



SERRC

Social Epistemology
Review & Reply Collective

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

Intellectual Virtues and Internet-Extended Knowledge

Paul R. Smart, University of Southampton, ps02v@ecs.soton.ac.uk; Robert W. Clowes,
Universidade Nova de Lisboa, robert.clowes@gmail.com

Smart, Paul R. and Robert W. Clowes. 2021. "Intellectual Virtues and Internet-Extended Knowledge." *Social Epistemology Review and Reply Collective* 10 (1): 7-21.
<https://wp.me/p1Bfg0-5AY>.

Introduction

We are grateful for Lukas Schwengerer's contribution to the topic of Internet-extended knowledge. We greatly enjoyed reading his paper, "Online Intellectual Virtues and the Extended Mind," which was the basis for many of the ideas rehearsed in the present paper.

Schwengerer's paper explores the extent to which the online environment (and, in particular, Google Search) might support extended knowledge. As noted by a number of scholars, claims about the extended mind suggest the possibility of extended knowers—individuals whose epistemic states are tied to the operation of extra-neural (and/or extra-organismic) processing loops (e.g., Bjerring and Pedersen 2014). Given the classic thought experiment presented by Clark and Chalmers (1998)—the one involving a neurologically impaired individual, named Otto—it seems that the exploitation of a bio-external resource (e.g., a paper-based notebook) might provide the basis for extended knowledge. Rupert (2004), for example, notes that if Otto's beliefs are allowed to extend to his notebook, then the same might be true of Otto's knowledge. Similarly, Palermos (2018) suggests that:

[...] from the commonsense functionalist point of view, it is a small step from claiming that there are dispositional beliefs to claiming that we know the contents of our smartphones, laptops, or websites even before looking them up (459).

In view of this, it seems that if we allow for the possibility of Internet-extended minds (i.e., minds that extend to the Internet), then we also open the door to Internet-extended knowledge.

Schwengerer highlights some of the problems confronting this proposal. His paper is framed by the claim that the Internet is a potential source of epistemic harm (e.g., misinformation, filter bubbles, echo chambers, and so on). Schwengerer suggests that online intellectual virtues are a potential solution to these epistemic harms. Online intellectual virtues are glossed as instances of the more general intellectual virtues (e.g., curiosity, intellectual autonomy, open-mindedness, etc.) that feature as part of virtue responsibilist accounts of knowledge (see Battaly 2008). Accordingly, Schwengerer embraces a virtue responsibilist approach to Internet-extended knowledge.

At this point, we encounter a problem in marrying virtue responsibilism with a particular conception of the extended mind that is attributable to Clark and Chalmers (1998).

According to Schwengerer, Clark and Chalmers provide a so-called first-wave account of the extended mind, which attaches considerable importance to the notion of automatic endorsement (i.e., the uncritical acceptance of bio-external information). This is deemed to be incompatible with a virtue responsibilist approach to the Internet, i.e., an approach that mandates the careful scrutiny of online (i.e., bio-external) information. Assuming that we want to maintain the possibility of Internet-extended knowledge, and we also want to hold

onto the notion of online intellectual virtues, then we appear to confront a choice: either we reject the possibility of Internet-extended knowledge, or we adopt a modified view of the extended mind. Schwengerer pursues the latter option. In particular, he calls for a multidimensional approach to cognitive integration, which yields an alternative way of evaluating claims about the extended mind.

Our response to Schwengerer's proposal is organized as follows.

First, we question the need for intellectual virtue in the case of Internet-extended knowledge. We suggest that not all online systems are necessarily prone to epistemic shortcomings, and that intellectual virtues are not the only means by which the epistemic shortcomings of online systems might be addressed.

Second, we question the extent to which a multidimensional approach to cognitive integration can help to resolve the apparent tension between virtue responsibilism and the extended mind. In particular, we question whether the tension can be restricted to a single dimension of the multidimensional framework countenanced by Schwengerer. Our view is that the tension applies to multiple dimensions of this framework and is therefore something of a multidimensional phenomenon.

Third, we suggest a means by which a virtue responsibilist account of knowledge might be reconciled with claims of Internet-extended knowledge. We suggest that the exercise of intellectual virtue plays a potentially important role in the developmental emergence of Internet-extended knowers. According to our proposal, the exercise of intellectual virtue can yield knowledge about the trustworthiness of online systems, as well as the circumstances in which the automatic endorsement of online content is epistemically justified.

Online Epistemic Hazards: Fake News about Fake News?

