

JUSTIFIED BELIEF IN A DIGITAL AGE: ON THE EPISTEMIC IMPLICATIONS OF SECRET INTERNET TECHNOLOGIES

BOAZ MILLER AND ISAAC RECORD

boaz.miller@gmail.com and isaac.record@gmail.com

ABSTRACT

People increasingly form beliefs based on information gained from automatically filtered internet sources such as search engines. However, the workings of such sources are often opaque, preventing subjects from knowing whether the information provided is biased or incomplete. Users' reliance on internet technologies whose modes of operation are concealed from them raises serious concerns about the justificatory status of the beliefs they end up forming. Yet it is unclear how to address these concerns within standard theories of knowledge and justification. To shed light on the problem, we introduce a novel conceptual framework that clarifies the relations between justified belief, epistemic responsibility, action and the technological resources available to a subject. We argue that justified belief is subject to certain epistemic responsibilities that accompany the subject's particular decision-taking circumstances, and that one typical responsibility is to ascertain, so far as one can, whether the information upon which the judgment will rest is biased or incomplete. What this responsibility comprises is partly determined by the inquiry-enabling technologies available to the subject. We argue that a subject's beliefs that are formed based on internet-filtered information are less justified than they would be if she either knew how filtering worked or relied on additional sources, and that the subject may have the epistemic responsibility to take measures to enhance the justificatory status of such beliefs.

I. INTRODUCTION

When we enter search terms into search engines such as Google, Bing, and Yahoo!, we typically expect them to give us relevant results. For the most part, we do not know what goes on behind the scenes – for example, how the search engine retrieves the results, what databases are accessed and what assumptions are made. Even if we know the general logic that is being implemented, we know neither the fine details of its implementation, which are often classified trade secrets, nor the content of the databases being processed.

Nowadays, people, especially of younger generations, increasingly rely on search engines and other online technologies as the primary and sometimes exclusive source of information on various subjects, such as politics, current affairs, entertainment, medicine and the environment, as well as about their close and remote social circles (Colombo and Fortunati 2011). Subjects form beliefs using information obtained from technologies

whose workings they don't understand; what are the epistemic implications of that ignorance for the justificatory status of those beliefs?

We argue that, all else being equal, people's ignorance of the workings of the technologies they use to obtain information negatively affects the justificatory status of the beliefs they form based on this information, with the possible consequence that their beliefs may fail to be justified. While this claim may seem trivial, existing theoretical frameworks in epistemology face difficulties explaining why it is so, and more generally, how a subject's reliance on technology in the formation of beliefs affects the justificatory status of those beliefs.

We introduce a novel analytical framework that conceptually connects belief justification, epistemic responsibility, action and technological resources. We argue that our framework can successfully explain why beliefs that are formed mostly on the basis of internet-filtered information are less justified than they would be if the subject knew how filtering worked or relied on additional sources. We focus on a particular type of internet technology – personalized information filtering, which is widely used by all major online services – to illustrate the epistemic worries that arise from subjects' excessive epistemic reliance on information technology whose mode of operation is unknown to them.

Our claim in this paper illustrates a more general principle that we call the *Practicable Responsibility Principle*. While in this paper we neither argue for it nor rely on it, we hope that the analysis of the case study lends it some indirect support. This principle states that a belief is justified inasmuch as it is responsibly formed, where a responsible subject is one who does, to the extent that she can, what is required of her in a particular situation to bring about true and rational beliefs. Responsibility here is delimited in part by role-expectations within particular situations, while practicability is delimited by facts about the subject's competencies as well as her technological, ethical and economic circumstances. Thus, both what is practicable and what a subject is responsible to do will vary from person to person and from situation to situation.

The remainder of the paper is divided into five parts. In the first part, we introduce a specific and widespread internet technology, personalized filtering, and describe one consequence of its use, the creation of 'filter bubbles' that can prevent users from properly assessing the biases and completeness of search results. We next elaborate an account of justified belief as responsible belief, and go on to discuss practicability as a constraint on responsibility and distinguish three types of justification failure. Then we analyze technological possibility in the context of personalized filters and argue that secret internet technologies detract from the justificatory status of the beliefs formed based on them. Finally, we argue that subjects can and should take measures to make up for the detraction in justification that results from hidden filters if their beliefs are to be justified.

2. FILTERING ALGORITHMS AND THE FILTER BUBBLE

The way that internet technologies such as search engines work is not only a technical matter, but a political one as well. The selection of a particular subset of web pages and their presentation in one order rather than another embodies and reflects a particular set of value judgments. For example, it has been argued that search results in mainstream search engines are biased toward the positions of the rich and powerful (Introna and

Nissenbaum 2000; Rogers 2004: 3–9). Similarly, reliance on search engines whose modes of operation are concealed from us may have negative political implications, such as the covert reinforcement of uneven participation. These political concerns have an epistemic dimension: biased or incomplete information, especially if it is thought to be unbiased or complete, can result in bad judgments. Although such worries are not unique to web technologies, we think the combination of their novelty and opaqueness creates an epistemic situation worthy of careful study.

To complicate things, different users of major contemporary online services encounter different output even if they enter the same input. For example, different users encounter different results for the same search query in Google. This is because online services such as Google search and the Facebook news feed personalize the content they provide to an individual user based on personal data they collect about him, such as his geographical location, previous search histories, and pattern of activity and use of various web services.

Pariser (2011) dubs such personalized, technologically filtered information ‘The Filter Bubble’. He writes:

Most of us assume that when we google a term, we all see the same results . . . But since December 2009, this is no longer true. Now you get the results that Google’s algorithms suggest is best for you in particular – and someone else may see something entirely different . . . In the spring of 2010, while the remains of the Deepwater Horizon oil rig were spewing crude into the Gulf of Mexico, I asked two friends to search for the term ‘BP’. They’re pretty similar – educated, white left-leaning women who live in the Northeast. But the results they saw were quite different. One of my friends saw investment information about BP. The other saw news. For one, the first page of results contained links about the oil spill; for the other, there was nothing about it except for a promotional ad from BP. (2011: 2)

Sunstein (2007) describes a dystopic vision of excessive content personalization on the internet called ‘The Daily Me’. He identifies four dangers the Daily Me poses to deliberative democracy. First, it denies citizens knowledge of the full range of choices available to them. Second, it exposes users to results that reaffirm their prior beliefs, therefore encouraging political dogmatism. Third, it enhances social fragmentation and consequently extremism. Fourth, it militates against the existence of a core of common experience and knowledge shared by the public, which is necessary for successfully carrying out democratic processes. Such worries about deliberative democracy have clear epistemic corollaries.

