

Google Epistemology

Hanna Gunn and Michael P. Lynch

Forthcoming: *The Routledge Handbook of Applied Epistemology*. D. Coady (Ed.) (Routledge: London) 2018.

Introduction

In a recent *New Yorker* cartoon, a man is fixing a sink. His partner, standing nearby skeptically asks, “Do you really know what you are doing, or do you only google-know?”

This cartoon perfectly captures the mixed relationship we have with googling, or knowing via digital interface, particularly via search engines. On the one hand, googling is now the dominant source of socially useful knowledge. The use of search engines for this purpose is almost completely integrated into many of our lives. On the other, the point the cartoon is making resonates with nearly all of us: users often recognize that there are risks and trade-offs associated with gaining certain kinds of information via online search.

These facts about googling make it particularly interesting to the applied epistemologist. Our practices involving search-engines not only have a distinctive character, that character puts some traditional epistemic questions in a new light. This paper will examine two of those questions. The first concerns the extent to which googling raises problems similar to familiar quandaries surrounding testimonial knowledge. The second—and more radical—concerns whether googling is a type of distributed or extended knowledge.

Googling as a Mode of Inquiry

Finding information via a search engine is a ubiquitous feature of modern life. Google, as a company, is now so dominant that the near-universal term for such activity is “googling”. For most people, googling is just how you acquire socially useful knowledge.

In this paper, we’ll use “googling” to cover a particular method or *mode of inquiry*. By a “mode of inquiry” we mean a process whose point is to answer a particular

question or questions. Such questions needn't be, and often aren't, explicitly stated, but implicit in the inquiry itself. Importantly, following common usage, we take "googling" to name the process of finding information online. As such, we take it to include not just the act of using a search engine, *but the whole process of acquiring information* over the Internet, including the search, the sites reached by that search, and the answer to the question(s) which is/are the target of the inquiry. Thus, to say you "googled it" is to say that you acquired the relevant answer online via a search.

The sheer ubiquity of googling hints at one of its more epistemically interesting features. Googling is deeply integrated into our epistemic life in at least three respects.

First, for many questions, googling is our primary or "go-to" way of finding answers. A quick way of illustrating this fact is to simply reflect how you might go about finding answers to any of the following questions, assuming normal working Internet access:

- What are the possible causes for your car not starting?
- Where is the closest auto-parts store?
- How late is that store open?
- What opinions do people have about local auto mechanics?

Whether or not these particular questions matter to you, it is clear to most users of the Internet that the first thing one would do if one wanted to know the answer to such questions is to engage in an Internet search. Indeed, for most people, searching online happens without much forethought. It is just the obvious, immediate first step in answering almost any question about the social world; for those questions, it has a kind of priority for many people.

Modes of inquiry have causal priority over one another when you can't engage in one mode without engaging in another. Thus one can't read without using vision or some other sense (touch, for example). Modes have epistemic priority over one another when you are justified in using the second only if you are justified in using the first. But modes can also enjoy what we can call *a priority of use* relative to certain classes of questions. Priority of use is a kind of statistical priority. *A mode has priority of use over another, relative a given class of questions, just when it is, other things being equal, the first mode or method used in order to answer questions of that class.* Consider, for example, perception with regard to questions about our immediate physical environment. If we wish to know whether it is safe to cross the road, we first try to look and see; failing that, we might ask someone nearby. If we want to know whether the phone is ringing, we first try listening for it, etc. Perception often has causal, epistemic, and use priority over other modes of knowing. But as the above experiment indicates, googling has priority of use for many ordinary questions. We

can use other modes of inquiry to answer these sorts of questions, but we typically try googling first.

Second, googling is cognitively seamless – that is, we look through the interface, or don't treat it as cognitively salient. This itself is a consequence of ubiquity and epistemic priority of use together with the sheer speed of Internet searches. In this way too, googling is like certain other modes of inquiry, like reading street signs or even perceiving the world around you. It is a process that happens easily enough that we “see” the information we are seeking without always noticing how we are going about it.

Third, and perhaps most strikingly, most people treat googling with default, *prima facie* trust¹. How many times has someone you know announced on some point of fact and everyone else in the room races to their phones to verify or falsify it? We routinely use googling to trump other forms of inquiry, even to question experts. Of course, almost no one will admit to trusting everything on the Internet, and while people are no doubt more credulous than they care to admit, no one really does. We all know that googling can lead us astray. But that doesn't stop us from using it routinely, nor regarding it as essentially reliable on a broad range of topics. Here again, googling is akin to perception. We know that perception can mislead us, but we nonetheless treat it with *prima facie* trust. So too with how we acquire information online; we used to say that seeing is believing. Now googling is believing.

The above three points suggest that googling, like perception, is a mode of inquiry that is deeply integrated into our epistemic life. Of course, googling is unlike perception in numerous ways, two of which will form the basis for the rest of the discussion of this paper. First, and perhaps most obviously, when googling, we are not interacting with physical objects but consulting a source for sources that are *distributed* across a virtual space. And second, googling, unlike perception, is radically dependent on other people and processes that are not our “own”. 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.

¹ As we will discuss later, it is also true that we trust different types of internet sources to different degrees. That is, if our searching took us to both a personal blog post and a peer-reviewed academic paper it is likely we would trust the latter more. However, given most people's searching behaviours, unless an individual is intentionally trying to corroborate or compare results when they search, we are much more likely to just go with the top result (Joachims, et al., (2005), Beitzel et al., (2007), King, (2008)). Thus, it is not inconsistent to suggest that search results will be treated with a *prima facie* level of trust, even when some situations might cause us to be more particular about how much trust we afford to different types of results.

Taken together, the above points indicate what we believe is distinctive about googling from an epistemological point of view. It is a mode of inquiry that is at once closely integrated (like perception) and yet distributed and reliant on others (like testimony). It is this combination of features that makes googling of particular interest to the social epistemologist.