Schwengerer's appeal to intellectual virtue is based on concerns about the epistemic hazards of the Internet. These include worries about filter bubbles and echo chambers (Nguyen 2020), deep fake capabilities (Fallis, in press), and various forms of misinformation. There is, to be sure, plenty of public concern about these supposed threats to our epistemic standing. But do these concerns adequately reflect the nature of the epistemic hazards we face online? Is the nature of our cognitive and epistemic contact with the online environment sufficiently hazardous as to warrant the exercise of intellectual virtue in securing Internet-extended knowledge?

We suggest that it is difficult to answer these questions in a general sense (i.e., to make general claims about the Internet as a whole). The Internet, we suggest, ought to be seen as a complex ecosystem, populated by a multiplicity of functionally diverse online systems that feature varying degrees of reliability. Schwengerer directs his attention to a particular online system, namely, Google Search, but it should be clear that the term "Internet" subsumes a much broader array of systems than Google Search. In this sense, it is hard to make general claims about the prospects for Internet-extended knowledge based on the analysis of a single

online system. It is also unclear whether the issues raised in respect of a particular system (in this case, Google Search) are sufficient to warrant sweeping claims about how we ought to engage with the Internet. Just because intellectual virtue is required in the case of Google Search (if, indeed, it is required), why assume that *all* our interactions and engagements with the online environment ought to be governed by the exercise of intellectual virtue?

While Google Search is a popular target of epistemological analysis, it is not the only system that is deserving of epistemological attention. (Nor, we might add, is it necessarily the most suitable system for evaluating claims of Internet-extended knowledge.) Some online systems feature various checks and balances that are intended to improve the reliability of online content. Wikipedia, for example, includes a so-called ‘immune system’ to mitigate the threat of malign interference (Halfaker and Riedl 2012). Inasmuch as these interventions work to ensure the reliability of online content, then it is not clear that *every* form of interaction and engagement with the Internet ought to be governed by the exercise of intellectual virtue. Is the exercise of intellectual virtue still required if an online system has already taken steps to ensure the reliability of its informational deliverances? Might this not add needless complexity to our modes of interaction and engagement with the online world? What, moreover, of the largely implicit ‘neo-liberal’ (or, at any rate, individualist) leanings of the virtue responsibilist position? The virtue responsibilist assumes that it is individual citizens who ought to be responsible for their own cyber-epistemic well-being. But it is contentious as to whether individual citizens should be burdened with these responsibilities in all cases. Don’t governments also have a responsibility to ensure the *epistemic safety* of the online environment? And what about the role of big technology companies in delivering innovative solutions to reliability-related problems? Google’s parent company, Alphabet, had an estimated research and development budget of \$16.2 billion in 2018 (Jaruzelski, Chwalik, and Goehle 2018), so it’s not as if there is a shortage of cash to fund reliability-related research.

Web-based systems are not the only systems that might support Internet-extended knowledge. In recent years, there has been growing interest in a class of systems known as *cyber-physical systems* (Rajkumar, Niz, and Klein 2017). These systems are typically built around Internet of Things (IoT) devices, which rely on the Internet for the purposes of information exchange. Issues of reliability are a core concern for such systems, especially when they are used to support safety-critical functions such as electricity distribution and traffic control. A particular concern relates to the potential for a cyber-epistemic attack, wherein a malign actor attempts to manipulate the readings of an IoT sensor device so as to disrupt the functional operation of a cyber-physical system.¹ In the wake of such concerns, engineers have sought to develop techniques that can be used to ensure the epistemic integrity of Internet-mediated information flows. These include the use of background knowledge to detect anomalous or inconsistent sensor readings, as well as techniques that rely on cryptographic, watermarking, and redundancy-based solutions.²

¹ One example of this sort of attack is a so-called acoustic injection attack, which uses sound waves to manipulate the activity of sensors that are otherwise assumed to be reliable (e.g., Trippel et al. 2017).

² See Pfeifer and Bongard (2007, 260–261), for a discussion of redundancy in an IoT context.

It is not our aim in the present paper to survey (or evaluate) the mechanisms that might be used to ensure the epistemic integrity of online information—although that is surely a topic worthy of further consideration. Our aim is merely to highlight the need for a more refined epistemological approach to the Internet—one that recognizes the genuine diversity of online systems from an epistemological standpoint. This is important, for the term “Internet-extended knowledge” is perhaps best interpreted as a claim about the capacity of certain online systems to support extended knowledge. This seems preferable to an interpretation that sees the term as making a claim about the capacity of the Internet as a whole to support extended knowledge. In this sense, it should be clear that the failure of a single online system (e.g., Google Search) to satisfy the criteria for extended knowledge does not impugn the possibility of discovering (or deliberately engineering!) systems that *do* meet these criteria.