One may argue that the filter bubble is not a new phenomenon, or that its epistemic characteristics are not unique. On such a view, the bubble is equivalent to hanging around only with like-minded friends or consuming information only from narrow interest- or agenda-oriented sources, such as the Food Network, Fox News or the *New York Times*. As Pariser argues, however, the filter bubble differs from these other practices. First, the filter bubble is truly personal, as opposed to traditional specialized information channels, which still broadcast to a population, however small. Second, the bubble is invisible. Many people do not know that they are in a bubble, and even if they do, they do not know how it works or how to escape. Last, the bubble is involuntary. As opposed to being in social groups or reading newspapers, people do not choose to be put into bubbles. Politically liberal Pariser notes, for example, that his conservative Facebook friends gradually disappeared from his news feed, although he did not

voluntarily choose to exclude them (2011: 5–10). As Pariser notes, a central feature of a filter bubble is its *secrecy*:

Google doesn't tell you who it thinks you are or why it's showing you the results you're seeing. You don't know if its assumptions about you are right or wrong – and you might not even know it's making assumptions about you in the first place ... From within the bubble, it's nearly impossible to see how biased it is. (2011: 10)

Indeed, there is no good reason to assume that the filtering criteria correlate with epistemic desiderata such as reliability, objectivity, credibility, scope and truth.¹ Filters are constructed by companies that compete for user attention and tend therefore to prefer content that is interesting and comforting rather than dull or challenging. If a person is interested in a particular topic, such as celebrity gossip, there is no good reason to suppose the filtering algorithms will choose only reliable stories. When it comes to politics, there is little reason to think that they will present troubling stories and views that are not aligned with the user's existing political orientation.

It is worth noting that Facebook has taken this criticism seriously, and responded with its own study purporting to show that

even though people are more likely to consume and share information that comes from close contacts that they interact with frequently ... the vast majority of information comes from contacts that they interact with infrequently. These distant contacts are also more likely to share novel information, demonstrating that social networks can act as a powerful medium for sharing new ideas, highlighting new products and discussing current events. (Bakshy 2012)

The claim is that inclusion of input from weak links counteracts bubble effects. The extent to which our use of Facebook or Google affects the quality of the beliefs we form is, in the end, partly an empirical question.

Our concern in this paper is that the algorithms of Facebook, Google and other content providers are opaque, which makes this determination of quality especially difficult for users. It may well be the case that web technologies have benefits, such as giving us easier access to more information than ever before, but the question is not whether we are better off epistemically with Facebook than without it; rather, it is whether and to what extent we would be better off without the blinders on the workings of filter bubbles than with them. In epistemic terms, the question on which we focus is: how does users' ignorance of the workings of their sources of information affect the justificatory status of their beliefs?

Mainstream theories of justification such as evidentialism and reliabilism seem to face difficulties dealing with this question. According to evidentialism, the justification of *S*'s belief that *p* at time *t* depends only on the evidence *S* possesses in *S*'s mind for *p* at *t* (Conee and Feldman 2004). If *S* has excellent evidence for *p*, then according to evidentialism, *S*'s belief that *p* is highly justified. But as we have seen, a possible reason that her evidence strongly supports *p* is that she has been exposed only to a biased subset of the available evidence because of the operation of filtering technology, about the workings

1 Simpson (2012) introduces an epistemic normative evaluative framework for search engines based on the objectivity of their results. He reaches a conclusion similar to ours regarding the epistemic hazards of personalization.

of which she has little or no evidence. Hence, it seems to us that *S*'s belief that *p* should be regarded as less justified than evidentialism would deem it. As we will argue in the next section, the difficulty with evidentialism partly lies in Conee and Feldman's too narrow notion of epistemic responsibility.

Reliabilism seems to face other challenges. According to process reliabilism (the leading version of reliabilism) justified belief is belief generated by a reliable cognitive process (Goldman 2011). Imagine two subjects with the same cognitive apparatus who perform the same Google search, and consequently form the belief that *p*. Suppose that they encounter different results because of personalized filtering, and that their respective beliefs therefore differ in their justificatory status. How should process reliabilists respond to this scenario? One option is saying nothing about it, because filtering is outside of the explanatory scope of their theory. We believe, however, that a robust theory of justified belief should aspire to explain such common scenarios. A second option is to extend 'cognitive process' outward beyond the human cognizer's bodily boundaries to include a discussion of filtering.² If reliabilists follow this route, however, they will need to address the host of problems it introduces, such as blurring the boundaries between the cognitive agent and her environment, and attributing epistemic agency to networks of humans and computers (Giere 2006, 2007; Preston 2010; Goldberg 2012). We propose a third option, which avoids such problems to begin with; and as we tentatively suggest in the next section, reliabilists need not necessarily find our framework objectionable.

Standard theories of testimony do not seem to offer much help in addressing the worries raised by secret filter bubbles either. We may think of online news reports, blogs, and perhaps search-engine results, as testimonies.³ Standard accounts in the epistemology of testimony analyze the recipient's justification to believe the testifier in terms of the testifier's sincerity and competence. Yet these factors are largely irrelevant to the question at hand. First, while sincerity and competence are relevant to the justificatory status of every individual testimonial belief obtained from the bubble, the question here concerns the accumulated influence over time of a set of testimonies on the justificatory status of the beliefs the recipient has formed based on them (cf. Goldman 2008: 117). Second, a major factor affecting such justification is the secret filtering algorithm that determines which testimonies the subject encounters; sincerity and competence are not necessarily the right terms, even metaphorically speaking, for discussing the properties of algorithms.

It seems that the difficulties of these theories with addressing secret filter bubbles lie in the tendency of standard epistemology to analyze knowledge in terms of human beings' properties. Despite our vast and deep dependence on technology for acquiring knowledge and justified belief, epistemology has not, for the most part, given serious thought to the role technology plays in the fabric of knowledge and justification.⁴ Our paper constitutes an initial step in remedying this situation. The solution we propose begins with the view

2 This mind-expanding move has been made by influential philosophers of mind (Clark and Chalmers 1998; Clark 2010) and STS scholars (Haraway 1991; Knorr-Cetina 1999; Latour 2005). However, they do not share reliabilists' commitments.

