching about hotly contested topics—we may encounter multiple pieces of content that contain conflicting assertions, but appear equally relevant to our inputted search string. In such cases, algorithmic bent testimony could play a larger role. Unless the links at the top of such searches are of obviously poor quality, or are hosted on websites known to contain conspiracy theories and inaccuracies, or are obviously irrelevant to our inputted search query, we would rely on algorithmic bent testimony to select pieces of content to engage with.

As such, although we do receive other potentially relevant pieces of information when we carry out a Google search (e.g. reputation of the source, quality of the previewed content, etc.), such information may not always override the influence of the

search engine algorithms' bent testimony. This epistemic relationship might be best captured by tweaking the analogy from earlier once again. Suppose you ask your friend to hand you a box containing a random insect. He then hands you a translucent box. You can make out a small, dark moving object that seems like an insect, but cannot make out what it exactly is. He then tells you that the box contains a beetle. If he had instead told you the box contained a butterfly, you might have disregarded his testimony. However, since his testimony accords with what you expect to find inside the translucent box—i.e., a small, dark insect—you believe him. In this case, your belief that the box contains a beetle is not sustained *purely* based on his testimony, since you have some (albeit imperfect) ways of perceiving what the box contains. However, you are still epistemically reliant, in some way, on his testimony.

Similarly, it could be the case that other pieces of information influence which content we engage with (and which beliefs we acquire) when we search. However, in many cases, it does seem that the ordering of content in search results also plays an important role. The empirical results discussed earlier—which describe how many Google users only select a few links from the top of a results page (see, e.g., Pan et al., 2007; Joachims et al., 2007)—attest to this fact. We cannot so readily dismiss our reliance on the bent testimony of search engine algorithms.

4 Implications

Thus far, we have developed an account of how search engine algorithms provide us with a kind of bent testimony, and we have shown that this account can withstand at least a few immediate objections. In this final section, we turn to one remaining question: how might this account help us be more epistemically responsible for the beliefs we acquire while searching online?

We have shown how users only come to accept this bent testimony within a particular interactional context: namely, when we treat algorithms as if they were recommending the content ordered at the top and when we treat the recommendation of a piece of content as an assertion of its claims. The implications of our account derive directly from this previous discussion of context and norms.

Firstly, and most directly: should we deem this necessary, if we decrease our credence in the notion that the content ordered towards the top of a list of search results is any likelier to contain true claims than the content later in the queue, we would, in turn, decrease our reliance on the bent testimony of search engine algorithms. Accordingly, we might look further down the list of search results and pay closer attention to other factors—such as the reputation of the source website or the quality of the content preview—while deciding which content to engage with. We might also like to seek a "second opinion." We might, for instance, consult another search engine, like DuckDuckGo, Yandex, or Bing, or alter our personal settings on Google—by changing our location or logging out of our account (or going "incognito")—to obtain search results that are personalized for a different person.

These are all, of course, commonly proposed as "digital media-literacy skills" we should all adopt while navigating the Internet. Various primers have been written about the skills necessary to search effectively online, and the strategies

necessary to ensure that larger sections of society acquire these skills (see, e.g., Nentwich & Konig, 2012; van Deursen & van Dijk, 2014). Even so, we provide a fresh reason to pursue these skills—for the sake of reducing our epistemic dependence on the bent testimony of search engine algorithms. Users are not, as is sometimes assumed, powerless, passive recipients of algorithmic recommendations from search engines. The ordering and filtering carried out by a search engine algorithm only amount to a "recommendation" under certain interactional contexts. And, importantly, we retain at least some control over these interactions. By adjusting our credence in our belief that search engine algorithms expose us to true claims, we might have better control over how we acquire beliefs online.

However, it seems infeasible to expect that everyone would be able or willing to adjust their personal search strategies and appropriately control their reliance on the bent testimony of search engine algorithms. It therefore serves to ask: how should we treat search engines (and the organizations that operate them), given that many of us acquire beliefs on the basis of the bent testimony of their algorithms?