This has implications for philosophical efforts pertaining to extended knowledge and the extended mind. Suppose we discover that Google Search is not a particularly good candidate for Internet-extended knowledge according to a first-wave account of the extended mind. Can't the proponent of a first-wave account simply accept that Google Search is not a particularly good candidate for extended knowledge and then move on to other systems? Indeed, why begin the search for Internet-extended knowledge with Google Search? It is, after all, a search engine—a system whose main (functional) objective is to retrieve pointers to information that match a set of query conditions. In short, the goal of Google Search is to execute queries and return information; it is not its goal to check the veracity of the returned information. This more epistemologically-involved task is not so much the responsibility of the search engine as it is the *responsibility of the agent who executes the query*. Accordingly, perhaps we should not be surprised to discover that an epistemologically-oriented analysis of Google Search is already biasing us towards a virtue responsibilist conclusion. This reveals an important worry: The worry is that certain kinds of online system (including Google Search) are apt to favour a virtue responsibilist approach to knowledge simply because they delegate matters of truth and reliability to the human user of the system. Given this, it should be clear that we cannot use the analysis of these systems to laud the merits of a virtue responsibilist approach to the Internet *as a whole*. Nor should the epistemic properties of these systems (by themselves) serve as the basis for a wholesale revision of our philosophical conception of the extended mind.

There are, of course, good reasons to put Google Search under the epistemological spotlight. It is undoubtedly an important point of contact with the online environment, at least in the West. Its functionality has also evolved over the years to the point where it is perhaps no longer thought of as a pure search engine. Consider, for example, that Google Search can be used in the manner of a question and answering (QA) system. By typing a specific question into the query field (e.g., “What is the capital of Indonesia?”), we gain access to (hopefully correct) answers (i.e., “Jakarta”).³ In contrast to the more traditional (search-related) function

³ To some extent, then, Google Search ought not to be seen as a unitary system organized around a single functional imperative; it is, instead, a complex collection of different functionalities that can be used for a variety of epistemic purposes. This point is important, for it highlights that even within the context of a single online system (in this case, Google Search), we can often discover multiple kinds of epistemic functionality.

of Google Search, this QA-related capability seems to be a much more appropriate target for epistemological analysis.

One reason for this relates to the opportunities for empirical evaluation. By subjecting a QA system to empirical scrutiny, we can determine whether or not it is sufficiently reliable to warrant claims about its epistemic credentials (or lack thereof). Another reason for directing attention to Google’s QA capability is that it is possible to discern an epistemic goal for the system. In its capacity as a QA system, we rely on Google Search to furnish us with factually correct responses; we do not expect it to lead our epistemic projects and activities astray. In this sense, we see something akin to a commitment, albeit one that is largely implicit. We expect Google Search (*qua* QA system) to operate in a reliable manner, and we judge its success or failure in this light. This contrasts with the more conventional (information retrieval related) use of Google Search where issues of epistemic responsibility and commitment are, at best, unclear.

Question	Google	Bing	Siri
What year was the French Revolution?	1789	1789	1789
Which king built the Palace of Versailles?	Louis XIV	*	Louis XIV
What year did Pitt the Younger become prime minister?	1783	1783	*
Who was Louis XVI’s wife?	Marie Antoinette	Marie Antoinette	Marie Antoinette
What year did Louis XVI ascend to the throne?	1774	1774	*
Who founded the Methodist movement?	John Wesley	John Wesley	John Wesley
When did the Montgomery bus boycott begin?	5 December 1955	5 December 1955	5 December 1955
Who was Lyndon Johnson’s republican opponent in his landslide election?	Barry Goldwater	Barry Goldwater	*
What was Lyndon Johnson’s domestic policy called?	The Great Society	The Great Society	*
Who wrote “The Other America”?	Michael Harrington	Michael Harrington	Michael Harrington

Table 1: Table showing top-level responses to historical trivia questions on Google Search, Microsoft Bing, and Siri. [Asterisks represent unclear (but not necessarily incorrect) responses.]

In order to gain some insight into the reliability of Google Search (*qua* QA system), we presented it with the questions listed in Table 1. We also compared the results obtained from Google Search with those delivered by Microsoft Bing and Apple’s conversational assistant, Siri. All the results are presented in Table 1. To the best of our knowledge, the answers returned by Google Search are correct.

Thus, just as it is a mistake to assume that the epistemic properties of a single online system merit blanket conclusions about the epistemic properties of the Internet as a whole, so it is also a mistake to assume that the epistemic properties of a single system function (e.g., informational retrieval) merit blanket conclusions about the epistemic properties of the system as a whole.

It is, of course, impossible to generalize from the results of such a small-scale study, but on the basis of these results there seems little reason to doubt the reliability of Google Search's QA capabilities. This highlights the need to subject epistemic claims about the Internet (and its associated online systems) to empirical scrutiny. In particular, we should not assume that public fears and anxieties about the epistemic sequelae of online systems are a *reliable* source of information about the actual epistemic properties of those systems. To do this, is to absolve ourselves of intellectual virtue at precisely the point where intellectual virtue is most required: it is to accept some claim about the Internet as factually correct without subjecting this claim to proper scrutiny and evaluation.