We should note at the outset that there is one further feature of googling that makes it particularly distinctive: it is a *preference-dependent* mode of inquiry. From a certain level of abstraction, Facebook, Google, and most of our Apps, search engines and social platforms all work in the same basic way, different algorithms aside. They attempt to track people's preferences by way of tracking their likes, their clicks, their searches, or their friends. That data is then analyzed and used to predict what a given person's current and future preferences will be. But it is also used to predict what sort of information you—and crucially, those similar to you—will find interesting, what posts you will like, and what links you will most click. The results of this preference-aggregation are then served up on websites you visit. That's how the magic happens. It is why those shoes you were thinking of buying are being advertised on your Facebook feed. And that's also what makes googling both one of the most efficient modes of inquiry humans have ever produced and one of the riskiest, from the epistemic point of view².

In the next section, we discuss how those risks bear on the question of whether it is rational for us to trust what we learn by googling.

In Googling We Trust

Trust in the testimony of others plays an important role in our epistemic lives, since the amount of knowledge we need to rely on in daily life far exceeds what we can acquire personally. With the Internet serving as a major source of testimonial knowledge (and seemingly equally for low-risk queries like “how to fold a fitted sheet” to high-risk queries like “can I feed my newborn peanuts?”) it is important to explore the epistemic vulnerabilities googling presents. Trusting, in general, is an activity that makes us vulnerable to other people, and trusting in googling as many of

² Preference-tracking can be manipulated while manipulating at the same time (Lynch, 2016). That's important because search prompts and the ranking of results effects what people know as much what turns up in the encyclopedia (back when there were such things). Here's a quick example: type, “is climate change...” into Google and you are likely to have Google's Complete program helpfully finish your thought with “...a hoax?” (That answer often comes up before “...a fact?”). Hit that, and you might get, e.g., “Top Ten Reasons Climate Change is a Hoax” from a nice little outfit calling itself “globalclimatescam”. (One can learn so much on the Internet!).

us do on a daily basis brings many of the same problems. That being said, there are several features of googling that should make us wonder if it poses unique epistemic vulnerabilities that testimony from other sources does not.

For one, the Internet makes it very easy for many sources to remain anonymous and so removes the risks to reputation that face-to-face testimony may bring (and the security that can provide for the inquirer). Another related problem is the credibility of the sources that Google directs us towards, and in particular our ability to verify the credibility of those sources. In an ideal world, we would be able to rely on genuine experts for the information that we need – but googling doesn't necessarily take us to experts for answer to our questions. When we google we are directed to forums, Internet encyclopedias, social media, news websites, online stores, blogs, videos, to name a few possible sources. We can immediately see two potentially unique problems with googling for testimony: we regularly don't know who we are talking to and we typically don't have access to information about their trustworthiness.

In the discussion that follows we explore a number of similarities between googling and verbal testimony, and point out where we think there may be unique issues raised by taking googling as a mode of inquiry. We consider in more depth the nature of the sources of testimony that googling directs us towards, and how these various sources present us with different epistemic problems. We also consider the issue of credibility markers available to us when we google and how this may be different from existing discussions in the literature on related discussions about recognising experts.

Googling and Prima Facie Trust

Placing trust in others' testimony is unavoidable. But it isn't always warranted. So that raises a natural theoretical question: under what circumstances is such trusting rational?

We mentioned earlier that googling is arguably treated by us with a default, *prima facie* trust – much like the testimony of others. Put more carefully, absent obvious defeaters for what someone says, we tend to take someone's testimony that *p* as *prima facie* evidence for the truth of *p*. And, we contend, we typically do the same thing when we google answers to our questions. If so, then we can extend the natural question above and ask whether are rational in putting *prima facie* trust into results of our googling³? The connection we draw here between googling and the philosophy of

³ Other approaches to this question have been taken in literature, i.e., one's that do not appeal directly to the existing literature on testimony. Simpson (2012) is a notable example. Simpson

testimony can also be found in a related topic of testimony and professional library and information scientists by Fallis (2004). Fallis' concern is to provide criteria for evaluating the accuracy of content on the Internet that information professionals can draw on to better teach people how to evaluate information. Our focus is primarily from the perspective of the individual user and the information she requires to make an assessment to trust or not when she googles.

There are two dominant positions in the existing literature on testimony that bear on this question: reductionism and anti-reductionism. According to reductionists, the epistemic merit of testimony is reduced to the epistemic merit of sources such as perception, memory, and inductive inference. Our trust in testimony is rational when it is based in these other sources. According to one such view, therefore, we are rational in trusting testimony by virtue of our past experiences of the *general* reliability of testimony – reliability that we've presumably tracked via other evidential sources ("global reductionism"⁴). Alternatively, it may be because we are justified in accepting the testimony of *particular* individuals on *particular* topics due to beliefs about their reliability on those topics ("local reductionism"⁵). Note, however, that in contrast to the common-sense understanding of trust as *prima facie* trust sketched above, reductionist positions require positive evidence for trustworthiness; that is, trust is earned, not assumed.

Anti-reductionists, by contrast, argue that testimony itself is a basic form of evidence that provides universal, *prima facie* justification for accepting the testimony of speakers. What enables us to trust testimony, understood typically as factual assertions, is an epistemic principle. One such principle draws a connection between our self-trust in our own opinions and faculties and the similar faculties and beliefs of others⁶. We don't typically require independent justification of the reliability of our faculties in order to form beliefs, we by-and-large trust them as reliable (putting aside cases where we have medical or otherwise reasons to distrust them). So, when faced with others whose beliefs have a similar aetiology and who also have similar

proposes three criteria for evaluating Google as an "epistemic tool": timeliness, authority prioritisation, and objectivity. These are used to expand on Goldman's (1999) proposal of the "veritistic value" of search engines as determined by how well they meet his conditions of recall and relevance. Our approach and Simpson's come together over concerns about what we have termed preference dependence. For Simpson the personalisation of our searches poses a threat to his objectivity requirement: "when their rank ordering represents a defensible judgment about the relative relevance of available online reports" (434). Put slightly differently, a search engine is objective when it provides a balanced selection of the relevant evidence for a query and does not only return results that support a particular side. This concern over objectivity and personalisation is applied to support an argument for government regulation of search engines.