On one view, search engines should be treated as *neutral* and should not be held responsible for the pieces of content their algorithms serve up. Google, for example, seems to think that this is the case—that search results are simply a "reflection of the content publicly available on the web" (Google Privacy & Terms, n.d.). We might better understand their position by thinking in terms of the disputed norm discussed earlier—about whether a recommendation of a piece of content counts as an endorsement of its claims. Those—like Google—defending neutrality must necessarily say that recommendations are *not* endorsements, at least not in the case of search engine algorithms. But this would be in stark opposition to how most of us actually interact with these algorithms. When we accept their bent testimony, we do in fact see algorithms as endorsing the pieces of content ordered at the top. The norm adopted during most interactions with search engine algorithms seems to be opposite to the norm that defenders of neutrality would have us adopt. This seems manifestly unreasonable.

Accordingly, since many of us continue to acquire beliefs on the basis of the bent testimony of search engine algorithms, we ought to treat search engines as testifiers. By adopting norms that are consistent with our ordinary interactions with its algorithms, we should insist that Google is not neutral, but rather an active participant in the process by which we acquire beliefs online. We should, then, expect them to behave as responsibly and sincerely as we would expect from any other testifier.

Having such expectations from Google is not just reasonable and consistent, but also normatively necessary. To see why, consider the two following illustrations of interactions with ordinary testifiers:

(1) Betty goes into a shopping mall, stops the first person she sees there, and asks him whether climate change is a real and immediate threat.
(2) Betty asks Ashish, her close and trusted friend, whether climate change is a real and immediate threat.

In (1), it would be unfeasible for Betty to check or confirm this stranger's reliability as a source of testimony. She has not met this person before and is unlikely to seek his testimony again—so it would be difficult for her to assess his reliability, and unnecessary for her to try to do so. So, if Betty accepts his testimony, she does so on the basis of *prima facie* trust. This is not the case for (2). Since Ashish is her close friend, and it is possible that Betty might seek his testimony again, it would be both desirable and feasible for Betty to come to a reasoned assessment of Ashish's trustworthiness.

In many ways, search engine algorithms as testifiers resemble the case in (2). Google has been, for the past decade, among the most visited websites in the world—with about 88 billion unique visits in 2020 (SimilarWeb, 2020). Google's algorithms are consulted widely and often. However, these algorithms operate in the background of our experience while searching online. Seemingly oblivious to the work they are doing to bring content to us, we reflexively accept their bent testimony. In this way, it seems like we treat Google's algorithms with the same *prima facie* trust that we would extend to the stranger in (1). Given how often we seek their testimony, and how much influence they have over our beliefs, this seems irrational.

This is especially so because, as several scholars have suggested, search engine algorithms may not always provide reliable testimony. Introna and Nissenbaum (2000) famously argued that search engines "systematically exclude (in some cases by design and in some, accidentally) certain sites and certain types of sites in favour of others, systematically giving prominence to some at the expense of others" (p. 169). Since their early contribution, we have learnt a great deal about what these systematic exclusions look like. Safiya Noble, for example, provides a comprehensive account of the biases against marginalized groups (especially women of color) that are embedded in search engine algorithms. She discusses how, for instance, a search with the term "black girls" is more likely to return sexually explicit content as compared to a similar search for "white girls" (Noble, 2018). Other scholars have also provided evidence that search engines disproportionately promote paid content (and make it harder to distinguish sponsored listings from non-sponsored ones), limit our exposure to points of view that differ from our own, and promote sensationalizing content that is more likely to "generate clicks" (Diaz, 2008; Grimmelmann, 2011; Vaidhyanathan, 2012; Verma et al., 2020). Furthermore, we know that many contemporary organizations devote considerable time and resources to "search engine optimization"—to manipulate their website designs such that they might feature higher up on Google's search results. All this suggests that our *prima facie* trust that search engine algorithms reliably bring forward true and relevant content may be unwarranted. Multiple, well-resourced studies have shown us how search engine algorithms get things wrong, systematically and often, and as such, the reliability of their algorithms as a source of testimony seems to be an open question.

It is often said that it is difficult for users to assess exactly how reliable or unreliable search engine algorithms are (Miller & Record, 2013; Pasquale, 2015). These algorithms are often kept secret, and even when they are not, one would need deep technical expertise to understand precisely how they work. The average user of Google cannot be expected to come to a full understanding of the inner workings of search engine algorithms, to determine whether or not their testimony is reliable.

