Re-assessing Google as Epistemic Tool in the Age of Personalisation

Tanya de Villiers-Botha[1][0000-0001-8790-9062]

[1] Stellenbosch University, Stellenbosch, 7600, South Africa
tdev@sun.ac.za

Abstract. Google Search is arguably one of the primary epistemic tools in use today, with the lion's share of the search-engine market globally. Scholarship on countering the current scourge of misinformation often recommends "digital literacy" where internet users, especially those who get their information from social media, are encouraged to fact-check such information using reputable sources. Given our current internet-based epistemic landscape, and Google's dominance of the internet, it is very likely that such acts of epistemic hygiene will take place via Google Search. The question arises whether Google Search is fit for purpose, given the apparent misalignment the general epistemic goal of promoting true beliefs and the greater online commercial ecosystem in which it is embedded. I argue that Google Search is epistemically problematic as it stands, mainly due to the opacity related to the parameters it uses for personalising search results. I further argue that in as far as an ordinary internet user is legitimately ignorant of Google's workings, uses it in an "ordinary manner", and is generally unable to avoid using it in the current information environment, they are not epistemically blameworthy for any false beliefs that they acquire via it. I conclude that too much emphasis is currently placed on individual epistemic practices and not enough on our information environment and epistemic tools when it comes to countering misinformation.

Keywords: Google Search, Epistemic tool, Personalisation, False beliefs, Blameworthiness, Veritistic value

1 Introduction

Search engines and other internet platforms currently constitute vital epistemic tools on a global scale.[1] Taddeo and Floridi [1] rightly point out that the large internet platforms are information gatekeepers in that they control access to and the flow of information in our societies. The focus in this paper will be on those platforms that primarily serve to link up users to information online. The main players here are search engines (with

[1] "Epistemic tools" here refers to tools by means of which we obtain information in order to form beliefs about the world.

Google Search the most widely used by a massive margin) and various social media sites (with Facebook dominating).[2] Currently, there is much focus on the problem of misinformation and disinformation that stems from our information environment. The internet is a major part of this information environment, and much of the discussion is centred around on the political dangers that the commercial internet holds for democracy (e.g. [2] [3] [4] [5]). While this is a grave concern, we face a more general epistemic danger as a result of the current workings of the commercial internet, in that our primary contemporary epistemic tools are owned and operated by a handful of commercial players whose interests are often misaligned with a fundamental epistemic good, namely truth-promotion. Thus, we face the possibility that a good-faith enquirer who sets out to obtain reliable information on any given topic faces an onerous task where the primary tools at her disposal may hinder rather than help her.

In the current literature, a common view is that one way to remedy the negative effects of misinformation on the internet is to promote individual internet users' digital, information, and critical-thinking skills (e.g. [6]; [7]). Internet users are enjoined to avoid social media or to fact-check information that they do obtain from such sources. The implication is that these are basic epistemic obligations and that an epistemic agent who fails to fulfil these duties is blameworthy for any false beliefs they may end up holding.[3] In contrast, Millar [8] argues that internet users who hold false beliefs obtained via social media are *not* epistemically blameworthy in that, i) they cannot reasonably be expected to fulfil their epistemic obligations, or ii) they are legitimately ignorant of the need to fulfil such obligations. He argues that even if social media users are aware of the fact that these are problematic sources of information, they cannot reasonably be expected to avoid using social media given its ubiquity and utility. In addition, the general user may be excused for not even realising that social media is epistemically problematic, given that the extent of their filtering and biasing of content is not generally known. My aim is not to assess Millar's arguments relating to social media use but to shift the focus to the internet information environment more generally. Presumably, those who direct social media users to fact check their information are working from the assumption that it is relatively straightforward to do so. Let us concede that in as far as one knows that one has epistemic obligations, such as needing to fact check specific information, and one is able to fulfil such obligations, one is epistemically blameworthy for holding false beliefs that stem from not fulfilling these obligations. A question that arises is whether there is good reason to believe that internet users who attempt to fulfil their epistemic obligations will be less likely to hold false beliefs. I will argue that this is not the case. In addition, I will argue that the individual is not well-placed to compensate for the epistemic shortcomings of our online information environment. Arguably, the general internet user who wishes to obtain or verify information will do so by making use of a search engine to look for reliable information

[2] According to one estimate, as of December 2021, Google had a global market share of (desktop) search engines of 85.55% [34].

[3] I will not delve into the vexed questions of whether there truly are epistemic obligations, nor whether and when epistemic blame might be warranted. For a useful recent overview of some of these issues see [35].

on the open internet.[4] And given Google Search's dominance, it is likely the search engine that will be used. However, I will show that not only is it not clear that a user will access reliable information via Google Search, but the average internet user is often not well-placed to assess the reliability of the information they obtain or even to know that this may be necessary in the first place. Hence, I conclude that internet users are often not blameworthy for false beliefs obtained from the internet. I also hold that countering misinformation will require greater focus on our information environment and related epistemic tools.