⁴ For example, Van Cleve (2006). See David Hume's discussions of testimony provide examples of a historical reductionist position (Hume, 2010 [1740], T:1.3.4).

⁵ For example, Elizabeth Fricker, (1994).

⁶ For example, Foley, (2001). See Reid (1975) for an early anti-reductionist position.

(reliable) faculties, we must, in order to avoid inconsistency, treat their testimony (the outcomes of their exercise of these faculties and opinions) with a similar level of default trust. Other anti-reductionists argue that we can afford *prima facie* trust to the testimony of others because, they argue, there is a special kind of epistemic warrant that comes with something being presented as true⁷.

To extend either reductionist or anti-reductionist positions to cases of googling will depend on the degree to which googling and verbal testimony are similar. And in order to evaluate that, we need to look more closely both at the nature of sources we find when we google (i.e., our candidate testifiers) and the testimonial contexts we come across.

Testifiers and Trolls

With these positions in mind, we can return to the question with which we began; namely, whether we are rational in giving *prima facie* trust to googling as a source of beliefs. What we find is that differences between traditional examples of testimony and googling really matter. In most discussions of testimony, philosophers explore cases of (a) one-on-one verbal testimonial exchanges, where (b) “testimony” is typically understood as the speech act of asserting. Both points are relevant when considering googling because neither are guaranteed (or, indeed, the norm) for many sources we find and then form beliefs from on the Internet.

Two common sorts of cases we find in the existing testimony literature involve novices looking for expert testimony in areas they know little about (e.g., anthropogenic climate change), and individuals needing to ask a stranger for advice (e.g., as a tourist asking for directions). While we may be engaged in the same sorts of inquiry on the Internet (i.e., for expert opinions or for directions) we are often not interacting with one other individual⁸. In many cases our testifiers may be community forums or community compiled encyclopedias⁹. In still more cases, we are not using a human source at all¹⁰. This range of possible testifiers—and in some cases a range of

⁷ Burge, (1995). Similar arguments are made by appeal to pragmatic principle like Grice’s “Cooperative Principle” (1989), the general idea behind these views being that the testimonial setting itself entitles the listener to believe what is said.

⁸ We leave aside here examples of trusting online stores about their own products. For one, the credibility of businesses touting their own wares is not new with the Internet and so doesn’t present new issues. On the other, concerns about the credibility of reviews are captured in the other cases we discuss here and in the next section on credibility markers.

⁹ For example, Wikipedia or Stack Exchange.

¹⁰ For example, search results delivered automatically by Google, or alternatively asking natural language questions to computational knowledge engines like Wolfram|Alpha.

candidate testifiers for the same query—is one of the unique features of google inquiry. We take it that these two features have significance for thinking about reductionist and anti-reductionist positions that would aim to show how we are justified in trusting our google searches¹¹.

On a reductionist reading, we might say that we are warranted in trusting our google results because of the general or local reliability of googling in the past. This would seem to square well with one feature of preference ranking in search engine results – the most popular results are the ones that our search engines direct us towards. That being said, and given the varied types of sources we come across when we google, a local reductionist position may be more appropriate. It seems plausible that one should not treat community maintained wikis in the same way as one treats a medical blog or a post on social media. In order to afford *prima facie* trust, then, it might be epistemically preferable to localise to either types of sources (e.g., wikis, scientific blogs, news websites) or to particular testifiers (e.g., Wikipedia, Phil Plait's astronomy blog, or The Guardian).

Whether we find reductionist positions plausible when applied to googling will hinge in part on whether we think that they are too demanding – that is, whether they presume too high an evidential standard for determining reliability. How much experience do we need to have with googling to give us rational trust, either locally or globally? If you think that we can, at least sometimes, be rational in trusting some online searches even when we lack wide experience, you might find the reductionist account too conservative. On the other hand, perhaps we *should* demand or desire that googling, even to a local level, have a potentially high evidential demand like this to justify affording default trust to their results. The nature (i.e., individuals or collectives/groups) and variety of sources that we come across when we google is one reason we might desire these high evidential demands. Reductionist accounts of testimony are open to the critique that they presume too great an evidential standard for trust: global reductionism would seem to require a great deal of leg-work in establishing the reliability of past-testimony given the ratio of the number of claims we are exposed to via testimony to the number to claim we check; and, similarly, the narrower evidential demand of the local reductionist is more demanding that it first appears.

Indeed, the anti-reductionist position seems to face some serious challenges in the

¹¹ As stated above, typically philosophers take assertions to be the speech act in testimony cases that lead to testimonial knowledge. While we may want to say that our non-human sources and community compiled encyclopedias, for example, are presenting information to us as true, it's not immediately clear that they (qua non-human or non-individual) are performing assertions. In addition, in many cases where the source may be both human and individual, there are a great many types of media on the Internet we may be directed towards from our googling and these need not involve any assertions.

case of googling. Appealing to general wisdom for a moment, we might wonder how we should square the common refrain of “You can’t trust what you read on the Internet” with the anti-reductionist view about googling that would give the verdict that we are justified *because* we found it googling. A self-trust view like that mentioned above does not appear to transfer over to the case of googling very easily. Recall that testimony has a *prima facie* level of trust on these views because of a presumption of trust in our own opinions and faculties. If we know enough about the author of the results we find when we google, then we may be warranted in trusting those results. The issue is that we often do not know much about the authors of such results (e.g., forum contributions or blogs) or the author may not be an individual (e.g., a community wiki). But more problematically for self-trust views, the author may not be a person at all (e.g., Google’s automated responses to many questions including currency conversions and mathematical questions).