3 The question of whether and on what conditions information from computers and other instruments constitutes testimony has been largely overlooked. For a few exceptions see Sosa (2006), Humphreys (2009), and Tollefsen (2009).

4 For exceptions, see Lehrer (1995), Baird (2004), Humphreys (2004), Rothbart (2007), and Simon (2010).

that justified beliefs are formed by subjects who have taken appropriate actions to bring about true and rational beliefs. We argue that available technologies play an important role in determining which actions are practicable and appropriate. In the next section we introduce our account of justified belief as responsible belief.

3. JUSTIFIED BELIEF AS RESPONSIBLE BELIEF

Our framework endorses a responsibilist account of justified belief. We generally follow Kornblith's account of justified belief as responsible belief, according to which an epistemically responsible subject desires to have true beliefs, and her actions are guided by this desire. Kornblith conceptually relates *justification*, *responsibility* and *action*:

Sometimes when we ask whether an agent's belief is *justified* what we mean to ask is whether the belief is the product of *epistemically responsible action*, i.e. the product of action an epistemically responsible agent might have taken. . . . When we ask whether an agent's beliefs are justified we are asking whether he has done all he should to bring it about that he have true beliefs. The notion of justification is thus essentially tied to that of action, and equally to the notion of responsibility. (Kornblith 1983: 34; emphasis in the original)

Whether a subject's belief that *p* is justified at time *t* may depend on whether he performed certain relevant actions prior to *t*. That is, sometimes responsible subjects are not only required to form beliefs responsibly, but also to have conducted some inquiry before they come to believe that *p*. This isn't just to say that a subject's ignorance of the shortcomings of his evidence is not always a valid defence against the charge of being epistemically irresponsible. Rather, it is to say that some forms of ignorance are culpable, and the subject may be blamed on epistemic grounds.

Our view is that responsibilities – and the means to fulfill them – often accompany the specific roles we take on. Being a doctor or judge is to be entrusted to make certain kinds of determinations on the basis of evidence gathered in a prescribed way and interpreted in light of a standard body of case knowledge using sanctioned modes of reasoning. Yet roles need not be formal. A casual request like 'George, does the restaurant take reservations?' assigns George the responsibility of finding out, and common experience supplies him with reasonable methods for doing so: stopping by, phoning or searching online listings, if he does not already know the answer. It might be argued that some areas of human activity are exempt from any epistemic responsibilities, but we think epistemic responsibilities are quite widespread, and at any rate we limit our discussion to cases in which subjects have accepted at least the basic epistemic goal of forming true and rational beliefs, which we think is enough to confer certain basic responsibilities. For example, internet users are responsible to insure that source information is unbiased and complete enough to underwrite the judgment being made.

Before we proceed to our argument, we want to address some potential objections to a responsibilist account of justified belief. The responsibilist account of justification has a venerable history in philosophy. For a long time, the view that justified belief is analyzable in responsibilist terms was widely accepted in Western epistemology.⁵ It became disputed

5 See Alston (1988: 294 nn. 4 and 5).

with the rise of reliabilist theories of justification. Reliabilists deny that a subject's reasons or evidence for believing are what ultimately justify her beliefs. For them, what determine if a subject's belief is justified are those things in the world that are causally responsible for the fact that the belief-forming subject ends up with a true rather than false belief. They argue *inter alia* that responsibilist conceptions of justification are too intellectually demanding of epistemic subjects, who cannot be expected to know all the rules of good epistemic conduct (Goldman 1999). We believe, however, that a responsibilist account of justification need not necessarily be in conflict with a reliabilist one. For example, Williams (2008) proposes a responsibilist account of justification in which a responsible subject is not required to explicitly follow a set of known rules in forming her beliefs. Williams argues that his account is compatible with reliabilism because a responsible subject's beliefs would be reliably formed.

In any case, readers who are not inclined like us to think that responsibilism is reconcilable with reliabilism may still read our argument as concerning responsible belief, where responsibility is a *constraint* on justification, namely, one justifiably believes that *p* only if one responsibly believes that *p*. This constraint seems intuitively plausible, at least in the context of obtaining justified belief from internet sources, where subjects' discretion and critical judgment are called for. In such cases, the standards of epistemic justification to which we commonly allude when we refer to beliefs we deem justified include a dimension of epistemic responsibility. Put differently, it seems hard to imagine how, for example, a subject who performs a web search in order to seek certain information may end up with justified beliefs without adhering to responsible epistemic conduct, that is, without aiming at obtaining rational and true beliefs.

Let us now address a potentially more serious objection to our account from within the responsibilist camp. Certain responsibilists want to keep separated the question of whether a belief is justified and the question whether a subject has made appropriate investigations. Conee and Feldman support the view that justified belief is responsible belief, but deny that there is an *epistemic* duty to gather evidence, or that the justificatory status of a belief may depend on whether the subject fulfilled such a duty. In their view, the justificatory status of a belief that *p* depends only on the evidence the subject has for *p*. If there is ever any duty that relates to the conduct of inquiry, for example, a duty to gather more evidence, it is only *moral* or *prudential* (Conee and Feldman 2004: 189).

We reject this view. Contra Conee and Feldman, there are *clear cases* in which a subject who does not gather evidence before forming a belief is being epistemically irresponsible. For example, a scientist studying the reasons that boys generally outscore girls on math SATs would be epistemically irresponsible if he allowed a desire to find a biological basis for these differences, for which he has some evidence, to prevent him from seeking or seriously considering evidence that the basis might be cultural (Simson 1993: 374). If he did so, he would allow prior prejudice and goal-directed biases to influence his belief formation. Such biases would likely obstruct him from reaching true or rational beliefs. Because an epistemically responsible subject aims at true and rational beliefs, such behaviour is clearly epistemically irresponsible, rather than merely morally or prudentially irresponsible. It obstructs the subject from achieving epistemic aims. The beliefs formed this way are therefore unjustified. The subject may have additional moral or prudential reasons to seek evidence, but the duty to gather evidence in such cases is clearly epistemic.