2 Veritistic value, expertise, and evaluating search engines

My interest is in information-seeking as it pertains to truth, i.e., instances where we are engaged in truth-seeking enquiry. Not all information-seeking relates to truth. One may seek out information that confirms one's pre-existing, for example. But it seems fair to say that most of us seek out information in order to establish truth at least some of the time. As Goldman [2] points out, "information seeking" is pervasive in our lives, both in terms of practical concerns (Will it rain?) and for satisfying our curiosity (Who won the game last night?). Much of human life will simply not be possible if truth, or something near enough, weren't often the object of our enquiries. Hence, my focus will be on instances of online information seeking aimed at obtaining knowledge as described by Goldman [2]: establishing true beliefs in a "weak" sense while avoiding error (false beliefs) or ignorance (no beliefs).

The internet is currently one of our principal sources of information and search engines are our primary means navigating it, making both vital epistemic tools. Goldman [2] provides prescient early analyses of the internet and of search engines as epistemic tools. He holds that modern societies dramatise the social dimension of knowledge where much of our truth-seeking behaviour is aimed at specialised agencies, tasked with gathering and disseminating knowledge. Writing in 1999, Goldman mentions the World Wide Web as one such agency, alongside newspapers and libraries [2]. Fast forward 25-odd years, and the picture looks remarkably different. Our primary current knowledge (information?) gathering and disseminating agency is the World Wide Web, or, more accurately, the major online platforms that dictate the architecture and traffic on the commercial internet [4]. This is especially true of the search engine-social media nexus formed by Google (and parent company, Alphabet) and Facebook (and parent company, Meta). Clearly, Goldman's useful analysis of, what he calls, the *veritistic value* of internet communication technology needs to be updated to accommodate the central role that today's internet occupies. "Veritistic value" refers to a practice or agency's propensity to affect the beliefs of those who engage with it. These are accorded positive verististic value in accordance with the degree to which they lead to true beliefs rather than false beliefs in their users. Communication technologies can be

[4] See, for example, [36].

assessed along veritistic lines as well. Roughly, technology that allows for the transmission of information in a way that allows for an increase in true beliefs in its users, and/or a decrease in false beliefs and/or ignorance in its users is veritistically valuable. It seems uncontroversial to claim that the internet and search engines, in particular, have significant veritistic value. Yet, remarkably few assessments of their veritistic value exist. Goldman's early and brief analysis is now mostly outdated. In 2012, Simpson [3] updates Goldman's analysis in accordance with the subsequent developments in internet technology and the role that it had come to play in society by then. I will briefly discuss these two foundational analyses before moving to update them in light of developments since 2012. My focus will be on Google Search as globally dominant search engine.

Goldman's analysis of what he terms "Computer-Mediated Communication" is prescient in its identification of how the internet environment can hinder the veritistic "knowledge enterprise" [2]. His main concerns were with navigating the vast amount of information available and the trustworthiness and reliability of that information, since anyone can post information online. His focus fell on search engines and blogs, and news- or chat rooms—early precursors to social media. Veritistically, he accorded low value to blogs and news- and chat rooms. The former he saw as being parasitic on newspapers and other traditional forms of media [9]. With the latter, he raised concerns about the possibility of what has come to be termed filter bubbles [10], or what he called "narrowly selected listening" [9]. An additional concern was the proliferation of "infojunk" (ibid), where anonymity and the absence of oversight and accountability meant that there were very little constraints on truth-telling. He also raised the spectre of commercial advertising, which could further call the motivation behind and reliability of posted information into question. He uses the example of the advertising of medicines masquerading as medical information. Of course, these problems are now associated with social media.

In contrast, Goldman sees search engines as necessary online epistemic tools with great potential for positive veritistic value. Search engines are tasked with addressing two major impediments to finding information online: i) indexing all of the information posted and ii) presenting information relevant to a specific enquiry in a useable form (i.e. in a way that the user can make use of in a realistic timeframe).[5] Hence, he suggests that search engines need to be evaluated along the dimensions of *precision* and *recall*. Precision is the ratio of relevant documents returned to a user's query over the total number of documents (both relevant and irrelevant) returned. Recall is the ratio of total relevant documents returned to the total number of [relevant][6] documents on the web. While seminal, Goldman's work on the epistemic valuation of search engines are perfunctory, perhaps indicative the epistemic role occupied by the internet in 1999. More than a decade later, the epistemic picture had changed drastically, paving the way for Simpson's more detailed analysis of Google Search as the by then dominant search engine. Simpson's analysis would show how search engines could potentially be

[5] Note that relevant information need not be accurate or authoritative. Relevant information is simply information that seems to pertain to a given query, whether accurate or not.

[6] See Simpson [3].

plagued by problems similar to those Goldman identified with blogs and news- and chat rooms.

Simpson's starting point is also the necessity of search engines. Given the vast amounts of information on the internet, a way is needed to link up a query to relevant information. In addition, all the potentially relevant information identified needs to be ranked, since no one can hope to sift through even a fraction of the relevant information available, let alone in a manageable timeframe. Simpson makes a useful distinction between the ways in which we generally use search engines, namely, navigationally and informationally. Navigationally, one uses a search engine to retrieve a particular bit of information one already knows about.[7] Informationally, we use search engines to find information on a topic where we do not have a specific bit of information in mind (e.g. information on the efficacy of a vaccine). In such enquiries, there may be many relevant sources of information. He concludes that in such informational searches contemporary search engines fulfil the role of *surrogate experts*. This points to additional dimensions along which search engines can be epistemically evaluated.