For similar reasons, appealing to Burge or Gricean style anti-reductionist position for googling is not straightforward either. The general idea of these positions is that we are warranted in trusting because of norms that apply to *verbal* testimonial settings, namely ones that are tied to truthfulness and speech acts of asserting. Unlike verbal testimony where we typically take testimony to be a voiced assertion, googling arguably directs us towards a much larger range of speech acts and speakers than just individuals offering up explicit opinions. We might find an answer to our query in a picture, a forum, a video, reading news websites, reading a blog, or from an automated system. In some such cases, the information we draw on is not being presented as true and thus it doesn’t seem we would be entitled in the way that Burge, for example, argues.

But even in cases where the results of googling can be taken as claims that would be verbally asserted by someone, it is not clear that the same conversational norms that guide these anti-reductionist accounts necessarily apply on the Internet.

Consider the phenomenon of Internet “trolls”. Trolls are individuals who partake in forums or comment sections (or anywhere where individuals can post messages to others) with the intention of angering, fooling, or otherwise making insincere contributions to the discussion. Trolling contributions vary in their obviousness, but in many cases they are intentionally deceptive in order to get the desired rise out of other participants in the discussion. Consequently, it can be difficult to detect trolls. If it is highly likely that trolls are contributing to the source our googling directs us towards, then we would be wise not to treat the information in the source with a default level of trust. In response to this concern, we might take trolls to be merely a new instantiation of liars or deceivers (or even bald faced liars), something that philosophers of language already have treatments of¹². This would make Internet

¹² See Bok (1978), Carson (2010), Saul (2012).

trolls not a new problem, but still one that poses as epistemic threat to us when we google.

In sum, there is some intuitive appeal for reductionist defenses of *prima facie* trust in our googling results, even with the potentially high evidential demand that these accounts bring. In cases of verbal testimony, there are factors that can bias us against (or in favour of) certain potential testifiers¹³. Consequently, our ability as individuals to perform the relevant tasks for assessing the reliability of potential testifiers can be biased in unreliable ways. These are serious issues for social epistemologists concerned about testimonial knowledge and trust. Likewise, the Internet brings with it a range of biasing factors that can thwart our ability to place trust appropriately.

Selecting Our Google Sources

A number of different factors determine the search results that we see when we look for information through a search engine, but it often includes our personal search history and the searches of others, i.e., facts about the preference-dependence that we introduced earlier¹⁴. It is plausible that the ranking of search results brings with it a degree of reliability: information that is accurate, true, or otherwise reliable will be more likely to be shared or cross-linked and thus will rise in the rankings. There are other factors, though, that can lead to a result being higher in the rankings including “viral” information, memes, or simply being the product of a celebrity (of, film, music, television, YouTube, etc., fame) that can cause some pages to rank highly when googling. In these cases we have popularity-because-popular, rather than popularity-because-reliable as in the former cases. These facts select for certain testifiers over others, and in ways that may not be tied to truth or reliability¹⁵.

Of course, we may not be googling for information via a search engine, but instead by posting on forums or through social media and these sources run the risk of being

¹³ See Fricker (2007), Medina, (2013), Mills, (2007).

¹⁴ Simpson (2012, 437-438) provides a useful discussion of two ways our searches can be personalised, what he terms “individual personalisation” and “profile personalisation”. The difference marks one we mentioned in introducing preference dependence. Search results can be personalised to an individual and their personal browsing history, i.e., individual personalisation. Alternatively, search results can be personalised in line with a “profile” that matches the browsing behaviour of sufficiently similar other people. In this case, the googler will see results tailored to people “like them” according to the profile they match, hence “profile personalisation”.

¹⁵ Indeed, research into the clicking behaviour of individuals using search engines supports the conclusion that individuals are likely to click the top search result even when the next result may be (quite plainly) the more relevant option. One conclusion of the study was that individuals have a “trust bias” for top results because of an assumption that search engines are presenting the most relevant results (Joachims et. al., 2005).

undermined by other biasing factors. Consider social media platforms and the way in which our networks are there are formed of our friends, and that our friends are typically people with a similar values and worldviews as ourselves. These facts about our friend circles mean that one is likely to find testifiers for opinions and beliefs that support what one already believes given facts about how we tend to form friendships¹⁶. This can bias the testimony that we find, but it can also lead to a biasing of the primary sources that we are advised to look at when we ask on social media¹⁷. For example, those on one side of the alternative medicine debate are likely to send recommendations and links to outside sources that support their own beliefs. So, the advice we get about what to look at to answer our queries on social media (and potentially forums) can be highly selective in ways that might bias us in one direction¹⁸.

Thus, googling poses us with a selection problems for testifiers because of the way in which searches are biased, and another in the way that we gather information from those we already trust.

Are these new problems with googling? We would suggest yes and no. Obviously prior to being able to google, the results of our inquiry could not be subject to the preference-dependence biasing effects that search engines cause¹⁹. That being said, other social factors likely bias the popularity of the information we hear about through other traditional media forms²⁰ and these are also not necessarily because of any epistemically virtuous qualities on behalf of the media in question. The social media discussion is also not all that new, friend circles (or family as the case may be) are naturally selective, and the biasing effects from social media are probably best understood because of pre-existing facts about non-social-media friendships and

¹⁶ See Quattrociocchi et al., (2016) for a discussion of social media and confirmation bias.

¹⁷ Many terms have come about for these kinds of interactions between automatic personalisation mechanisms and social dynamics on social media in particular, including both “filter bubbles” and “echo-chambers”. See Pariser (2011) for a discussion of filter bubbles and Quattrociocchi et al., (2016) on echo chambers.

¹⁸ One medium we do not discuss here in detail are Internet blogs. Given the sheer number of blogs created by a vast range of individuals (and thus, representing a potentially equally large number of perspectives and opinions) blogs are an important area for epistemic consideration for googling. For a discussion of the interplay between and epistemic comparison of blogs and professional (particularly, political) journalism see Coady (Chapter 6, 2012).