While true, this is beside the point. We frequently make assessments about the reliability of other people's testimony, even when we cannot fully understand the inner workings of their mind. For instance, if someone gave us testimony that we later found to be false, or if we learned that their testimony is biased—we would at least update our beliefs about their reliability, and we may be suspicious about accepting their testimony again in the future. Google seems to be immune to such suspicions. When Google gives us an inaccurate answer, we would just seek testimony from them again, perhaps this time with different search input. When evidence about embedded biases in Google's search algorithms comes to light, we all but ignore it. It is this double standard, in particular, that needs to be remedied.

This paper will stop short of making specific and detailed policy or legal recommendations about how we ought to treat search engines. In part, this is because social and computer scientists still have much to tell us about how exactly search engine algorithms are designed and how we interact with them. Even so, we would suggest that such recommendations might derive directly from the normative and legal expectations we have from ordinary testifiers. To illustrate, if we think that a particular testifier is consulted widely and often, we might want to alert them to any errors we find and help improve their reliability. For Google, this might entail allowing users to give feedback and suggestions when they find the content at the top of a search result to be untrustworthy or inaccurate. For widely consulted testifiers, we might also want to invest time and resources towards improving our understanding of how and why they get things wrong. If so, we would want to better support the various scholars and activists working to uncover the biases embedded in search engine algorithms, and those working to develop policy and technical interventions that mitigate these biases and their social consequences.

In sum, the implications of our view are underpinned by two general ideas. First, understanding how exactly the testimony of search engine algorithms is bent can help us control the influence that this algorithmic testimony might have on our beliefs. Rejecting this bent testimony, if we deem that necessary, entails rejecting the notion that search engine algorithms bring to fore content that is more likely to contain true claims. Second, and perhaps more importantly, we should treat the bent testimony of Google's algorithms like we would the testimony of anyone else—especially when it comes to handling new evidence that bears on our appraisal of how reliable these algorithms are as testifiers.

5 Conclusion

We have argued here that search engine algorithms provide us with a kind of testimony when they bring to fore some pieces of content for us to engage with and push behind others. This testimony is "bent," because:

(1) We treat these algorithms as if they are *recommending* to us the content that they feature at the top of a search results list, trusting that this content is more likely to contain true claims.

(2) There are disputed norms of communication about whether the recommendation of a piece of content counts as an assertion of its claims.

An understanding of this mechanism of bent testimony shows us how to control our reliance on it, if we so desired. Decreasing our reliance on this bent testimony entails decreasing our credence in the belief that the content ordered at the top of a search engine is any likelier to contain true claims. Further, we have argued that we ought to treat search engines as if they were testifiers. By having comparable expectations between search engines and ordinary testifiers, we would be able to pursue policy and legal interventions that befit the outsized role that these search engines seem to play when we acquire beliefs online.

Finally, although we restricted our discussion to the content-curation algorithms used in search engines, they are far more ubiquitous. Similar algorithms operate in the background of our experiences almost throughout the Internet—on entertainment platforms (such as YouTube or Netflix), shopping platforms (like Amazon or Book Depository), and social media platforms (like Facebook or Twitter). When we acquire or alter beliefs by engaging with content on these platforms, these beliefs might be similarly reliant on the bent testimony of the algorithms that bring that content to us. We think it is possible to extend our account here to other types of content-curation algorithms as well. For this reason, we hope that future researchers will embark on pursuing such an extension, and in turn, develop a more general account of the testimony provided by algorithms that curate and order content.

References

Adler, J. (2017). Epistemological problems of testimony. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017). Metaphysics Research Lab, Stanford University. Retrieved 15 Jan 2021, from https://plato.stanford.edu/archives/win2017/entriesestimony-episprob/

Bermúdez, J. L. (2007). *Thinking without words*. Oxford University Press.

Boutin, P. (2011, May 20). Your Results May Vary. *Wall Street Journal*. Retrieved 15 Jan 2021, from https://online.wsj.com/article/SB10001424052748703421204576327414266287254.html

Chaney, A. J. B., Stewart, B. M., & Engelhardt, B. E. (2018). How Algorithmic confounding in recommendation systems increases homogeneity and decreases utility. *ArXiv:1710.11214 [Cs, Stat]*. https://doi.org/10.1145/3240323.3240370

Coady, C. A. J. (1992). *Testimony: A philosophical study.* Clarendon Press, Oxford, England.