There is another problem with Conee and Feldman's view. Suppose Jones is a headstrong self-admiring young physicist. Jones presents a novel theory, which is supported

by the evidence that he has, to his colleagues. The theory is harshly criticized by a senior colleague, but Jones is too preoccupied and lost in thought, privately indulging in admiring his own work, to take notice of the fact that his theory is being criticized, to say nothing of considering the criticism. The evidence he has for his theory is the same as it was before. Kornblith (1983: 36) argues that Jones's belief in his theory is not justified because he is being epistemically irresponsible. By contrast, Conee and Feldman argue that it is. They argue that if his evidence supporting his theory is the same as it was prior to the presentation, his belief in his theory is justified (if it was before). They write: 'It may be true that the young physicist ... lacks intellectual integrity ... But the physicist's character has nothing to do with the epistemic status of his belief in his theory' (2004: 90).

In our view, one problem with Conee and Feldman's analysis of this example is that it fails to acknowledge our epistemic dependency and reliance on others, which is twofold. First, other people, including experts, often possess the best available evidence for our beliefs, and we acquire justification for many beliefs by believing their testimony (Hardwig 1985). Second, the justificatory status of our beliefs typically improves after they undergo critical scrutiny and evaluation, as in peer review, in which errors and unwarranted background assumptions are exposed (Longino 2002). Jones fails to use his colleagues as an epistemic resource to justify his beliefs; thus he is being epistemically irresponsible. Conee and Feldman's analysis fails to note this dimension of justification.⁶

In this section, we have elaborated the view of responsibilist justification on which our argument rests. In the next section, we turn to the practical limitations on what can reasonably be expected from a responsible subject, with a focus on what we take to be a hard limit: technological possibility.

4. PRACTICABILITY, TECHNOLOGICAL POSSIBILITY AND JUSTIFICATION FAILURE

In the previous section, we presented a responsibilist account of justification. A responsible subject does what she can to bring about true and rational beliefs. The idea is to reflect a subject's particular epistemic position: not only what she already knows, but also what she can find out. When I call my office to ask a colleague if she knows whether a letter I am expecting has arrived, I expect her to check the incoming mail before she answers. The point is that epistemic responsibility reflects not only what knowledge a subject already has, but also what knowledge the subject should be expected find out, given her roles and abilities. But reasonable expectations have limits: I won't blame my colleague for not noticing that the letter has slipped behind a desk or was delivered to the wrong recipient. The task in this section is to begin to delimit which actions are the appropriate ones for an epistemically responsible subject to undertake in acquiring justified beliefs.

Responsibilities can require a subject to take action, but those demands are not without limits; to be carried out, they have to be *practicable*, that is, possible in practice. A given subject cannot undertake just any investigation. She will be competent to perform only some investigations, and her technological, economic and ethical circumstances will

6 This dimension is captured by contextualist theories of justification, according to which being justified is being able to adequately respond to relevant challenges by one's epistemic peers (Longino 2002; Williams 2001; Annis 1978).

allow still fewer. Here, we focus on technological possibilities as a hard constraint on responsibility. Technological possibility depends on both material and conceptual resources. For example, the possibility of spanning a river with an iron bridge turns on both what the world is like (i.e. that iron is available and has certain properties) and how our concepts fit together (i.e. that we think iron has certain properties that we can put to use in making trusses). Without the physical possibility, the bridge would fail. Without the conceptual possibility, it would never be attempted (Record forthcoming).

What is the relationship, then, between technological possibility and practicability? A subject's technological tools are a clear determinant of what is practicable. An action is only practicable if it is technologically possible. Not having access to a microscope (together with the competencies required to use it) means not being able to observe tiny things, not having X-ray equipment means not being able to detect some fractures, and not having a web-enabled device means not being able to use search engines.

How does technology help or hinder a subject's attempts to gain knowledge? Technological resources make certain activities technologically possible, where without the technology the activity would not be practicable. When those actions are knowledge-apt, they enable us to take on more epistemic responsibilities or to satisfy responsibilities we could not before. For example, it would be impossible to know that the surface of the moon is not smooth without observing it with a telescope or some other enabling technology.

By allowing or limiting subjects' attempts to gain knowledge, the availability or lack thereof of certain technologies effectively changes standards of epistemic responsibility and justified belief. Recall that the responsibility criterion provides grounds for saying when a belief is justified – namely, when it is responsibly formed. When a technology enables subjects to conduct certain inquiries or make determinations that would be impracticable otherwise, they may reasonably be expected to perform them in order to fulfill their epistemic responsibilities. For example, at the beginning of the modern Olympic Games, judges visually determined sprint-race winners, but photo-finish cameras became required for this determination in close races soon after the technology became available; in 1991, vertical line-scanning video replaced human judges altogether (McCrory 2005). Technology may also limit action, for example by blocking access to previously available information, thus lowering standards of epistemic responsibility.

To be clear, practicability provides both *upper* and *lower limits* for deciding when a subject has done enough to satisfy the responsibility criterion. With respect to the upper limit, a responsible subject should obviously not be expected to do more than what is practicable. With respect to a lower limit, as in the photo-finish example, when a certain activity becomes practicable, performing it might become a minimal requirement for forming justified beliefs on certain matters. As a rule of thumb, the more practicable a putative justification-increasing investigative activity is, the likelier that it should be included among the subject's responsibilities. The exact content of a subject's epistemic responsibilities will depend on her role and other relevant circumstances of the case. Our point is that practicability constitutes a rigid upper bound on justification standards for beliefs, and is also a prominent factor in setting minimal standards of justified belief.

It is important to stress that, having merely completed all relevant practicable investigations regarding *p*, a subject is *not* necessarily justified in believing that *p*. For example, in rare cases, even a photo-finish is not accurate enough to determine the winner of a race. Rather, once the plausibly fruitful avenues have been exhausted, the subject is free to stop

investigating, whether or not she has gathered enough evidence to reach a solid conclusion. That is, justified belief is responsible belief, but not all responsibly conducted putative belief-producing actions result in belief.