To illustrate, let us say that we want to establish whether p or not-p. An expert is someone who already reliably knows whether p or not-p. An effective (i.e. reliable and quick) way of establishing whether p or not-p is to consult an expert. Note that expertise comes in degrees. Whereas a "shallow" expert may only be able to authoritatively answer whether p, a "deep" expert will be able to contextualise p in terms of a bigger domain of knowledge and point out other relevant and reliable sources. Hence, a deep expert can evaluate the *relevance* and *reliability* of sources of information on p. The navigational use of search engines corresponds with the functions of a shallow expert, while the informational use corresponds to that of a deep expert. In presenting search results, a search engine is in effect making a judgement on the relevance of the information in the linked pages. And placing the results in a particular order on the search results page (SERP) implies that these are ranked in accordance with relevance (if not necessarily reliability). Here, search engines fulfil part of the function of a deep expert—pointing an enquirer to relevant sources of information. We can value them veritistically in accordance with this function.

Simpson thus adds timeliness—how long it takes a user to find relevant links on the SERP—and, when there is more than one relevant result, distributed timeliness—how all the relevant sources are distributed over the SERPs. Crucially, he also adds *authority prioritisation* and *objectivity*. Needless to say, all webpages containing relevant information are not necessarily reliable or trustworthy. Given the growth of the web, Simpson argues that a search engine that is able to distinguish between (seemingly) relevant pages with truthful content and those peddling falsehoods would be extremely veritistically valuable. In fact, we cannot do without this function. Hence, he adds the dimension of authority prioritisation—the ability to rank *reliable*, rather than merely relevant

[7] Simpson uses the example of finding a particular quote by a politician to verify who made the remark.

sources of information higher on the SERP.[8] One difficulty here is to identify computable markers of epistemic authority.

Simpson also adds the criterion of *objectivity*. Even when successfully identifying and prioritising reliable information, search engines could potentially skew the results by ranking some reliable sources higher than others. A practical constraint on using search engines informationally is that users tend to only consult the first few listed results. Hence, it is theoretically possible to rank all available links to reliable information in such a way so as to provide accurate but biased information. Simpson uses the example of a query regarding important philosophers. Let us assume there are three sets of important philosophers: German, French, and neither German nor French. Even if a SERP successfully prioritises all the authoritative sources on important philosophers available, it is still possible for these sources to be grouped so that those pertaining to important German philosophers are clustered at the top of the list, while all those pertaining to neither German nor French are grouped in the middle, and all those pertaining to French philosophers are grouped towards the end. Such results may count in the thousands. Practically speaking, a good faith and (mostly?) conscientious enquirer will still come away with the impression that there are no important French philosophers. Hence, Simpson adds an objectivity criterion—where equally reliable results are randomly distributed over SERPs to counter potential bias. Thus, Goldman/Simpson give us five dimensions along which to judge search engines as epistemic tools: precision, recall, timeliness/distributed timeliness, authority prioritisation and objectivity. These need not be the only relevant assessment criteria, but they are certainly essential. Our discussion will be confined to the latter two dimensions.

On Simpson's own analysis, search engines, and specifically Google Search, fall short on objectivity. This was due to the then relatively recent practice of the personalisation of results. Simpson discusses personalisation in terms of using algorithms to rank results in accordance with a specific user's past browsing habits (individual personalisation) or with the browsing habits of other users deemed similar to that user (profile personalisation). This causes search results to be ordered in terms of an algorithmic "judgement" of what that user would likely find relevant, based on pages that that user (or similar users) has visited before. Simpson argues that personalisation thus falls foul of his objectivity criterion, since the ranking of information relevant to a query is not done on epistemically defensible grounds. Potentially, individual and more general biases could be reinforced, leading to the epistemically disvaluable result of a decrease in understanding, if not in true belief. Understanding the distinction here is important, as my claim is that current personalisation practices are in fact potentially epistemically worse than Simpson recognises. What he doesn't consider is that results personalisation could also fall foul of his authority prioritisation/relibility criterion.

Simply put, Simpson's main concern with search engines and personalisation is that, although two users may enter the exact same query, the search engine will rank relevant, *reliable* sources differently for those two users, depending on their own past browsing habits and those of users like them. Thus, even though both users will be presented with

[8] Reliable here refers to bearers of truthful testimony that answers an enquirer's informational need.

reliable information and can form *true beliefs* about the object of their enquiry, they will lack the *understanding* that arises from an objective overview of the available reliable information. Going back to our important philosophers, the results pages, while containing the same reliable sources relevant to the query, may be ranked such that user A comes away with true beliefs about only important German philosophers, while user B comes away with true beliefs about only important French philosophers. Both will have true beliefs but will lack understanding relating to the topic of important philosophers as a whole. The idea is that objectivity gets a user from true belief to understanding and knowledge and hence personalisation threatens knowledge. What Simpson fails to recognise is the impact that personalisation potentially has on authority prioritisation and hence on true belief simpliciter. We may contest Simpson's claim that true knowledge entails understanding or argue that his thought experiment is contrived and that such non-objective SERPs will be marginal cases. However, falling short on authority prioritisation and, by extension, reliability, is much more serious in an epistemic tool. It is also a failing that a conscientious user cannot easily recognise or compensate for.