¹⁹ Recall that preference dependence is used to describe the processes by which companies use our online behaviour to shape our online experience so that we have a personalised Internet browsing experience. For example, using targeted advertising, altering search engine result listings (order and type), and the content of our social media feeds. This is distinct from one’s personal preferences causing one to buy a particular newspaper or believe a particular story and not another. Though these personal preferences also lead to many of the same epistemic problems we raise, they have very different sources and thus require different explanations of their epistemic import.

²⁰ Because it sells, because of institutional/discipline related trends or social make-ups, etc.

associations. In both cases, however, the near-ever-present ability to google, the massive body of information that makes up the Internet, and the speed of results, present us with a medium that has the potential to greatly exaggerate these problems.

It is this final point that makes the case of testimonial knowledge on the Internet unique. While the types of biases may not be exclusively tied to the use of googling, the potential for these biases to have a bearing on our epistemic lives is much greater given the way that preference-dependence intentionally shapes this epistemic environment. Part of the bind, of course, it is just these same processes of preference-dependence that help make the Internet possible to navigate and useful at all.

Credibility Markers

The risk of false information, liars, or bullshit (ala Frankfurt²¹) is always present when dealing with testimony. As we have just discussed, it's not necessarily that the Internet and googling create new epistemic problems for us but they run the risk of exaggerating or making us open to existing threats more often. This is because of the easy availability of googling and the huge amount of information on the Internet to support nearly any position one wants. For instance, it is just as easy to find very complicated explanations of why the moon-landings were a hoax as it is to find complicated explanations of how they couldn't have been a hoax. The consequence is that our need to assess the credibility of our sources when googling is especially important.

The relevant existing problem in social epistemology explores when we—as novices—are rationally justified in trusting experts. By definition, as novices we are not in a position to judge the truth of what our purported expert tells us²². Sources on both sides of the moon-landing debate (if we can call it such) will claim expertise, and there may even be debates internal to each side about the details and evidence that best supports their side's position. Insofar as there are debates between two experts (regardless of their stance on the moon-landing), we as novices can be faced not only with novice/expert situations but novice/2-expert situations²³. In the former, we find ourselves trying to evaluate the credibility of one purported expert. In the latter, we are faced with trying to adjudicate between conflicting putative experts on the topic at hand.

The sheer availability of reports poses a problem for something that has been suggested as a useful tool for evaluating testimony: corroboration. Ideally,

²¹ Frankfurt (2005).

²² See Plato's *Charmides* for an (particularly) early discussion of this problem.

²³ We use here Goldman's (2001) descriptions of these expert/novice problems.

“corroboration” means that we find the same information from a number of experts and this affords credibility to the claim we are investigating²⁴. However, as Fallis notes, information on the Internet can be easily copied and duplicated in a number of places, giving a veneer of corroboration, but ultimately it is all of one original source. It is not the mere fact of duplication that makes the information dubious, to draw from Fallis again, “over thirty different web sites (including the National Library of Medicine 2002) provide exactly the same information on how to treat children with fever. In this case, the information seems to be accurate. However, the fact that all of these sites corroborate each other still does nothing to help us verify that the information is accurate” (2004, 10*). What we really need to know is where the source we are using has got its information from.

One solution to this problem points to institutional markers of expertise. These are such things as earned degrees, support from other experts on the same side of the debate in question (best when this support is from the majority), performance in presenting their testimony, evidence pertaining to the purported expert’s past credibility, or evidence of questionable biasing factors (for example, funding from groups or individuals with a vested interest in the results).

Let us limit ourselves to cases where we feel confident saying that the google results are presented as true and where we are not experts on the issue at hand²⁵. With the cases leftover, we will still be able to find ourselves in many novice/expert situations and novice/2-expert (or more) situations when we google. The question now is whether we can appeal to any institutional markers of expertise that might assist us in such a situation, or whether there are non-institutional but reliable markers of expertise available to us when we google. We take it that a key consideration here, once again, is the nature of the source one has found on the Internet. The trustworthiness of someone who writes their own blog will need to be assessed differently from that of someone who has answered a question on a forum post.

Facts about the website itself can often aid us in assessing credibility, but this requires that we have sufficient knowledge about different types of websites, the role of advertising, and the nature of the testifiers to assess credibility²⁶. We discuss three examples here to demonstrate what types of credibility markers we can have when we google: publicly accessible feedback for individuals, advertising, and community

²⁴ See Goldman (1987) for a discussion.

²⁵ We put aside cases where people are using non-assertions to base their beliefs on, and cases where the searcher has enough knowledge to adjudicate between competing purported experts.

²⁶ Research looking at what users self-report as indicators of credibility when they arrive at particular websites seems to show that the design look (46.1% of participants) and organisation (28.5% of participants) are the most frequently mentioned features (Fogg, et al., 2003). If a website looks professionally made and pleasing, and it also well organised then individuals take this as a sign of credibility.

enforced standards.

One source of questionable help in these cases is biographical information provided by the authors themselves, which in some cases may be able to be verified (e.g., by checking their alleged institutional affiliations or other affiliations). On many sites, though, there are internal systems that report on users in ways that can indicate their credibility. For many forums, participants are ranked with titles that indicate their expertise (i.e., ‘moderator’, ‘administrator’), and receive publicly accessible feedback either for themselves as posters of information (e.g., systems that award visual tokens for the number of contributions made), or for the particular piece of information (e.g., “upvoting”²⁷). Examples of these forums include StackExchange (user titles, visual tokens (bronze/silver/gold medals), and upvoting), Reddit (visual tokens (karma), upvoting), and the Ubuntu Forums (titles, visual tokens (coffee beans)). In order to make use of these credibility markers effectively, though, requires that the user who has been directed to the website through googling knows both what credibility markers are in place, and what markers like visual tokens or upvotes indicate on that website. It is in some cases it is also useful to have an understanding of the risk of trolling behaviours on particular forums (e.g., Reddit), and how these can influence the upvoting of results.