We are now in a position to distinguish three types of justification failure, i.e. three types of cases in which a putative belief that p will fail to acquire justification:

- (JF₁) The subject has not adequately fulfilled her epistemic responsibilities. For example, the subject ought to have conducted certain inquiries, but did not, and consequently has not found evidence that might have supported (or ruled out) p .
- (JF₂) The subject has fulfilled her epistemic responsibilities, but p falls short of justification. For example, the subject has adequately sought evidence for p , but the evidence does not support p . A common case will be when the subject can competently find out the truth about matters such as p , but p is false, thus she will fail in her attempts to justify it.
- (JF₃) An activity required to justify p is not technologically possible, that is, the required material and conceptual resources are not available; for example, the confirmation of a theory in particle physics may require a more powerful particle accelerator than the ones available.

A subject may not know which type of failure she is encountering, or be wrong about it. Suppose a scientist makes every reasonable effort to confirm a theory and fails. She believes that, if the theory is true, she can confirm it. She therefore forms the belief that the theory is false. Thus, she believes she is in JF₂. In fact, the theory is true, but the equipment and method required to confirm it have not yet been invented. The scientist is really in JF₃.

Before we continue, let us briefly address an objection to our view according to which tying responsibility to practicability would make justification depend on practicability, and therefore competence, implying a lower standard of justification for an incompetent subject than for a competent one with the same role-responsibilities. We acknowledge that this may sometimes be the case, but we do not think that such dependence of justification on subjects' competence is so problematic. Let us first stress that the point of the practicability criterion is not to say that, once a subject has done what is practical, any judgment he reaches is justified. Rather, it is to say that practicable actions are the limit of what can be expected of him before he attempts a judgment. The question for the responsible subject is not, 'what investigations might shed light on this situation', but 'given my competencies and resources, what investigations can I undertake to increase my likelihood of coming to a responsible conclusion?' It may well be that a particular subject has neither the competency nor resources to come to a satisfactory conclusion. In this case, her only responsible recourse (belief-wise) is to suspend judgment.

Our framework states that subjects may vary in the standards they need to meet to have a belief justified on a given matter because one subject has access to a particular technology that enables her to conduct certain inquiry, which then may be required of her to reach justified belief, while the other subject has no access to such technology, hence is not required to perform this inquiry to achieve justification. We may similarly imagine a case in which two subjects are both sufficiently competent to perform a certain role and have access to the same technological and other resources, yet one is more competent than the other, and this competence allows her to perform an inquiry that the other

subject cannot. In such a case, in some circumstances, the more competent subject may indeed face higher standards of justification than the less competent subject, because she might be expected to perform the relevant inquiry to acquire justification. Such a case is consistent with our framework. But as we have stressed, we do not argue that any subject is competent to perform any role, and we do not say that, once a subject has responsibly conducted the relevant inquiry with respect to p , she may *always* legitimately believe that p . The responsible attitude may still be disbelief or suspension of judgment. In other words, the individual subject's competence may in some cases play a role in determining the standards of justified belief she needs to meet, but we do not completely relativize these standards to the individual subject's competencies. Such standards are a function of other factors as well, particularly her role and the available technological means at her disposal.

In order to analyze the implications of the reliance on personalized technologically filtered information for the justification of beliefs formed based on it, we presented a conceptual framework that connects justified belief, epistemic responsibility, action and technological possibility. We argued that the availability of particular technologies enables or restricts particular inquiry-related actions, thereby effectively changing minimal and maximal requirements of epistemic responsibility in a given situation. Therefore, changing technology effectively changes standards of justified belief.

Our framework echoes other works that also stress the pragmatic dimensions of knowledge and justification.⁷ However, existing accounts all focus on subjects' *stakes* regarding a proposition, particularly the risk of wrongly accepting or believing a false proposition or wrongly rejecting or disbelieving a true one. By contrast, we identify another pragmatic dimension of knowledge – the role of available conceptual and material resources, particularly technology, as defining and constraining responsible epistemic action required for achieving justification. In the next section, we discuss the relevant features of filtering technology to epistemic justification.

5. EPISTEMIC FEATURES OF THE FILTER BUBBLE

Now that we have laid out our conceptual framework, let us return to the specifics of the case at hand. The hidden algorithms behind search engines and other filtering technologies have clear implications for what actions a subject can practicably undertake to vindicate a belief, and hence to its justificatory status. Let us focus on three characteristics of internet information sources: discernment in aggregation, transparency in generating results and representativeness of the database.

7 Regarding justified belief, our account resonates with Annis's (1978) and Foley's (2005) theories of justification, which both stress the role of contextual pragmatic elements, e.g. subjects' roles, in defining standards of justified belief. Regarding knowledge, according to a recent view dubbed 'pragmatic encroachment', whether a subject has knowledge partly depends on her interests; specifically, if she has high stakes with regard to a proposition, she is *ceteris paribus* in a worse position to know it than if she has low or no stakes regarding it (Fantl and McGrath 2009; Stanley 2005). Douglas (2009) similarly argues that evidential thresholds for accepting and rejecting scientific theories inevitably involve value judgments because social values determine the inductive risks we are willing to take in a given context.

First, discernment in aggregation: how are different contributions sorted and combined together into the list of results displayed by Google? Is it possible to tell the difference between contributions from experts, biased parties, paid-for inclusions or hooligans? The more control subjects have over the process of aggregation and the more knowledge they have about the likelihood of bias or capture by special interests, the better they can protect themselves from undiscerning websites, like ad-baiting networks.

Second, transparency of generating results: does the user know how the site works, particularly, how relevance is calculated, and what assumptions about his needs and interests are made? Most algorithms are opaque or even trade secrets. Most users are not aware of these facts, as we observed at the start. Users typically believe that they are getting the same results anyone else would get, and they typically overestimate the completeness of those results. On the other hand, search logic opacity makes it harder for site owners to manipulate search results to their advantage. This is why Google frequently changes its algorithm. A by-product of this policy is that certain sites gain and fall out of favour in a seemingly arbitrary fashion (Pasquale 2011).

Third, representativeness of the database with respect to its content and user access to it: do the results cover the whole world? The developed world? The English-speaking world? Is the site giving users access to all the results, or just a subset? Does the database change depending on where a user logs on or who she is? Knowing who is contributing and which set of results we are seeing can help us evaluate the trust we should put in the algorithm.