3 Assessing Google Search

3.1 Commercial incentives

Key to understanding Simpson's assessment of Google Search is his assumption that along the four dimensions of epistemic assessment other than objectivity, the interests of search engine operators, users, and society are aligned. It is worth quoting him in this regard [3]:

> Search engines' core business models are structured around advertising; Google provides a free service to enquirers, making money by providing sponsored links. Each time an enquirer clicks on a sponsored link, a small amount of income is generated for Google. The higher the number of enquirers who click on sponsored links, the higher Google's revenue. So, it is in Google's interest to provide as excellent a service as possible to the enquirer, to maximise the number of enquirers who use the search engine. Sheer volume of traffic is the strategy. Given that precision, recall, timeliness, generalised timeliness, and authority promotion are all dimensions of search engine performance that enquirers desire, it is in Google's interest to perform well on these. There is no reason to suppose that these outcomes are anything but publicly desirable (p. 440).

While Simpson is right that Google Search' core business model is structured around advertising and that the aim is to maximise the number of users, he is wrong in supposing that this necessarily provides an incentive to always deliver *reliable* rsults. Instead, the complex internet advertising ecosystem that has taken shape on the commercial internet potentially skews Google's workings away from primarily delivering *reliable* content towards primarily delivering *ostensibly relevant* content, i.e. content that the user "wants". In short, with personalisation that draws on the troves of information that

Google has on users (and users like them) from across the internet, relevance and reliability may come apart when ranking results. Hence, authority prioritisation/reliability can also become a casualty of current personalisation practices.

To understand how the business model behind the commercial internet potentially impacts reliability in search result rankings, one needs a basic understanding of the online advertising market. Much of this market, as well as the digital infrastructure of the internet, is controlled by two companies, Google and Meta [5]. A massive amount of internet traffic goes through platforms, websites, and apps owned by or affected by these two companies, making them among the most influential players in shaping our current online information environment. One does not need to be a user of Google's products, services, and apps[9] (other than Google Search), or of any of Meta's suite of products and services[10] to be affected by their dominance. The main source of revenue for both Google and Meta is advertising [11] [12]. More accurately, Google and Meta make the bulk of their revenue from collecting enormous amounts of data on internet users which they use to sell advertising opportunities to other companies and entities, both on their own platforms and on real estate that they own across the internet [13] [4] [5] [14] [15]. The kinds of data collected can include anything from IP addresses, time spent on page content, interaction with content (clicks, likes, retweets, watches, etc.), time of day, device type used, browser used, internet connection, etc., to highly personal information, such as location, name, telephone number, social connections, contact lists, transaction data, relationship status, interests, and browsing habits [4] [16]. The key to these companies' dominance in digital advertising is their ability to use the data they collect to target ads at those users who are thought to be most susceptible to what is being peddled, based on the analysis of this data. The colossal amounts of data collected is used to develop highly-granular profiles if internet users, which allow for highly specific targeted advertising.

This is where the front-end of these platforms' operations come in. Firstly, the frontend of platforms such as Facebook and Google Search offers advertising space where users can be targeted. They also serve as vital sources of internet-user data. Just about any interaction with an internet platform is a useful bit of data that can be transformed into information on that user and others like them. Hence, these platforms have an incentive to draw users to and keep them engaged on their platforms as much as possible. Müller [15] puts it succinctly when he states that "[t]he primary focus of social media, gaming, and most of the Internet in this "surveillance economy" is to gain, maintain, and direct attention—and thus data supply". User data is thus also used to tailor platforms and deliver individualised content to keep the user engaged on the platform for as long as possible. To do this, artificial intelligence is used to classify any given user in terms of a given machine-learning model to recommend or deliver content that they are most likely to engage with.

[9] These include Google Search, Gmail, Google Maps, YouTube, Google Scholar, Google Calendar, and Google Docs, among many others.

[10] These include Facebook, Instagram, and WhatsApp, Messenger, and Oculus, among others.

Although the proprietary nature of various online platforms' recommender algorithms makes it difficult to determine on what basis, exactly, specific content is recommended, it is clear that the system often favours the proliferation of content that Meta CEO, Mark Zuckerberg, describes as "sensationalist and provocative" [17]. "Sensationalist and provocative" content tends to elicit a lot of engagement on social media and keeps users interacting with a given site, app, or service [16]. As it turns out, controversial, highly emotive, and outlandish content—including misinformation and disinformation—tends to lead to greater engagement. Hence, AI recommender systems tend towards recommending such content. It should be emphasised that this is not a bug of the current system but a feature. There is very little incentive to reduce the amount of "engaging" content recommended and much incentive to keep recommending it. Arguably, contra Miller's claim above, it is now generally well-known that (overtly) social media feeds are epistemically suspect sources of information due to the dynamic described here. Facebook is an especially egregious example.[11] In terms of our criteria, it should be clear that social media feeds fare badly as veritistic tools, especially on the dimensions of objectivity and authoritativeness. What is less appreciated is the extent to which Google Search results are also personalised, on the one hand, and affected by the above commercial dynamic, on the other. Hence, even though Google Search may sometimes have an incentive to deliver reliable results high on SERPs, this is not necessarily always the case. Personalisation and the architecture of the commercial internet can undermine this incentive.