A second potential credibility marker is the role of advertising on the websites that we find when googling. A healthy scepticism towards advertising is generally warranted in our epistemic practices as companies typically have monetary gain in mind rather than our epistemic interests. On the Internet such general scepticism is similarly warranted and, in particular, when the articles we are directed toward are “sponsored content”. Sponsored content, in contrast to traditional editorial content, are paid for articles that are used as advertisements or promotions, but are presented as though they are actually editorial content (usually with “sponsored content” written somewhere on the page). This is a way for advertisers to take advantage of the credibility we typically afford to editorial articles and instead direct it towards their product. We take it that the general rules we apply in being sceptical of advertised claims should apply when looking at sponsored content, the epistemic problems arise from our ignorance of what sponsored content is and a matter of being sensitive to seeing that a source is actually sponsored content²⁸.

²⁷ Upvoting is a way for users to indicate support for the information in a post or contributed to an Internet forum. We use “support” here because upvoting is used across a wide variety of websites and may be an indication of truth, reliability, aesthetic value, etc. Typically, upvoting systems will rank the contributions to the forum so that those with the highest votes are at the top.

²⁸ There is a growing body of literature looking at our ability to detect sponsored content and our general to assess credibility of internet sources. For a comprehensive bibliography of research on Internet credibility see Abdulla et. al. (2007). In one study looking at how users evaluate the credibility of websites, the presence of advertising was taken as an indicator of credibility by some and not by others (Fogg et al., 2003). The authors note that, “In 13.8% of

The third credibility marker is less a set of markers or signals but the commitment of some Internet sources to a set of epistemic norms upheld by the source itself. One major source of information that googling directs us toward is Wikipedia, and understanding how Wikipedia functions is an important part of understanding the credibility of it as a source²⁹. Entries on this Internet encyclopedia are created, edited, and managed by the community of users, importantly the majority of users are novices and not experts on the topics they contribute to (or at least, are not required to have qualifications to contribute or be a community moderator). As P. D. Magnus describes the situation in an article on trusting Wikipedia, “When we read *Wikipedia* entries, we read the uncredited, collective work of individuals whose only qualifications for contributing were an internet connection and an interest in doing so” (2009, 74)³⁰. The community does have rules for contributing to entries (the “five pillars”) that are enforced by the community. Consequently, Wikipedia is “self-consciously” not a normative epistemic source (i.e., one that tells us what we ought to believe about a topic) but rather a descriptive one (i.e., one that describes what the present consensus or majority view is on a topic)³¹. All entries are required to meet the Wikipedia standard of verifiability³², and a main task of community users who edit and maintain entries is to check-up on edits made by others (in part facilitated with discussion boards for each entry).

Given that there isn’t a structure in place to ensure that those who edit Wikipedia are in fact credible with regard to the contributions that they make this seems the rational move – publish what is already accepted in the literature³³. Entries are required to

the comments, users referred to advertising, usually negatively. But at times, study participants mentioned judicious use of advertising in a positive way” (2003, 7). While this indicates that some googlers are also thinking about what advertising might indicate on a website, the other 86.2% of the participants (a sample of 2,684) did not mention advertising as something that they take as an indicator of credibility in their responses.

²⁹ See Fallis (2008), “Toward an Epistemology of Wikipedia”, for discussions of the epistemic evaluation of Wikipedia in particular. Fallis (2008) also draws the connection between social epistemology and testimony with the nature of Wikipedia.

³⁰ In the paper, Magnus presents a detailed discussion of empirical evidence supporting Wikipedia’s reliability, corrective mechanisms, its similarities and differences with traditional encyclopedias, how we evaluate claims online, and makes some normative recommendations concerning Wikipedia usage.

³¹ As the website itself states, “We strive for articles that document and explain major points of view, giving due weight with respect to their prominence in an impartial tone. We avoid advocacy and we characterize information and issues rather than debate them. In some areas there may be just one well-recognized point of view; in others, we describe multiple points of view, presenting each accurately and in context rather than as “the truth” or “the best view”. All articles must strive for verifiable accuracy, citing reliable, authoritative sources, especially when the topic is controversial or is on living persons. Editors’ personal experiences, interpretations, or opinions do not belong.” (Wikipedia:Five Pillars (2017)).

³² See the community guidelines for more details (Wikipedia:Verifiability (2017)).

³³ It is possible that all contributors to Wikipedia are novices, in which case we might think it is also rational for them to accept only material accepted by the relevant group of experts. Even if it is merely a large proportion of contributors who are novices, we might take this to be epistemically advisable. For a defense of the position that novices ought believe with the larger group of experts, see Coady (Chapter 2, 2012).

come supplied with citations and other credibility markers (e.g., “citation needed” or comments at the top of the entry indicating the page needs work to meet the community standards). What this does mean, though, is that experts may not be able to update entries with our best information until sufficient time has passed for new information to be widely accepted by the relevant community of experts. This can mean that when our googling directs us to Wikipedia, our information can be out of date or false (until such a time as accepted by the majority of the relevant expert community)³⁴.

All of these measures for assessing the credibility of sources are taxing, they require having a reasonably comprehensive knowledge of different platforms one comes across when one googles and of the particular sources that one finds (i.e., to know which credibility markers to look out for). While googling provides us with easy and immediate access to answers for any question we could think to ask, our ability to rationally evaluate the quality of the answers we find is a hard ask—even for those who have grown up in the Internet age³⁵. What this suggests is a distinction between comfort and competence when googling. While many of us are perfectly comfortable asking our preferred search engine for the answers to our problems, it is apparent that many of us are far less competent at evaluating the credibility of what we find. The ease with which we google and readily accept answers despite this distinction raises interesting questions about our relationship with the information we come across when we google.