In short, we worry about the reliability of some of the claims that fuel epistemological concerns about the Internet. If such claims should turn out to be nothing more than fake news—or fake news about fake news—then it is hard to see why intellectual virtue would be required to exploit the epistemic offerings of the Internet. We are, of course, not denying the existence of online epistemic hazards; we are merely calling for a more thorough (virtuous?) assessment of these hazards and (crucially) the circumstances under which they occur. It is only once we 'know' something about the epistemic properties of online systems that we will be in a position to advocate an epistemological response to them. Until then, we perhaps ought to adopt a position of open-mindedness. This, at least, seems like the virtuous thing to do.

It would, of course, be unfair to accuse Schwengerer of failing to provide any empirical support for his position. In respect of Google Search, Schwengerer refers to the work of Lynch (2016) who reports a creationist response to a question about the fate of the dinosaurs. Specifically, when Lynch presented Google Search with the question “What happened to the dinosaurs?”, he obtained the following result:

The Bible gives us a framework for explaining dinosaurs in terms of thousands of years of history, including the mystery of when they lived and what happened to them. Dinosaurs are used more than almost anything else to indoctrinate children and adults in the idea of millions of years of earth history (Lynch 2016, 66).

Schwengerer and other proponents of virtue responsibilism appear to regard this result as evidence of the epistemic shortcomings of Google Search (Heersmink 2018; Schwengerer, in press; Heersmink and Sutton 2020). We are not so sure. Firstly, it is difficult to see why a single erroneous result ought to be a cause for much concern. No doubt if we tried hard enough, we could extend the list of questions in Table 1 to the point where we too observed an incorrect response. But so what? We often rely on our bio-memory systems to service our epistemic interests, but this does not mean that they always do so. Just like many online circuits, our biologically-based information retrieval circuits sometimes let us down, but this does not mean that such circuits are incapable of sustaining states of knowledge.

A second worry relates to the nature of Lynch's query result. In particular, it is not immediately obvious that we are being presented with any sort of answer to our original

question. Given the response from Google Search, are we any the wiser about what happened to the dinosaurs? It seems that we are being referred to another source of information, namely, the Bible, but the Bible doesn't mention dinosaurs. What we seem to have here is a null result, as opposed to a false result—it is not so much an incorrect answer to our question as it is no answer at all.

In the interests of evaluating Lynch's query, we ran the same question ("What happened to the dinosaurs?") on our own machine. Here is the result we obtained:

Dinosaurs roamed the earth for 160 million years until their sudden demise some 65.5 million years ago, in an event now known as the Cretaceous-Tertiary, or K-T, extinction event.⁴

Evidently, the result is not the same.⁵ Why the difference? One possibility relates to the operation of personalized search algorithms.⁶ Perhaps, then, our own results reflect the nature of our (specifically, Smart's) search history. Note, however, that we obtained a similar result when we ran the query on a newly installed virtual machine with no search history.

A second possibility relates to the passage of time. Lynch ran his query circa 2016, whereas ours was performed on 20th October 2020. Perhaps, then, the difference is attributable to changes in Google's algorithms or the back-end database that drives query responses. This highlights the importance of revisiting past claims about the epistemic integrity of online systems. Just because a system is unreliable today, there is no guarantee that it will be unreliable tomorrow (the reverse is, of course, also true). Ideally, what we need, for the purpose of cyber-epistemological research, is a means of periodically checking the reliability of online systems. Such a capability could be used to evaluate the reliability of online systems and assess the need for intellectual virtue. Evidence of poor epistemic performance (low reliability) would, of course, support the need for epistemic caution (and thus the exercise of intellectual virtue) in respect of specific online systems. In addition, to this, however, it is also possible that variable performance across time (i.e., variable reliability) would also be a cause for concern.

If, for example, it turns out that the reliability of Google Search (*qua* QA system) varies from one year to the next, then it is, at best, unclear that we can absolve ourselves of the need to verify query results as part of our epistemic interactions with the system. In this case, there would surely be some merit in treating the informational deliverances of Google Search with a degree of caution, and any commitment to automatic endorsement (as per a first-wave

⁴ See <https://www.history.com/topics/pre-history/why-did-the-dinosaurs-die-out-1>.

⁵ On our machine, the result referred to by Lynch appears some way down the list of search results. In particular, it appears halfway down the second page of search results. It is hard to see why this would prove problematic, however (at least according to a first-wave approach to the extended mind). We are, after all, assuming people will endorse the first entry they see. Someone a bit more 'open-minded' might ferret around and find the Bible entry, but that is hardly an argument for being more open-minded!