3.2 Personalisation and authoritativeness

To understand how the reliability of search results may be affected, we need a basic understanding of how Google Search works. As mentioned, vast amounts of results need to be ranked in accordance with their estimated relevance on the SERP. One of the main strengths of Google Search is the pioneering way in which it initially accomplished this ranking. As its creators point out in their seminal 1998 paper, Google was designed to deliver "high precision" results by posting documents deemed highly relevant to the query "in the top tens of results" [18].[12] This meant that Google not only had to identify documents containing information linked to the search query but had to filter those documents for quality and/or other indicators of relevance. For this, Brin and Page utilised the linked structure of the web, i.e. the fact that web documents can link to one another via hyperlinks. To assess the quality of web documents, these hyperlinks were treated analogously to academic citations—the more links to a page there were, the higher its "citation importance" [18]. Moreover, not all pages' links to a particular page were weighted equally. A link from a page that was itself highly ranked, was given more weight than a link from a lower-ranked page. Pages that were linked

[11] See [29] for quantitative analyses relating to its role in the spread of dis- and misinformation relating to the 2016 US election.

[12] Users rarely consult results lower down the list [36]. As the amount of information on the web grows, the problem of precision becomes more acute.

to too profligately, however, were downgraded, to counter obvious attempts at gaming the system. Hyperlinking was taken to be an objective measure of quality, and it meant that their search engine generally did better than rivals in finding and making accessible information relevant to queries. Already in their 1998 paper the authors point out that results can potentially be made more relevant, or personalised, by taking a user's "proximity information" such as "location, home page and bookmarks" into account. Subsequently, personalisation expanded exponentially, thanks to the insight that the vast amounts of metadata and other data that users generate online (and offline) can be used to make inferences about them, leading to ever-more refined possibilities for personalisation. Numerous internet platforms now make use of such personalisation to target content and advertisements. Personalisation also serves to improve search engine results, since it helps to narrow down the unimaginably vast numbers of potentially relevant results, e.g. by taking a user's general location into account. Hence, personalisation of some form is indispensable for an effective search engine. Personalisation became the Google Search default in 2009 [10]. Currently, Google uses "over 200" parameters or "relevance signals" to rank search results [19]. What these are remain proprietary. Nevertheless, despite changes to its algorithms over the years, it seems safe to assume that content linked to most will generally appear higher up on the SERP [20]; [21]; [22]. Google also explicitly states that country, location, past search history, search settings, and "recent activity in your Google account" are some of the signals it uses [22].[13] Crucially, users do not know what parameters go into personalising any given result.

Note that Google states that it uses information from a user's Google account to personalise results. This includes information from its social media platforms, such as YouTube [23]. Over and above its own platforms, Google has extensive tracking abilities via third-party tracking, using its advertising/analytics network (e.g., DoubleClick and Google Analytics) [23], with which it can gather data on users from across the internet, meaning that it does not only have data on Google account holders. In addition, profile personalisation also occurs, where results are tailored to a user based on an analysis of the data of "similar" users. This means that the kind of content one (or those deemed similar to one) has accessed on the internet in the past, on social media and other sites, can also influence what content one encounters when executing a Google search. Hence, the SERP of two different users of Google Search to the same query can look very different. A study by Le Huyen et. al [23], for example, found that search results on Google News for various politically contentious issues in the US were significantly different (i.e. showed significant political partisanship) for fresh user profiles, distinctly trained on different browsing histories. These "users" received results that were slanted in accordance with their apparent political leanings as inferred only from their "browsing histories" and nothing else (as no other information about them existed). This means that filter bubbles and echo chambers (see [24]) are not necessarily confined to overtly social media or to Google account holders. Information gained from Google Search may have more in common with information gained via other social media sites than is often appreciated. To some extent, at least, Google Search tends to

[13] It should be noted that Google denies "personalising" search result rankings, despite making use of such signals [36].

give a user "what they want", as inferred from a legion of data points (ostensibly) about them. Problematically, these data points may not indicate that a user "wants" *reliable* search results.

The extent to which personalisation affects Google Search results is a contentious issue in the literature, partly due to the difficulty in designing a study that measures differences in search results and partly due to the difficulty in determining what, exactly, constitutes "differences" [23]. There is also disagreement on the relevance of such personalisation (e.g. whether it significantly skews the information environment of any one user). In essence, assessing Google Search as an epistemic tool takes place in an information vacuum. From the users' perspective, they may know that their search results have been personalised but not what has gone into determining particular results. This is already undesirable from an epistemic standpoint. Moreover, information that is publicly available paints a picture that gives us reason for concern.

Simpson was right in his argument that taking a user's past browsing behaviour into account when compiling a SERP potentially compromises objectivity, even if all of the results are from reliable sources. What Simpson failed to appreciate was that it may not always be in Google Search's interest (or power) to deliver *reliable* personalised results. Trivially, some users may want relevant but unreliable information, e.g. information that supports their favourite conspiracy theory. However, it turns out that determining whether or not a user "wants" reliable information may not be straightforward. As explained, what users "want" is inferred from their past online behaviour and that of users "like" them. Such behaviour may skew towards unreliable content, but this need not indicate that this is what the user is after. A user or type of user may spend time browsing outlandish conspiracy theories, for example, but this may be an artefact of the commercial internet business model rather than indicative of preference. We have seen that personalisation tends to skew towards unreliable content on more overtly social media platforms, such as Facebook and YouTube. If the content they (or those "like" them) access here is taken to account by Google Search, this will steer them towards more such content. In addition, users often browse the web via social media platforms, thus adding more data points on the content that they "want", and they may become trapped in a vicious feedback loop. So, the first point of concern is that we do now know what parameters go into a given result nor how they are weighted. There is also reason to think that past browsing behaviour may skew results away from reliability, even for good faith enquirers who may be after reliable information.