Googling as Extended Knowledge

Imagine “neuromedia” – technology that allows the capabilities of your smartphone to be encoded on technology directly linked to your neural network: a brain-chip, if you will³⁶. Suppose everyone in a particular community had access to this technology. They can query Google and its riches “internally”; they can comment on one another’s blog posts using “internal” commands. In short, they can share knowledge—they can review one another’s testimony—in a purely internal fashion. This would have, to put it lightly, an explosive effect on each individual’s “body of knowledge.” That’s because whatever I “post” mentally would then be mentally and almost instantly accessible by you (in a way that would be, we might imagine, similar to accessing memory). We’d share a body of knowledge by virtue of being part of a

³⁴ For an interesting case of an expert finding themselves unable to edit a Wikipedia page on their own area of expertise, see (Messer-Kruse, 2012).

³⁵ See Hargitti et. al (2010) and Wineburg, S. & McGrew, S., (2016) for two examples of studies looking specifically at young people’s ability to conduct reliable credibility judgements on the Internet (and what inhibits them). Studies of this sort also raise important questions about epistemic responsibility on the Internet and whose job it is to make us safe when we google.

³⁶ Lynch (2014).

network. But neuromedia also raises a further, more radical possibility: it is possible that neuromedians' knowledge would be "extended", in that individual neuromedians would be sharing not just the content of what is known, but the actual process or act of *knowing* itself.

While neuromedia makes this point vivid, current technology arguably already extends our knowledge. In this last section, we'll briefly examine this possibility.

Traditionally, humans have known about the world via processes such as vision, hearing, memory and so on. These modes of inquiry are internal; they are in the head, so to speak. But if you had neuromedia, the division between ways of forming beliefs that are internal and ways that are not would no longer be clear. The process by which you access posts on a webpage would be as internal as access to your own memory. So, plausibly, if you come to know, or even justifiably believe, something based on information you've downloaded via neuromedia, that's not just a matter of what is happening in your own head. It will depend on whether the source you are downloading from is reliable—and that source will include the neural networks and cognitive processes of other people. In short, were we to have neuromedia, the difference between relying on yourself for knowledge and relying on others for knowledge would be a difference that would make less of a difference.

David Chalmers and Andy Clark's "extended mind" hypothesis suggests that, in fact, our many of our intentional mental states are *already* extended past the boundaries of our skin³⁷. When we remember what we are looking for in a store by consulting a shopping list on our phone, they argue, our mental state of remembering to buy bread is spread out; part of that state is neural, and part of it is digital. The phone's notes app is part of my remembering. If Chalmers and Clark are right, then neuromedia doesn't extend the mind any more than it already is extended. We already share minds when I consult your memory and you consult mine.

The extended mind hypothesis is undoubtedly interesting, and it may just be true. But we don't actually have to go so far to think that when employing googling as a mode of inquiry, our knowledge is extended. That is because even if we don't literally share mental states, according to other philosophers, it is possible that we share the processes that ground or justify what our individual minds believe and think.

Sandy Goldberg has argued, for example, that when I come to believe something based on information you've given me, whether or not I'm justified in that belief doesn't depend just on what is going on in *my* brain. Part of what justifies my belief is whether *you*, the teacher, are a reliable source. What justifies my receptive beliefs on the relevant topic—what grounds them—is the reliability of a process that includes

³⁷ Andy Clark and David Chalmers (1998).

the teacher's expertise. So whether *I* know something in the sense already can depend as much on what is going on with the teacher as it does the student³⁸.

Goldberg's hypothesis seems plausible when applied to googling if you accept two conditions. The first condition is that when we form beliefs via googling—by relying on TripAdvisor, or Google Maps for example—we form beliefs by a process that is essentially socially embedded—a process the elements of which include not just chips and bits but aspects of other people's minds, social norms and my own cognition and visual cortex. The second condition is that a broadly reliabilist view of knowledge is correct. According to this position, you know that, e.g. Denver is in Colorado just when that belief is the result of a reliable belief forming process. Importantly, you don't have to know that the process in question is reliable—it just has to *be* reliable. If that is how knowledge, or at least knowledge via googling works, then if these processes are themselves extended, then the grounds for belief are extended, and arguably our Google knowledge is as well.

Conclusion

Googling is a mode of inquiry that is not going away; it is, for many people, the most common way to form beliefs about the social world. The epistemological questions it raises are old questions—but questions given a new, and pressing form. In a similar fashion, applied epistemologists (and ethicists) are beginning to look in more detail at issues of privacy and surveillance on the Internet, topics which dovetail with the phenomenon of preference-dependence discussed here. As googling continues to occupy such a central role in our lives, new concerns over epistemic responsibility and epistemic agency on the Internet are also rising topics of concern. We expect that applied epistemologists (including ourselves) will continue to explore these topics as our lives become ever more integrated with the digital world.

Bibliography

Abdulla, R.A., Aiken, K.D., Anderson, C.C., Boush, Casey, D., Eysenbach,
Flanagin, A., Metzger, M., & Rainey, J.G. (2007). Bibliography on Web /
Internet Credibility. Retrieved February 09, 2017, from

³⁸ Goldberg (2010).

<https://www.semanticscholar.org/paper/Bibliography-on-Web-Internet-Credibility-Flanagin-Metzger/042d8e5813f3d70ffbc7b0d9e99e4b4b99842372>

- Beitzel, S. M., Jensen, E. C., Chowdhury, A., Frierder, O., & Grossman, D. (2007). Temporal Analysis of a Very Large Topically Categorized Web Query Log. *Journal of the American Society for Information Science and Technology*. 58(2):166-178.
- Bok, S. (1978). *Lying: Moral choice in public and private life*. New York: Pantheon Books.
- Burge, T. (1995). Content preservation. *Philosophical Issues*, 6, 271.
doi:10.2307/1523046
- Carson, T. L. (2010). *Lying and deception: Theory and practice*. Oxford: Oxford University Press.
- Clark, A., & Chalmers, D. (1998, 01). The extended mind. *Analysis*, 58(1), 7-19. doi:10.1093/analys/58.1.7
- Cleve, J. V. (2006, 06). Reid on the Credit of Human Testimony. *The Epistemology of Testimony*, 50-70.
doi:10.1093/acprof:oso/9780199276011.003.0003
- Coady, D. (2012). *What to believe now: Applying epistemology to contemporary issues*. Chichester, West Sussex: Wiley-Blackwell.
- Fallis, Don. 2004. "On Verifying the Accuracy of Information: Philosophical Perspectives", *Library Trends*. 52:463-87.
- Fallis, D. (2008) Toward an epistemology of Wikipedia. *Journal of the American Society for Information Science and Technology*, 59(10), 1662-1674.