⁶ Personalised search algorithms tailor search results to specific users based on a user's search history (see Simpson 2012, for more details).

account of the extended mind) would seem to serve as a rather poor foundation for Internet-extended knowledge. We could, at that point, seek to defend claims of Internet-extended knowledge by modifying our theoretical approach to the extended mind, or we could simply stick with our original (first-wave) conception of the extended mind and accept that certain kinds of online system (i.e., those exhibiting low or variable reliability) are not particularly good candidates for Internet-extended knowledge. Schwengerer adopts the first of these options, but it is not clear (to us at least) that the second option ought to be removed from the table.

The Multidimensional Nature of the Extended Knowledge Dilemma

From an epistemic standpoint, Schwengerer regards the online environment as sufficiently hazardous as to mandate the need for intellectual virtue, at least in the case of Google Search. Schwengerer then goes on to identify an apparent tension between virtue responsibilism and a first-wave approach to the extended mind. The problem centres on the notion of automatic endorsement, which features as one of the “rough-and-ready” criteria used to evaluate putative cases of cognitive extension (Clark 2010). These criteria have come to be known as the trust+glue criteria, and the criterion that interests us in the present paper is what we will dub the trust criterion. Here is how Clark (2010) presents the trust criterion:

[...] any information [...] retrieved [from a bio-external resource should] be more or less automatically endorsed. It should not usually be subject to critical scrutiny (unlike the opinions of other people, for example). It should be deemed about as trustworthy as something retrieved clearly from biological memory (46).

As noted by Schwengerer, this criterion presents a problem for a virtue responsibilist account of Internet-extended knowledge. According to Schwengerer, the exercise of intellectual virtue requires the critical evaluation of online content, but this is difficult to reconcile with an approach that is committed to the automatic endorsement of such content:

Crucially, for information provided by a tool to be part of my extended mind, it has to be automatically endorsed [...] However, that stands in a clear conflict with how an intellectually virtuous agent ought to interact with Google Search. Online virtues demand a critical eye on information on the Web! [...] In particular, an intellectual [sic] virtuous agent using the Internet will not take Google results at face value. They will keep the conditions of the search results in mind and evaluate which result deserves to be trusted. Moreover, they will not merely take the first result but be open-minded to other results that are not ranked as highly (Schwengerer, in press, 5).

We thus appear to have a tension between the notion of automatic endorsement (which is certainly a feature of first-wave accounts of the extended mind) and the possibility of a virtue responsibilist approach to Internet-extended knowledge. This tension has been the subject

Schwengerer.⁷ Figure 1a depicts a state-of-affairs in which we encounter high levels of trust (and thus automatic endorsement) in an online system, whereas Figure 1b depicts a state-of-affairs in which we see only moderate levels of trust. Despite these differences, we arguably have the same level (or degree) of cognitive integration in both cases. This is because the score on the durability dimension in Figure 1b is elevated relative to that seen in Figure 1a. An increase in durability thus compensates for a reduction in the trust dimension, yielding more or less the same level of cognitive integration for both systems.

The merits of this approach for a virtue responsibilist account of Internet-extended knowledge should be relatively clear. A system such as Google Search, for example, might qualify as being cognitively integrated on account of the fact that it scores highly on the accessibility and individualization dimensions, despite the fact that it scores low on the trust dimension. Accordingly, we can absolve ourselves of the commitment to automatic endorsement without renegeing on the idea that Google Search might serve as the basis for Internet-extended knowledge. As noted by Schwengerer (in press):

While the application of intellectual virtues to the online world might give us restrictions on some dimensions, it has little impact on others. Hence, there might be room for combinations of agents and artifacts that still count as showing an overall high degree of cognitive integration, while being virtuous epistemic agents (8).

Can this approach be made to work? Unfortunately, we doubt this to be the case. One problem concerns the *multidimensional* nature of the extended knowledge dilemma—the idea that the extended knowledge dilemma might apply to multiple dimensions of Heersmink’s framework.

To help us understand this problem, let us direct our attention to the individualization dimension. According to Heersmink and Sutton (2020) individualization refers to the extent to which the informational deliverances of a bio-external resource have been tailored to suit the specific needs, concerns, and cognitive practices of an individual agent. This is deemed to be important because it “often streamlines a cognitive task, making it easier and faster to perform” (Heersmink and Sutton 2020, 154).