A further concern is, even if reliability/authoritativeness features strongly as a parameter irrespective of a user's profile, we do not know how reliability/authoritativeness is determined.[14] As Goldman [2] appreciated, finding a calculable marker of authoritativeness is difficult. It seems safe to assume that the incoming-link ranking system described above still features, but this has limitations. Arguably, a high volume of incoming links from influential sources is a measure of *popularity* or *notoriety* more than of authoritativeness or reliability. This has become more problematic than it might have been in 1998, given the proliferation of content on the internet, the influx of mis-

[14] Google states that it uses signals to identify expertise, authoritativeness, and trustworthiness, but does not specify what these are [22].

and disinformation, and the tendency to promote provocative content on social media. In a sense, Google themselves concede this problem in that they make use of human quality controllers to assess various proposed changes to the search engine to ensure "better quality" results [25]. Human quality controllers are tasked with assessing (seemingly randomly) the quality of websites delivered on SERPS in response to queries. According to the guidelines given to these raters, when it comes to pages with what Google calls "Your Money or Your Life" (YMYL) content, special attention needs to be paid to assessing the "expertise", "authority", and "transparency" of those pages [26]. YMYL content "could potentially impact a person's future happiness, health, financial stability, or safety" and includes content relating to news, health, finance, government, and "other". Inter alia, quality raters are encouraged to determine whether the content of the pages they encounter was created by authoritative or reputable sources and whether it is "factually accurate" across the range of YMYL topics just mentioned. Hence, this human quality control system is partly meant to assess the workings of the search engine along our dimension of authoritativeness ([26]). Nevertheless, the guidelines on how to determine "expertise", "authority", and "trustworthiness" rely very heavily on external, "offline" markers of authority—reputation, institutional recognition, and the like. Raters also need to assess YMYL pages in terms of "accuracy and well-established medical/scientific/historical consensus where such consensus exists", but how these are to be determined other than by referring to the reputation of page content creators is not made clear. The expertise of the quality controllers is also unknown.

In terms of assessing Google Search along the dimension of reliability/authoritativeness, the epistemic picture, so far, is at best opaque. We know that search results are personalised, but we do not know what exactly goes into such personalisation. We have to trust that Google Search gives a high ranking to reliable sources, but we do not know to what extent this is the case. We also do not know what its metrics for establishing reliability/authoritativeness are nor how accurate these are. These are major obstacles to developing a fair assessment of the veritistic value of Google Search. It also does not bode well in terms of individual users' epistemic obligations. When attempting to fact check a given piece of information, a simple informational Google Search may or may not result in a SERP where the most reliable sources of information are most highly ranked. Problematically, a user may be directed towards unreliable sources of information if their data points seem to suggest that these sources are most relevant to them, even if they are, in fact, after reliable information.

3.3 Manipulation

The problem is exacerbated when one considers that Google Search rankings can be gamed. Famously, one week after the 2016 US election, the top news listing in a Google search for "final election results" was a link to a blog called "70 News", which falsely reported that Donald Trump had won the popular vote [27]. Similarly, in 2017, search terms relating to a report on Russian interference in the 2016 US election and other politically sensitive issues in the US yielded top links to RT, a Russian, state-sponsored TV-network and online site said to feature state propaganda and found to have spread

misinformation relating to the 2016 US election [28]; [29]; [30]. While outside observers cannot definitively show that such incidents are due to successful manipulation of Google Search, the likelihood is strong. Moreover, if such incidents were the result of the organic working of Google's system instead of manipulation, it would be worse from an epistemic point of view, as it would mean that Google's ranking systems sometimes fails dismally in filtering for reliability, even while working as intended. It is likely that these are examples of successful manipulation. As Ghosh and Scott [30] explain, whereas "white hat" search engine optimisation entails trying to move up the SERP through website architecture, content formatting, and getting other sites to link to yours, "black hat" search engine optimisation attempts to trick Google's algorithms into putting certain content high on SERPs for a short period, such as a news cycle, to influence opinion. This can be done, inter alia, with rich, regularly updated content, coordinated backlinking by a set of domains, promotion through social media and advertising spends [30].

The potential interplay between Google Search results and the business model of the rest of the commercial internet became most apparent in 2016 when in Veles, Macedonia, a cottage industry in generating online advertising income via cobbled-together websites experienced a massive windfall due to the United States presidential elections. Velesian teenagers set up websites on which they posted mostly fabricated, outrageous content relating to the US election.[15] They then posted links to their websites to various fake accounts on Facebook, from where the post could go viral and drive traffic to their websites and the waiting Google advertisements [31]. Such deceptive, "junk news" sites were widely shared on social media and achieved, for a while at least, high rankings on Google Search [32] [33]. Whereas junk sites dropped down the rankings after August 2017, when Google announced changes to its recommender algorithms, Bradshaw has subsequently detected an upwards trend in their ranking again, indicating an adaptation in gaming strategies [33]. Depending on the sophistication of the manipulation campaign, it is potentially difficult for a user to determine whether they have been served such manipulated content.