- Fogg, B., Soohoo, C., Danielson, D. R., Marable, L., Stanford, J., & Tauber, E. R. (2003) How do users evaluate the credibility of Web sites? A study with over 2,500 participants. Conference on Designing for User Experiences, San Francisco, California, USA.
- Foley, R. (2001). Intellectual trust in oneself and others. Cambridge, UK: Cambridge University Press.
- Frankfurt, H. G. (2005). *On bullshit*. Princeton, NJ: Princeton University Press.
- Fricker, E. (1994). Against Gullibility. *Knowing from Words*, 125-161.
doi:10.1007/978-94-017-2018-2_8
- Fricker, M. (2007). *Epistemic injustice: Power and the ethics of knowing*. Oxford: Oxford University Press.
- Goldberg, S. (2010). *Relying on others: An essay in epistemology*. Oxford: Oxford University Press.
- Goldman, A. I. (1987). Foundations of social epistemics. *Synthese*, 73(1), 109-144.
- Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford University Press.
- Goldman, A. I. (2001, 07). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 63(1), 85. doi:10.2307/3071090
- Grice, H. P. (1989). *Studies in the way of words*. Cambridge, MA: Harvard University Press.
- Hargitti, E., Fullerton, L., Menchen-Trevino, E., & Yates Thomas, K. (2010). Trust Online: Young adults' evaluation of web content. *International Journal of Communication*. 4: 468-494.
- Hume, D., Norton, D. F., & Norton, M. J. (2010). *A treatise of human nature*. Oxford: Oxford University Press.

- Joachims, T., Granka, L., Pan, B., Hembrooke, H., & Gay, G. (2005).
Accurately interpreting clickthrough data as implicit feedback.
*Proceedings of the 28th Annual International ACM SIGIR Conference
on Research and Development in Information Retrieval - SIGIR '05.*
doi:10.1145/1076034.1076063
- King, A. (2008). *Website Optimization: Speed, search engine and conversion
rate secrets*. Sebastopol, California: O'Reilly Media.
- Lynch, M. P. (2014), Neuromedia, extended knowledge and understanding.
Philosophical Issues, 24: 299–313. doi:10.1111/phis.12035
- Lynch, M. P. (2016). *The Internet of us: Knowing more and understanding
less in the age of big data*. New York: Liveright Publishing Corporation,
a division of W.W. Norton & Company.
- Magnus, P. D. 2009. "On Trusting Wikipedia." *Episteme* 6:74-90.
- Medina, J. (2013). *The epistemology of resistance: Gender and racial
oppression, epistemic injustice, and resistant imaginations*. Oxford:
Oxford University Press.
- Messer-Kruse. (2012). The 'undue weight' of truth on Wikipedia.
Retrieved February 08, 2017, from
[http://www.chronicle.com/article/The-Undue-Weight-of-Truth-
on/130704/](http://www.chronicle.com/article/The-Undue-Weight-of-Truth-on/130704/)
- Mills, C. (2007). White ignorance. In Tuana, S. S. N., editor, *Race and
Epistemologies of Ignorance*, 11-38. State University of New York
Press.
- Pariser, E., (2011). *The Filter Bubble: How the New Personalized Web Is
Changing What We Read and How We Think*. Penguin Publishing Group.

- Quattrociocchi, W., Scala, A., & Sunstein, C. R., (2016). Echo Chambers on Facebook. *The Harvard John M. Olin Discussion Paper Series*. Available at SSRN: <https://ssrn.com/abstract=2795110>
- Reid, T., Beaublossom, R. E., & Lehrer, K. (1975). *Inquiry and essays*. Edited by Keith Lehrer and Ronald E. Beaublossom. Indianapolis: Bobbs-Merrill.
- Saul, J. M. (2012). *Lying, misleading, and what is said: An exploration in philosophy of language and in ethics*. Oxford: Oxford University Press.
- Simpson, T. W. (2012). Evaluating Google as an epistemic tool. *Metaphilosophy*, 43(4), 426-445.
- Wikipedia:Five pillars. (2017a). Retrieved February 08, 2017, from https://en.wikipedia.org/wiki/Wikipedia:Five_pillars
- Wikipedia:Verifiability. (2017b). Retrieved February 08, 2017, from <https://en.wikipedia.org/wiki/Wikipedia:Verifiability>
- Wineburg, S. & McGrew, S. (2016). Why students can't google their way to the truth. *Education Week*, 36(11), 22-28.

Websites Cited

- Wikipedia [Internet]. [cited 2017 Feb 9]. Available from: <http://www.wikipedia.org/>
- Google [Internet]. [cited 2017 Feb 9]. Available from: <http://google.com/>
- The Guardian [Internet]. News, sport and opinion from the Guardian's US edition. [cited 2017 Feb 9]. Available from: <http://theguardian.com/>
- Plait P. Bad Astronomy. Slate Magazine [Internet]. [cited 2017 Feb 9]. Available from: http://slate.com/blogs/bad_astronomy.html

Stack Exchange. Hot Questions - Stack Exchange [Internet]. [cited 2017 Feb 9]. Available from: <http://stackexchange.com/>

Ubuntu Forums RSS [Internet]. [cited 2017 Feb 9]. Available from: <http://ubuntuforum.com/>

Wolfram|Alpha: Computational Knowledge Engine [Internet]. [cited 2017 Feb 9]. Available from: <http://wolframalpha.com/>

reddit: the front page of the internet [Internet]. [cited 2017 Feb 9]. Available from: <http://www.reddit.com/>