De Cremer, D., McGuire, J., Hesselbarth, Y., & Mai, M. (2019). Social intelligence at work: Can AI help you to trust your new colleague. *Harvard Business Review*. 4 June. Retrieved 10 Jan 2021, from https://hbr.org/2019/06/can-algorithms-help-us-decide-who-to-trust

Diaz, A. (2008). Through the Google goggles: Sociopolitical bias in search engine design. In *Web search* (pp. 11–34). Springer, Berlin, Heidelberg.

Epstein, R., & Robertson, R. E. (2015). The search engine manipulation effect (SEME) and its possible impact on the outcomes of elections. *Proceedings of the National Academy of Sciences, 112*(33), E4512–E4521. https://doi.org/10.1073/pnas.1419828112

Freiman, O., & Miller, B. (2020, May 7). *Can artificial entities assert?* The Oxford Handbook of Assertion. https://doi.org/10.1093/oxfordhb/9780190675233.013.36

Google Privacy & Terms. (n.d.). *Frequently asked questions*. Retrieved 24 January 2022, from https://policies.google.com/faq?hl=en-US

Google Search. (n.d.-a). *Ranking results*. Retrieved 24 January 2022, from https://www.google.com/intl/en_in/search/howsearchworks/how-search-works/ranking-results/#context

Google Search. (n.d.-b). *Rigorous testing*. Retrieved 24 January 2022, from https://www.google.com/intl/en_in/search/howsearchworks/how-search-works/rigorous-testing/

Google Search. (n.d.-c). *Search quality rating guidelines*. Retrieved 24 January 2022, from https://static.googleusercontent.com/media/guidelines.raterhub.com/en//searchqualityevaluatorguidelines.pdf

Google Search. (n.d.-d). *Features*. Retrieved 24 January 2022, from https://www.google.com/intl/en_in/search/howsearchworks/features/

Grimmelmann, J. (2011). Some Skepticism About Search Neutrality. In B. Szoka, A. Marcus (Eds.), *The Next Digital Decade: Essays on the Future of the Internet* (pp. 435 - 459). TechFreedom, Washington DC., USA. Retrieved 15 Jan 2021, from https://papers.ssrn.com/abstract=1742444

Gunn, H. K., & Lynch, M. P. (2018, September 3). *Googling*. The Routledge Handbook of Applied Epistemology; Routledge. https://doi.org/10.4324/9781315679099-4

Hillis, K., Petit, M., & Jarrett, K. (2012). Google and the Culture of Search. *Routledge*. https://doi.org/10.4324/9780203846261

Introna, L. D., & Nissenbaum, H. (2000). Shaping the Web: Why the politics of search engines matters. *The Information Society, 16*(3), 169–185.

Jansen, B. J., Spink, A., & Saracevic, T. (2000). Real life, real users, and real needs: A study and analysis of user queries on the web. *Information Processing & Management, 36*(2), 207–227. https://doi.org/10.1016/S0306-4573(99)00056-4

Jansen, B. J., & Spink, A. (2006). How are we searching the World Wide Web? A comparison of nine search engine transaction logs. *Information Processing & mAnagement, 42*(1), 248–263.

Jansen, B. J., Zhang, M., & Schultz, C. D. (2009). Brand and its effect on user perception of search engine performance. *Journal of the American Society for Information Science and Technology, 60*(8), 1572–1595. https://doi.org/10.1002/asi.21081

Joachims, T., Granka, L., Pan, B., Hembrooke, H., Radlinski, F., & Gay, G. (2007). Evaluating the accuracy of implicit feedback from clicks and query reformulations in web search. *ACM Transactions on Information Systems (TOIS)*, 25(2), 7-es.

Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199219162.001.0001/acprof-9780199219162

Latour, B. (2005). *Reassembling the social: An introduction to actor-network-theory*. Oxford University Press.

Latour, B. (2012). *We Have never been modern*. Harvard University Press.

McKinsey (2011). *The impact of internet technologies: Search*. Retrieved 24 January 2022, from https://www.mckinsey.com/~/media/mckinsey/dotcom/client_service/high%20tech/pdfs/impact_of_internet_technologies_search_final2.aspx

Miller, B., & Record, I. (2013). Justified belief in a digital age: On the epistemic implications of secret Internet technologies. *Episteme, 10*(2), 117–134. https://doi.org/10.1017/epi.2013.11

Nentwich, M., & König, R. (2012). *Cyberscience 2.0: Research in the age of digital social networks*. Campus Verlag GmbH, Frankfurt-on-Main.