From the above, it should be clear that we have reason to question to veritistic value of Google Search as an epistemic tool. This is largely due to the lack of information available to end-users on its functioning. End users cannot easily establish what factors went into determining any given SERP. A user's own browsing history, and/or that of others deemed "similar" to them may affect both the objectivity and the reliability of results. Moreover, even if reliability features strongly as a ranking criterion, users do not know how it is measured or whether the metrics used are effective. Finally, we know that there are attempts to manipulate Google Search and that patently false content has appeared high up on SERPs in the past. All of these considerations count against Google as epistemic tool. It potentially does badly along two of the most important dimensions of assessment for veritistic value, namely objectivity and authority promotion.

[15] This also highlights another epistemically perverse incentive of the dominant web platform business model—creating junk content for commercial gain.

Clearly, a search engine that makes it easier to determine how strongly reliability features as a parameter in a given result and that is forthcoming on how it determines reliability will have greater veritistic value than Google Search. Hence, I suggest that the dimension of *transparency*, as it relates to the reliability of search results as an additional dimension along which the veritistic value of search engines should be assessed. At a minimum, a user should have some indication of how reliability featured in the ranking of results to any given search. Ideally, a user should have some, easily implemented, control over what goes into personalising a search, especially along the dimension of reliability. In addition, independent researchers need access to more information on Google Search's functioning is as far as it is necessary to determine its veritistic value. This is where epistemic and commercial interests potentially converge. It may be epistemically desirable to know what signals go into determining search results, but it may not be in Google's commercial interests to divulge this.[16]

I further contend that Google Search can be deemed *ethically* blameworthy in as far as it is (epistemically) illegitimately opaque in terms of how it identifies and ranks search result and in as far as it allows commercial considerations to outweigh epistemic considerations. The reason for this is firstly its presentation of itself as a *trustworthy epistemic* tool. Of course, as a commercial entity, Google Search need not function as an effective epistemic tool as measured along our criteria, as other social media platforms clearly do not do. However, it should not present itself as such. In as far as it does, it can be accused of being deceptive. In addition, in as far as it needlessly obscures its workings, making it difficult to independently determine its trustworthiness, it can also be held ethically blameworthy. Other technologies are required to both warn against and mitigate possible harms that may arise from their use. This should be the case here too. Google Search's dominance of the search engine market is also reason for epistemic and ethical concern, but space constraints preclude us from exploring this complex issue here.

4 Blameworthiness?

A question that remains to be addressed is the extent to which end users may be epistemically blameworthy for any false beliefs resulting from their use of Google Search. The above analysis suggest that users may often not be blameworthy, both i) on the basis of being legitimately ignorant of the need to fact check information so obtained, and ii) in that they cannot reasonably be expected to fulfil their epistemic obligation of fact checking all information so obtained. Firstly, whereas a user may be more obviously blameworthy for false beliefs that they acquire via social media, Google Search is generally considered to be trustworthy. It is less widely known that Google Search results are personalised and that this may affect the objectivity and reliability of search results. A user can plausibly be legitimately ignorant of having an epistemic obligation

[16] In as far as secrecy around signals protects the system from being gamed, this *may* be epistemically justified, provided that the epistemic damage incurred by such gaming outweighs the reduced epistemic agency of the user in not knowing the specifics of given search results.

to fact check information gained via Google Search, especially when it is presented high up on a SERP. In addition, even if a user were to be aware that personalisation potentially leads to epistemically problematic results, it is not always obvious when this occurs. It is impossible to know what factors have gone into the personalisation of any given result, and it is practically impossible to fact check all results to all one's queries. The user akso needs to trust that Google Search's metrics for reliability are accurate and effective *and* that the system has not been gamed in any given instance. Moreover, even if it becomes clear that a given search result needs fact checking, it is not clear that further online informational searches will fare any better. A user will have to use other, independent markers of authoritativeness and reliability. Finally, there are very few viable alternatives to Google Search, and it is not at all clear that they fare any better in terms of our criteria. And the internet remains the largest and most easily accessible deposit of general information. It is difficult for a user to avoid the internet when looking for information. All of this leaves very little room for according blame to those who adhere to false beliefs obtained from using Google Search in good faith. It also suggests that the excessive focus on the individual and their epistemic duties in countering the spread of misinformation is misplaced. Attention needs to be paid to the design and functioning of the primary tools that allow users to access our main repository of information, namely search engines.

References

1. Taddeo, M., Floridi, L.: New Civic Responsibilities for Online Service Providers. In: Taddeo, M., Floridi, L. (eds) The Responsibilities of Online Service Providers. Law, Governance and Technology Series, vol 31. Springer, Cham. (2017).
2. Goldman, A.I.: Knowledge in a Social World. Oxford University Press, Oxford (1999).
3. Simpson, T.W.: Evaluating Google As An Epistemic Tool. Metaphilosophy 43(4), 426-445 (2012).
4. Ghosh, D.: Terms of Disservice. How Silicon Valley is Destructive by Design, Brookings Institution Press, Washington (2020).
5. Simons, J., Ghosh, D.: Utilities for democracy: Why and how the algorithmic infrastructure of Facebook and Google must be regulated. Brookings Institution, Washington (2020).
6. Heersmink, R.: A Virtue Epistemology of the Internet: Search Engines, Intellectual Virtues and Education. Social Epistemology 32(1), 1-12 (2018).
7. UNESCO: Recommendation on the Ethics of Artificial Intelligence. United Nations Educational, Scientific and Cultural Organization, Paris (2022).
8. Millar, B.: The Information Environment and Blameworthy Beliefs. Social Epistemology 33(6), 525-537 (2019).
9. Goldman, A.I.: The social epistemology of blogging. In Van den Hoven, J., Weckert, J. (eds.) Information Technology and Moral Philosophy, pp. 111-122, Cambridge University Press, Cambridge (2008).
10. Pariser, E.: The Filter Bubble: What The Internet Is Hiding from You. Penguin, New York (2011).