In the context of Google Search, however, there seems to be a direct parallel between individualization and personalized search, with personalized search algorithms tailoring search results to specific users.⁸ This looks to be problematic, given that personalized search has been denigrated on account of its potential to cause epistemic harm. Simpson (2012), for example, suggests that personalisation is apt to lead to biases that undermine the epistemic standing of Internet users (specifically, the users of Google Search). The upshot is that we cannot appeal to individualization as a means of circumventing the commitment to

⁷ According to Schwengerer (in press, 7), the trust dimension measures “the degree to which one takes the information provided by an artifact to be correct.”

⁸ See Clowes (in press) for a discussion of the relationship between personalization and individualization.

automatic endorsement. That is to say, we cannot compensate for a reduction in trust by elevating the importance of individualization, at least in the case of Google Search.

There are reasons to think that this sort of problem applies to more than the individualization dimension, for a similar sort of concern arises in respect of the notion of procedural transparency. Schwengerer interprets the dimension of procedural transparency as indicating “the degree of fluency and effortlessness in interacting with an artifact” (7). But this appeal to fluency raises worries about a so-called truth bias in which the ‘truth’ of some body of external information is judged relative to the subjective ease with which it is processed (see Smart 2018, 297).⁹ As with individualization, what we seem to confront here is a basic worry about the extent to which low scores on the trust dimension can be substituted for high scores on another dimension (in this case, the dimension of procedural transparency).

The upshot is a general concern about the extent to which a multidimensional approach to cognitive integration can be used to resolve the extended knowledge dilemma. This concern is captured by a generalized version of the extended knowledge dilemma, which we have discussed in previous work (see Smart, Clowes, and Heersmink 2017, 80–81):

The Extended Knowledge Dilemma (Generalized)

The properties that work to ensure that an external resource can be treated as a candidate for cognitive incorporation are also, at least in some cases, the very same properties that work to undermine or endanger the positive epistemic standing of the technologically-extended agent.

This generalized version of the extended knowledge dilemma undermines the extent to which a multidimensional approach can be used to secure a virtue responsibilist approach to Internet-extended knowledge. For even if we accept the idea that a multidimensional approach reduces the need for automatic endorsement (or trust) we still need to show that compensatory adjustments to other dimensions do not cause problems for the epistemic standing of would-be knowers. If virtue responsibilism requires us to reduce the scores we assign to multiple dimensions of Heersmink’s framework (e.g., trust, transparency, and individualization), then the prospects for Internet-extended knowledge (or, indeed, any form of extended knowledge) may start to look a little dim.

Internet Extended Knowledge and Virtue Responsibilism

In contrast to Schwengerer, we doubt that a virtue responsibilist approach to Internet-extended knowledge will be forthcoming.¹⁰ This, however, does not mean that there is no room for intellectual virtue in our understanding of extended epistemic systems, including

⁹ See Wheeler (2019) for some additional worries about transparency-related conditions.

¹⁰ See Pritchard (2018) for a discussion of some of the problems confronting the attempt to provide a virtue responsibilist account of extended knowledge.

those that are built around our interactions and engagements with the online environment. Our own proposal relates to the developmental emergence (and perhaps maintenance) of Internet-extended knowers. In particular, we suggest that the exercise of intellectual virtue is likely to play an important role in guiding our decisions about which online systems are to be trusted, as well as the circumstances in which such trust is warranted.

In short, we see a role for intellectual virtue in the acquisition of a particular kind of knowledge, namely, knowledge about the trustworthiness (or reliability) of particular online systems for particular epistemic purposes. It is this knowledge that makes the automatic endorsement of online content epistemically justified. In short, if we know that a system is reliable, courtesy of the prior exercise of intellectual virtue, then we can proceed to automatically endorse its informational deliverances, and it is at this point that issues of Internet-extended knowledge start to come to the fore. By adopting this sort of ‘developmental’ perspective, we avoid many of the problems associated with the extended knowledge dilemma: We can accept that intellectual virtue plays a causal (developmental) role in the formation of extended epistemic systems, but we do not need to see the exercise of intellectual virtue as an intrinsic part of the operation of such systems.

There is an interesting parallel here with issues of trust and trustworthiness in the context of interpersonal relationships. In general, we do not trust people we do not know (strangers) with important matters, and even people we do know well (friends and family members) are not trusted with regard to every matter (see Hardin 2002).¹¹ The same perhaps applies to online systems. We trust some online systems, but not others, and even those systems we do trust are not trusted with regard to every matter. From an epistemic standpoint, I may trust Google Search in its capacity as a QA system, but I may not trust it in its capacity as an information retrieval system (recall the earlier distinction between these functionalities). Accordingly, when I use Google Search as a QA system, I automatically endorse its informational deliverances, but I am much more circumspect about its informational deliverances when I use it as an information retrieval system. In the former case, I form a cognitively-potent bio-technological union with Google Search that serves as the basis for Internet-extended knowledge, while in the latter case I do not.