Nickel, P. (2013). Artificial speech and its authors. *Minds and Machines, 23*(4), 489–502.

Noble, S. U. (2018). Algorithms of oppression: How search engines reinforce racism. *NYU Press*. https://doi.org/10.2307/j.ctt1pwt9w5

Oulasvirta, A., Hukkinen, J. P., & Schwartz, B. (2009). When more is less: The paradox of choice in search engine use. *Proceedings of the 32nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 516–523. https://doi.org/10.1145/1571941.1572030

Page, L. (2006). *Method for node ranking in a linked database* (United States Patent No. US7058628B1). https://patents.google.com/patent/US7058628B1/en

Pagin, P. (2016). Assertion. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2016). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/win2016/entries/assertion/

Pan, B., Hembrooke, H., Joachims, T., Lorigo, L., Gay, G., & Granka, L. (2007). In Google We Trust: Users' Decisions on Rank, Position, and Relevance. *Journal of Computer-Mediated Communication, 12*(3), 801–823. https://doi.org/10.1111/j.1083-6101.2007.00351.x

Pasquale, F. (2015). *The black box society*. Harvard University Press.

Rescorla, M. (2009). Chrysippus' dog as a case study in non-linguistic cognition. In R. Luiz (Ed.), *The Philosophy of Animal Minds* (pp. 52–71). Cambridge University Press.

Rini, R. (2017). Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal, 27*(S2), 43–64. https://doi.org/10.1353/ken.2017.0025

Roesler, P. (2021, October 16). *How Google's Latest Search Redesign May Impact Your Business*. Inc. Com. Retrieved 22 Jan 2022, from https://www.inc.com/peter-roesler/how-googles-latest-search-redesign-may-impact-your-business.html

Rothman, L. (2018). *20 Years of Google has changed the way we think. Here's how, according to a historian of information*. Time. Retrieved 24 Jan 2022, from https://time.com/5383389/google-history-search-information/

Searle, J. R. (1969). Speech acts: An essay in the philosophy of language. *Cambridge University Press*. https://doi.org/10.1017/CBO9781139173438

SimilarWeb. (2020). *Google.com Traffic Statistics*. Retrieved 15 November 2020, from http://similarweb.com/website/google.com/

StatCounter. (n.d.). *Search engine market share worldwide*. StatCounter Global Stats. Retrieved 8 November 2020, from https://gs.statcounter.com/search-engine-market-share

Tollefsen, D. P. (2009). *Wikipedia* and the epistemology of testimony. *Episteme, 6*(1), 8–24. https://doi.org/10.3366/E1742360008000518

Tuomela, R. (1992). Group beliefs. *Synthese, 91*(3), 285–318.

Vaidhyanathan, S. (2012). *The googlization of everything: (And why we should worry)*. University of California Press.

Van Deursen, A. J., & Van Dijk, J. A. (2014). *Digital skills: Unlocking the information society*. Springer

Verma, S., Gao, R., & Shah, C. (2020). Facets of fairness in search and recommendation. In L. Boratto, S. Faralli, M. Marras, & G. Stilo (Eds.), Bias and Social Aspects in Search and Recommendation (pp. 1–11). Springer International Publishing. https://doi.org/10.1007/978-3-030-52485-2_1

Westerwick, A. (2013). Effects of sponsorship, web site design, and Google ranking on the credibility of online information. *Journal of Computer-Mediated Communication, 18*(2), 194–211. https://doi.org/10.1111/jcc4.12006

Weisberg, J. (2011, June 10). *Eli Pariser's the filter bubble: Is web personalization turning us into solipsistic twits?* Slate Magazine. Retrieved 15 Jan 2021, from https://slate.com/news-and-politics/2011/06/eli-pariser-s-the-filter-bubble-is-web-personalization-turning-us-into-solipsistic-twits.html

Wheeler, B. (2020) "Reliabilism and the Testimony of Robots." *Techné: Research in Philosophy and Technology, 24*(3), 332–356.