There is a dispute in the epistemology of testimony as to whether a recipient of a testimony has a default epistemic right to believe it without positive reasons to do so; all parties to the debate agree, however, that justified belief formation requires the recipient to consciously or subconsciously monitor the testimony for trustworthiness indications (Goldberg 2007: ch. 6). Whether it constitutes testimony or not,⁸ we advocate that a similar requirement be placed on justified belief formation based on internet-obtained information.

But as the analysis in this section reveals, the opacity of filtering algorithms leaves users without much indication of the quality of information they encounter. In brief, we are concerned with cases in which subjects must rely mostly or only on automatically filtered internet information, and justified belief formation requires from them relatively rigid monitoring or critical assessment. In such cases, it is impracticable for subjects to directly assess the bias and completeness of their information sources because of the secret nature of internet algorithms. That is, because they cannot see how or why results are ordered the way they are, cannot know which, if any, results within the database have been excluded, and cannot know which potential results were never in the database to begin with, *it is not technologically possible* for the subjects to adjudicate internet search results in ways we might think are appropriate. Subjects in such scenarios will fail to form justified beliefs where, under our taxonomy, their failure is of type JF3, that is, resulting from the impracticability of the investigative activity required for achieving justification.

To be clear, it is not always the case that bubbled internet information results in JF3. In cases like the one posed above, subjects must fulfill relatively rigid monitoring or critical assessment requirements before a belief can be considered responsible. When internet

⁸ See n. 3 above.

filters exclude information that would be valuable for monitoring, such as minority positions, it is clear that the technology is preventing subjects from acquiring justification. But in cases where monitoring is not required, subjects can acquire justification even when filters provide only partial or otherwise epistemically problematic evidence.⁹

Beliefs obtained from internet sources would be more justified, then, if (1) results were not bubbled, (2) users had more control of how filters worked, (3) users were made aware of the concrete limitations and strengths of search and features of algorithms, or (4) internet sources could be checked against sources for which epistemic safety measures were practicable. When none of these four options is the case, beliefs formed based on internet-filtered sources face a genuine danger of being unjustified. As we will argue in the next section, option (4) is in fact a typical one. Luckily, we do not live in a bubbled world, thus subjects could be reasonably expected to make further investigations, either on- or offline, and this would add strength to the justification of beliefs formed on the basis of internet use – even if the conclusions reached were eventually the same.

6. JUSTIFIED BELIEFS IN A DIGITAL AGE

As mentioned in the previous section, we are concerned with cases in which justified belief formation requires subjects to assess the bias and completeness of their information sources, but it is impracticable to check them directly because of the secret nature of internet algorithms. We have argued that when a subject's beliefs are formed mostly on the basis of internet-filtered information, her lack of knowledge of how filtering works detracts from their justificatory status. In this section, we argue that subjects can typically compensate for this lack of knowledge and enhance the justificatory status of their putative beliefs by other means. In accordance with our claim that practicability plays a major role in defining the minimal standards of justified belief, we argue that because it is both practicable and relatively easy to take these extra steps, completing them will typically be part of subjects' epistemic responsibilities. Therefore, belief formation without completing them will result in justification failure of type JF_I, failing to fulfill epistemic responsibilities.

As we have argued, even casual surfers have some responsibility to ascertain whether information is biased or incomplete, but their responsibility may vary depending on the case. For example, Goldman (2008) argues that traditional news sources, such as newspapers, magazines and TV news shows, tend much more than blogs to filter reports at the stage of reporting by employing methods such as professional fact checking. This means that the stories they publish have already passed some quality tests. Within our framework, such practices relieve traditional news-source consumers from some of their responsibilities to check the facts for themselves, whereas blogs tend much more to delegate this responsibility to their readers.¹⁰

As we saw at the start, and as Goldman's analysis suggests, internet sources often present to their readers information that is biased or incomplete. This problem is aggravated when this information is automatically personalized and filtered in a way that does not

9 We thank an anonymous reviewer for raising this point.

10 Goldman (2008) is skeptical about readers' ability to successfully do that, while Coady (2011) is more optimistic about it and less impressed by the positive effects of journalists' professionalism.

necessarily aim at satisfying epistemic desiderata such as reliability, objectivity, scope and truth. Moreover, as we have argued, internet filtering algorithms are often opaque, hiding from their users key details of their workings that are relevant to their epistemic assessment. Thus, they prevent subjects from knowing whether information is biased or incomplete, and worse may give the appearance of unbiased and complete information. This means that sole reliance on internet sources, especially bubbled sources, for acquiring beliefs without making an additional effort to validate and evaluate them, will typically not satisfy the requirements of epistemic responsibility required for justification.

If users cannot know the relevant details of the workings of internet algorithms on which they rely, particularly filtering algorithms, what can they do to improve the justificatory status of their beliefs? Quite a lot, actually. First, subjects need not know the fine details of the operation of filtering algorithms to be epistemically secure. They need only be aware that bubbles and bias are possibilities they must protect against (cf. Williams 2008).¹¹ Furthermore, in verifying some kinds of information, subjects can use existing competencies for gaining information from traditional media such as newspapers to supplement internet-filtered information and therefore at least partly satisfy the responsibility to determine whether it is biased or incomplete. With respect to news and politics, there are many available non-algorithmic sources or sources that have both print and online versions, such as newspapers and television, which users can use to supplement and validate internet-filtered sources. Blog readers can thus make it an occasional habit to read the news headlines on major news sites, or watch the evening TV news. Not only is this practicable, it is easy to do. Therefore, a responsible epistemic agent should actively expose herself to a variety of sources.

What about social networking sites such as Facebook? Social networking enables us to stay informed about a large group of people, in a way that would be impracticable without it. People do not normally have the time to stay in touch or keep up-to-date with a group of friends as large as they typically have on Facebook using traditional methods such as face-to-face meetings, phone calls and emails. Social networking makes it technologically possible to obtain information that would be technologically impossible without it.

As we saw at the start, however, the Facebook news feed can be biased, tending to show stories from like-minded friends. Can a responsible epistemic subject do better than rely solely on it? Yes, he can. To improve his epistemic standing, a subject is not required to start phoning his friends who are not like-minded and ask them for their opinions on various matters. Instead he can casually visit their Facebook profiles and see whether they have posted an interesting story that the automatically generated news feed missed. Such a practice has an added benefit. While the fine details of the implementation of filtering algorithms is unknown and subject to change, there are good reasons to think that if a person clicks on such stories enough times, they will increasingly appear in his news feed, and perhaps in the news feed of users with similar profiles as well.¹² If Facebook only had an automatically generated news feed, without the option to independently browse through friends' profiles,

11 Similar reasoning has long been accepted in statistics. One can measure and even partially correct for errors in measurement even in cases where the cause of the error is unknown.