11. Alphabet: Alphabet Annual Report 2021. https://abc.xyz/investor/static/pdf/2021_alphabet_annual_report.pdf?cache=3a96f54, last accessed 2022/10/14.
12. Meta: SEC Filings 2022. https://d18rn0p25nwr6d.cloudfront.net/CIK-0001326801/14039b47-2e2f-4054-9dc5-71bcc7cf01ce.pdf., last accessed 27/05/2022.
13. Zuboff, S.: The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power. PublicAffairs, New York (2019).
14. Véliz, C.: Privacy Is Power. Penguin (Bantam Press), London (2020).
15. Müller, V.C.: Ethics of Artificial Intelligence and Robotics. 2021. <https://plato.stanford.edu/archives/sum2021/entries/ethics-ai/>, last accessed 12/05/2022.
16. Alfano, M., Fard, A.E., Carter, J.A. et al.: Technologically scaffolded atypical cognition: the case of YouTube's recommender system. Synthese 199, 835–858 (2021).
17. Zuckerburg, M.: A Blueprint for Content Governance and Enforcement.. https://www.facebook.com/notes/751449002072082/, last accessed 24/05/2022.
18. Brin, S., Page, L.: The Anatomy of a Large-Scale Hypertextual Web Search Engine. Computer Networks and ISDN Systems 30(1-7), 107- 117 (1998).
19. Google Search Central: What Crawl Budget Means for Googlebot. https://developers.google.com/search/blog/2017/01/what-crawl-budget-means-for-googlebot?hl=en, last accessed 03/06/2022.
20. Hoffmann, S., Taylor, E., Bradshaw, S.: The Market of Disinformation. Oxford Information Labs, University of Oxford, Oxford (2019).
21. Genot, E., Olsson, E.J.: The Dissemination of Scientific Fake News : On the Ranking of Retracted Articles in Google. In Bernecker, S., Flowerree A., Grundmann, T. (eds.) The Epistemology of Fake News, pp. 228-242, Oxford University Press, Oxford (2021).
22. Google: How results are automatically generated. https://www.google.com/search/howsearchworks/how-search-works/ranking-results/, last accessed 03/06/2022.
23. Le Huyen, L., Maragh, R., Ekdale, B., High, A., Havens, T., Shafiq, Z.: Measuring Political Personalization of Google News Search. In: The Web Conference 2019 - Proceedings of the World Wide Web Conference, WWW 2019. Association for Computing Machinery, Inc, pp. 2957-2963, New York (2019).
24. Nguyen, C.T.: Echo Chambers and Epistemic Bubbles. Episteme 17(2), 141-161 (2020).
25. Sullivan, D.: An overview of our rater guidelines for Search. https://blog.google/products/search/overview-our-rater-guidelines-search/, last accessed 08/07/2022.
26. Google: General Guidelines. https://static.googleusercontent.com/media/guidelines.raterhub.com/en//searchqualityevaluatorguidelines.pdf, last accessed 08/07/2022.
27. Bump, P.: Google's Top News Link for 'Final Election Results' Goes to a Fake News Site with False Numbers. Washington Post (14 November 2016).
28. Waddell, K.: Kremlin-Sponsored News Does Really Well on Google. The Atlantic (25 January 2017).
29. Benkler, Y., Faris, R., Roberts, H.: Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics. Oxford Scholarship Online, Oxford (2018).

30. Ghosh, D, Scott, B.: #Digitaldeceipt: The Technologies Behind Precision Propaganda on The Internet. New America Foundation and the Shorenstein Center at Harvard Kennedy School (2018).
31. Subramanian, S.: Inside the Macedonian Fake News Complex. Wired (15 Febrruary 2017).
32. Allcott , H., Gentzkow, M.: Social Media and Fake News in the 2016 Election. Journal of Economic Perspective 31(2), 211–236 (2017).
33. Bradshaw, S.: Disinformation optimised: gaming search engine algorithms to amplify junk news. https://policyreview.info/articles/analysis/disinformation-optimised-gaming-search-engine-algorithms-amplify-junk-news, last accessed 01/06/2022.
34. Johnson, J.: Statista. https://www.statista.com/statistics/216573/worldwide-market-share-of-search-engines/, last accessed 01/07/2022.
35. Boult, C.: Epistemic blame. Philosophy Compass 16(8) (2021).
36. Ross Arguedas, A., Badrinathan, S., Mont'Alverne, C., Toff, B., Fletcher, R., Nielsen, R.K.: Snap Judgements: How Audiences Who Lack Trust in News Navigate Information on Digital Platform. Reuters Institute for the Study of Journalism, Oxford (2022).