This complicates our understanding of cognitive and epistemic incorporation as it applies to the online realm, for it should be clear that we cannot make blanket assertions about the extent to which a particular online system is a candidate for cognitive incorporation; this will vary according to the nature of an individual’s interaction with the system in particular circumstances, which will in turn depend on the individual’s background beliefs about the trustworthiness (or reliability) of the system in those circumstances. Whether an individual’s trust-related beliefs amount to knowledge is, of course, unclear, although individuals clearly have an interest in ensuring that they *do* amount of knowledge. If an individual’s trust-related beliefs are incorrect, then there is a chance that their thoughts and actions will not be

¹¹ Trust is typically conceptualized as a three-place relation, such that X (trustor) trusts Y (trustee) to do Z (activity) (Hardin 2002). In the case of Google Search, we might say that Paul (X) trusts Google Search (Y) to deliver correct responses to historical trivia questions (Z).

coordinated with respect to the factive structure of reality. Accordingly, they may be led astray, just as they may be led astray if they believe that another human individual is trustworthy when in fact they are not.

Given the stakes involved, it clearly pays to pick and choose our relationships with significant others wisely (regardless of whether those significant others are lovers, friends, or online systems). The exercise of intellectual virtue is perhaps important in delivering that wisdom, but it need not be a feature of the interactions and exchanges that occur as part of an established, trusted relationship. Once we ‘know’ that a friend or lover is trustworthy, we do not need to monitor and evaluate their every move. We just trust them. If monitoring and evaluation were required to be a feature of every social exchange, then our social interactions and engagements would be a great deal more complicated than they actually are. The same may be true of our relationships with online systems. It clearly pays to know something about the trustworthiness of online systems before we expose ourselves to the promise and peril of a particular bio-technological union, and the exercise of intellectual virtue is no doubt one of the means by which we are able to acquire such knowledge. But once the bio-technological bond is established, and the extended cognitive circuits are in play, the intellectual virtues are perhaps more of a hindrance than a help. As in the social realm, critical scrutiny is not so much the glue that holds a trusted relationship together; it is more a sign that a trusted relationship has not yet formed, or perhaps that a previously trusted relationship is about to come to an end.

Conclusion

Schwengerer’s article marks an important contribution to the emerging literature on Internet-extended knowledge, as well as the more general debates pertaining to the epistemic impact of the online environment. Despite the critical tenor of the present article, we encourage further research into these issues, including the role of virtue-theoretic conceptions of knowledge in guiding cyber-epistemic policy. Global governments are currently embarking on regulatory and legislative programmes that will influence our modes of interaction with the Internet, as well as define the roles and responsibilities of various social actors in mitigating the threat of epistemic and other online harms (e.g., Government of the United Kingdom 2019). Epistemologists clearly have an important role to play in this debate, and we hope that Schwengerer’s article, in conjunction with our own contribution, provides epistemologists with some interesting focus areas for future research.

It should also be clear that work in this area has a deeply interdisciplinary flavour. It draws on the expertise of scientists, engineers, and philosophers whose interests span a variety of topics. As per the focus of discussion in the present paper, these topics include the mechanics of cyber-security protocols, the provision of research tools to evaluate the reliability of online systems, a consideration of the political leanings of different epistemological positions relative to global cyber-policy initiatives, a better understanding of the role of trust and trustworthiness in debates about the extended mind and extended epistemology, and the role of intellectual virtue in evaluating claims about the epistemic

properties of the online environment. (All of this, of course, needs to be aligned with philosophical work regarding the nature of knowledge and the possibility of extended minds.) Research in these areas is particularly timely, given the continuing growth of the Internet and the (perhaps not entirely coincidental) concern with issues of misinformation, truth, and various forms of fakery. In all likelihood, the COVID-19 pandemic will accentuate existing worries and fears about the epistemic impact of the Internet. At the time of writing, a number of COVID-19 vaccines are available, but there is considerable concern about the potential of anti-vaccination messages to thwart global vaccination efforts (e.g., Puri et al. 2020). If vaccine hesitancy should prove to a major problem moving forward, then 2021 is likely to be a bumper year for cyber-epistemological research.

Acknowledgments

Paul Smart: This work is supported by the UK EPSRC as part of the PETRAS National Centre of Excellence for IoT Systems Cybersecurity under Grant Number EP/S035362/1.
Rob Clowes: This work is supported by research funding contract DL 57/2016/CP1453/CT0021). The authors would like to thank Ms. Anaïs Ansari for her assistance in the preparation of this manuscript.