12 Sites like Facebook implement collaborative filtering algorithms. There are two main kinds of such algorithms: memory-based collaborative filtering, where a user is being recommended items that were liked by similar users, and model-based collaborative filtering, where machine learning techniques are used to predict a user's preferences from her past choices. For a technical overview, see Su and Khoshgoftaar (2009); for an epistemic analysis, see Origg (2012).

and given that the workings of the news-feed filtering algorithm are secret, it might be that beliefs formed based on the feed would be as justified as they could be, because it would be impracticable for a subject to run further checks.

If none of these avenues is available, it is of course permissible for subjects to make some of the usual epistemic backpedals: information that is insufficient to justify a belief *tout court* may nevertheless suffice for updating a subject's degree of belief. Alternatively, subjects may decide to formulate a different kind of belief, of the form, 'according to X , p ,' (which is not to form a belief about p at all). These suggestions amount to cautioning subjects to simply be more careful in formulating their beliefs when they are relying primarily on internet sources, which is ultimately not very satisfying.

One last possibility is worth mentioning. The practices discussed in this section assume that Google's and Facebook's algorithms stay secret or non-transparent. If they were made available for public scrutiny, experts could answer, with some degree of certainty, how much and under what conditions filtered results are sufficient. Indeed, some of this work can be done without the cooperation of Google and Facebook – including the research that was the basis for many of the worries expressed in section 2. Content providers may also enhance the transparency of the result-generating processes to their users. For example, Simon (2010: 353–4) suggests that search engines have two search buttons, one for personalized search and one for non-personalized search, and users can compare the differences between the two sets of results. Additionally, as Sunstein (2007: 208–10) suggests, internet sites, such as political blogs, may refer their readers to alternative views, for example, by linking to opposing sites, out of a commitment to pluralism.

To conclude the argument in this section: users who rely on automatically filtered information, particularly personalized information, to form beliefs lack knowledge about the workings of the filtering methods that is relevant to their epistemic assessment. They can, however, partly compensate for this lack of knowledge by obtaining information from other sources, whether online or off. Because obtaining such information is practicable and relatively easy, it will typically be part of the responsible epistemic conduct minimally required for justified belief.

One may raise the following objection: many subjects are unaware of the filter bubble or that they are in a filter bubble. Thus, they may not have a reason to suspect that they need to take further action to justify their beliefs. Hence, they are not responsible for taking these actions. As we noted at the start, however, ignorance may be culpable, and we think it typically will be in this case. Our point is that people do not know enough about the way the technologies they use to obtain information work to assess the beliefs they form based on them, whereas responsible agents are required to ensure, within the limits of what is practicable and required for the role they assume, that the information is unbiased, sufficiently reliable, etc. 'I read it in the first site that appeared in my Google search results' is normally not a good reply to the question 'how do you know?' The filter bubble just proves the point. It shows that there are indeed real and unforeseen epistemic dangers in blindly relying on internet sources for forming beliefs.

7. CONCLUSION

The responsibilist view of justified belief presented here states that a subject who aims to form beliefs that are true or rational will have greater or lesser epistemic responsibilities

according to her particular situation, but at a typical minimum has the responsibility to do what she can to ensure that the information upon which she forms her beliefs is unbiased and complete enough to underwrite the judgments she aims to make. We argued that epistemic justification bears close conceptual relations to practicable action, which depends in turn on the features of the particular technologies available to the subject. Drawing on these conceptual relations allowed us to outline the standards of justified belief contemporary internet users need to meet.

We have observed that people, especially young people, increasingly form beliefs on the basis of information gained primarily or exclusively from internet sources. The inner workings of such sources are often opaque, preventing subjects from knowing whether information is biased or incomplete. Furthermore, the phenomenon of filter bubbles gives us good reason to think that some internet-derived information actually is biased and incomplete. Thus, beliefs formed mainly or merely based on filtered internet sources face a genuine danger of being unjustified; hence it may not be atypical that subjects' epistemic responsibilities will include taking reasonable precautions against this hazard.

Our intention is not to add an unreasonable burden to belief-forming activities; indeed, our theory explicitly states that subjects can only be responsible to fulfill investigations that are practicable. Among the practicable investigations open to most subjects is the use of existing competencies for gaining information from traditional media to supplement internet-filtered information in order to help determine whether it is biased or incomplete. Another approach is impracticable at present, but could be implemented by designers of search engines: to make available to users information about filters, algorithms and database scope.¹³

REFERENCES

- Alston, W. 1998. 'The Deontological Conception of Epistemic Justification.' *Philosophical Perspectives*, 2: 257–99.
- Annis, D. B. 1978. 'A Contextualist Theory of Epistemic Justification.' *American Philosophical Quarterly*, 15(3): 213–19.
- Baird, D. 2004. *Thing Knowledge: A Philosophy of Scientific Instruments*. Berkeley, CA: University of California Press.
- Bakshy, E. 2012. 'Rethinking Information Diversity in Networks.' www.facebook.com/notes/facebook-data-team/rethinking-information-diversity-in-networks/10150503499618859.
- Clark, A. and Chalmers, D. 1998. 'The Extended Mind.' *Analysis*, 58(1): 7–19.
- Clark, A. 2010. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Oxford: OUP.

13 We would like to thank audiences at the Canadian Communication Association Annual Meeting (Waterloo, Canada, 2012), the Episteme Conference (Delft, the Netherlands, 2012), and the Seminar at the Cohn Institute for the History and Philosophy of Science and Ideas (Tel Aviv, 2012) for their comments on portions of this paper. We also owe debts to Anat Ben David, Arnon Keren, Jacob Stegenga, Eleanor Louson, Galit Wellner, and an anonymous referee for comments and suggestions that have substantially improved the paper. We thank Sandy Goldberg, Alvin Goldman, Andy Rebera, Jeroen van den Hoven and Ehud Lamm for helpful discussion and suggestions. This paper was partly written while Boaz Miller was an Azrieli Post-Doctoral Fellow at the Dept of Philosophy, University of Haifa. Boaz is grateful to the Azrieli Foundation for an award of an Azrieli Fellowship.

- Coady, D. 2011. 'An Epistemic Defence of the Blogosphere.' *Journal of Applied Philosophy*, 28(3): 277–94.
- Colombo, F. and Fortunati, L. (eds). 2011. *Broadband Society and Generational Changes*. Frankfurt am Main: Peter Lang.
- Conee, E. and Feldman, R. 2004. *Evidentialism: Essays in Epistemology*. Oxford: Clarendon Press.
- Douglas, H. 2009. *Science, Policy, and the Value-Free Ideal*. Pittsburgh, PA: University of Pittsburgh Press.
- Fantl, J. and McGrath, M. 2009. *Knowledge in an Uncertain World*. Oxford: OUP.
- Foley, R. 2005. 'Justified Belief as Responsible Belief.' In M. Steup and E. Sosa (eds), *Contemporary Debates in Epistemology*, pp. 313–26. Malden, MA: Blackwell.
- Giere, R. N. 2006. 'The Role of Agency in Distributed Cognitive Systems.' *Philosophy of Science*, 73(5): 710–19.
- . 2007. 'Distributed Cognition without Distributed Knowing.' *Social Epistemology*, 21(3): 313–20.
- Goldberg, S. C. 2007. *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge: CUP.
- . 2012. 'Epistemic Extendedness, Testimony, and the Epistemology of Instrument-Based Belief.' *Philosophical Explorations*, 15(2): 181–97.
- Goldman, A. I. 1999. 'Internalism Exposed.' *Journal of Philosophy*, 96(6): 271–93.
- . 2008. 'The Social Epistemology of Blogging.' In J. van den Hoven and J. Weckert (eds), *Information Technology and Moral Philosophy*, pp. 111–22. Cambridge: CUP.
- . 2011. 'Reliabilism.' In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2011 edn): <http://plato.stanford.edu/archives/spr2011/entries/reliabilism>.
- Haraway, D. 1991. 'A Cyborg Manifesto: Science, Technology, and Socialist-Feminism in the Late Twentieth Century.' In *Simians, Cyborgs and Women: The Reinvention of Nature*, pp. 149–81. New York: Routledge.
- Hardwig, J. 1985. 'Epistemic Dependence.' *Journal of Philosophy*, 82(7): 335–49.
- Humphreys, P. 2004. *Extending Ourselves: Computational Science, Empiricism, and Scientific Method*. New York: OUP.
- . 2009. 'Network Epistemology.' *Episteme*, 6(2): 221–9.
- Introna, L. and Nissenbaum, H. 2000. 'Shaping the Web: Why the Politics of Search Engines Matters.' *Information Society*, 16(3): 1–17.
- Knorr-Cetina, K. 1999. *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge, MA: Harvard University Press.
- Kornblith, H. 1983. 'Justified Belief and Epistemically Responsible Action.' *Philosophical Review*, 92(1): 33–48.
- Latour, B. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: OUP.
- Lehrer, K. 1995. 'Knowledge and the Trustworthiness of Instruments.' *The Monist*, 78(2): 156–70.
- Longino, H. 2002. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.
- McCroy, P. 2005. 'The Time Lords: Measurement and Performance in Sprinting.' *British Journal of Sports Medicine*, 39: 785–6.
- Origgi, G. 2012. 'Designing Wisdom through the Web: The Passion of Ranking.' In H. Landemore and J. Elster (eds), *Collective Wisdom: Principles and Mechanisms*, pp. 38–55. Cambridge: CUP.
- Pariser, E. 2011. *The Filter Bubble: What the Internet is Hiding from You*. London: Penguin.
- Pasquale, F. 2011. 'Restoring Transparency to Automated Authority.' *Journal on Telecommunications and High Technology Law*, 9: 235–53.
- Preston, J. 2010. 'The Extended Mind, the Concept of Belief, and Epistemic Credit.' In R. Menary (ed.), *The Extended Mind*, pp. 355–69. Cambridge, MA: MIT Press.
- Record, I. Forthcoming. 'Technology and Knowledge.'
- Rogers, R. 2004. *Information Politics on the Web*. Cambridge, MA: MIT Press.
- Rothbart, D. 2007. *Philosophical Instruments: Minds and Tools at Work*. Urbana, IL: University of Illinois Press.
- Simon, J. 2010. 'The Entanglement of Trust and Knowledge on the Web.' *Ethics and Information Technology*, 12(4): 343–55.
- Simpson, T. W. 2012. 'Evaluating Google as an Epistemic Tool.' *Metaphilosophy*, 43(4): 426–45.
- Simson, R. S. 1993. 'Values, Circumstances, and Epistemic Justification.' *Southern Journal of Philosophy*, 31(3): 373–91.

- Sosa, E. 2006. 'Knowledge: Instrumental and Testimonial.' In E. Sosa and J. Lackey (eds), *The Epistemology of Testimony*, pp. 116–23. New York: OUP.
- Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: OUP.
- Su, X. and Khoshgoftaar, T. M. 2009. 'A Survey of Collaborative Filtering Techniques.' *Advances in Artificial Intelligence*: <http://dx.doi.org/10.1155/2009/421425>.
- Sunstein, C. 2007. *Republic.com 2.0*. Princeton, NJ: Princeton University Press.
- Tollefsen, D. P. 2009. 'Wikipedia and the Epistemology of Testimony.' *Episteme*, 6: 8–24.
- Williams, M. 2001. *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: OUP.
- . 2008. 'Responsibility and Reliability.' *Philosophical Papers*, 37(1): 1–26.

BOAZ MILLER is a Postdoctoral Fellow at the Edmond J. Safra Center for Ethics, Tel Aviv University, and a teaching fellow at the Science, Technology, and Society (STS) program at Bar Ilan University. He has a PhD from the Institute for the History and Philosophy of Science and Technology (IHPST), University of Toronto. His main research interests are in social epistemology.

ISAAC RECORD is a Postdoctoral Research Fellow at the Semaphore Lab, Faculty of Information, University of Toronto. He has a PhD from the Institute for the History and Philosophy of Science and Technology (IHPST), University of Toronto. His main research interests are in philosophy of technology and epistemology.