References

- Andrada, Gloria. in press. "Mind the Notebook." *Synthese*, 1–20.
- Battaly, Heather. 2008. "Virtue Epistemology." *Philosophy Compass* 3 (4): 639–663.
- Bjerring, Jens Christian, and Nikolaj Jang Lee Linding Pedersen. 2014. "All the (Many, Many) Things We Know: Extended Knowledge." *Philosophical Issues* 24 (1): 24–38.
- Clark, Andy. 2015. "What 'Extended Me' Knows." *Synthese* 192 (11): 3757–3775.
- Clark, Andy. 2010. "Memento's Revenge: The Extended Mind, Extended." In *The Extended Mind* edited by Richard Menary, 43–66. Cambridge, Massachusetts, USA: MIT Press.
- Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.
- Clowes, Robert W. forthcoming. "The Internet Extended Person: Exoself Or Doppelganger?" *LÍMITE Interdisciplinary Journal of Philosophy & Psychology*.
- Fallis, Don. in press. "The Epistemic Threat of Deepfakes." *Philosophy & Technology*.
- Government of the United Kingdom. 2019. *Online Harms White Paper*. London, UK: Her Majesty's Stationery Office.
- Halfaker, Aaron, and John Riedl. 2012. "Bots and Cyborgs: Wikipedia's Immune System." *Computer* 45 (3): 79–82.
- Hardin, Russell. 2002. *Trust and Trustworthiness*. New York, New York, USA: Russell Sage Foundation.
- Heersmink, Richard. 2018. "A Virtue Epistemology of the Internet: Search Engines, Intellectual Virtues, and Education." *Social Epistemology* 32 (1): 1–12.
- Heersmink, Richard. 2015. "Dimensions of Integration in Embedded and Extended Cognitive Systems." *Phenomenology and the Cognitive Sciences* 14 (3): 577–598.
- Heersmink, Richard, and John Sutton. 2020. "Cognition and the Web: Extended, Transactive, or Scaffolded?" *Erkenntnis* 85:139–164.

- Jaruzelski, Barry, Robert Chwalik, and Brad Goehle. 2018. "What the Top Innovators Get Right." *strategy+business*, no. 93, 1–24.
- Lynch, Michael Patrick. 2016. *The Internet of Us: Knowing More and Understanding Less in the Age of Big Data*. New York, New York, USA: W. W. Norton.
- Nguyen, C Thi. 2020. "Echo Chambers and Epistemic Bubbles." *Episteme* 17 (2): 141–161.
- Palermos, Spyridon Orestis. 2018. "Epistemic Presentism." *Philosophical Psychology* 31 (3): 458–478.
- Pfeifer, Rolf, and Josh Bongard. 2007. *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, Massachusetts, USA: MIT Press.
- Pritchard, Duncan. 2018. "Extended Virtue Epistemology." *Inquiry* 61 (5–6): 632–647.
- Puri, Neha, Eric A Coomes, Hourmazd Haghbayan, and Keith Gunaratne. 2020. "Social Media and Vaccine Hesitancy: New Updates for the Era of COVID-19 and Globalized Infectious Diseases." *Human Vaccines & Immunotherapeutics* 16 (11): 2586–2593.
- Rajkumar, Raj, Dionisio de Niz, and Mark Klein, eds. 2017. *Cyber-Physical Systems*. Boston, Massachusetts, USA: Addison-Wesley.
- Rupert, Robert D. 2004. "Challenges to the Hypothesis of Extended Cognition." *Journal of Philosophy* 101 (8): 389–428.
- Schwengerer, Lukas. 2020. "Online Intellectual Virtues and the Extended Mind." *Social Epistemology* 1– 11. doi: 10.1080/02691728.2020.1815095.
- Simpson, Thomas W. 2012. "Evaluating Google as an Epistemic Tool." *Metaphilosophy* 43 (4): 426–445.
- Smart, Paul R. 2018. "Emerging Digital Technologies: Implications for Extended Conceptions of Cognition and Knowledge." In *Extended Epistemology* edited by Adam J Carter, Andy Clark, Jesper Kallestrup, Orestis Spyridon Palermos, and Duncan Pritchard, 266–304. Oxford, UK: Oxford University Press.
- Smart, Paul R, Robert W. Clowes, and Richard Heersmink. 2017. "Minds Online: The Interface between Web Science, Cognitive Science and the Philosophy of Mind." *Foundations and Trends in Web Science* 6 (1–2): 1–232.
- Trippel, Timothy, Ofir Weisse, Wenyuan Xu, Peter Honeyman, and Kevin Fu. 2017. "WALNUT: Waging Doubt on the Integrity of MEMS Accelerometers with Acoustic Injection Attacks." In *2017 IEEE European Symposium on Security and Privacy*, 3–18. Paris, France: IEEE.
- Wheeler, Michael. 2019. "The Reappearing Tool: Transparency, Smart Technology, and the Extended Mind." *AI & Society* 34 (4): 857–866.