

FILOSOFI *ifilosofi* 哲学 *Filosofie* Philosophie  
Φιλοσοφία *Filosofia* *Filosofia* Общее соображение  
PHILOSOPHY *ifilosofi* *Filosofia*

November 3rd-5th

# EPISTEMOLOGY OF THE INTERNET WORKSHOP

The African Centre for Epistemology and  
Philosophy of Science  
Department of Philosophy  
University of Johannesburg



# A B S T R A C T S

November 3rd

---

## Metaepistemic Negotiation and Epistemically Toxic Online Behaviour:

**Chris Blake-Turner (University of Alabama at Birmingham)**

The internet is an integral part of modern epistemic environments. While it can bring great epistemic benefits, *epistemically toxic behavior* also pollutes the internet. Epistemically toxic behavior is a catchall for a variety of behaviors that individuals and groups perform, behavior that degrades an epistemic environment. Epistemically toxic behavior, then, makes it harder for those in effected epistemic communities to gain positive epistemic statuses, such as knowledge, and understanding. In this talk, I try to make progress in the fight against epistemically toxic online behavior. I do this in three steps.

First, I give a new diagnosis of some of our most entrenched epistemic disagreements. Rather than disagreeing about the truth of a first-order claim, participants are better understood as disagreeing about attributions of trustworthiness. For instance, a disagreement about what to believe (e.g., whether the election was rigged or not) is better framed as a disagreement about what sources are trustworthy (e.g., the far-right media or the mainstream media). Call this phenomenon *metaepistemic negotiation*: negotiation not about what to believe, but about *who to believe about what to believe*. (I draw inspiration from the literature on metalinguistic negotiation. See Plunkett and Sundell 2013.)

Second, I apply the concept of metaepistemic negotiation to epistemically toxic behavior. Drawing on work on echo chambers (Nguyen 2020), I argue that some epistemically toxic behavior is best made sense of less as an attempt to propagate first-order epistemic claims, and more as an attempt to establish and secure a certain distribution of trustworthiness in an epistemic environment. This helps us to understand some of the particular difficulties in challenging epistemically toxic behavior. Suppose "The election was rigged" is less about the election and more about who the speaker thinks is trustworthy (at least with respect to relevantly similar claims). Typically, the speaker's attributions of trustworthiness to some sources (e.g., the far-right news media) will be complemented by attributions of untrustworthiness to others (e.g., main stream news media). Thus, attempts to challenge "The election was rigged" by appealing to sources that the speaker is metaepistemically impugning will be difficult.

Third, I explore practical upshots of the preceding discussion. In particular, I highlight some pitfalls of a strategy that naturally suggests itself: directly challenging the implicit attributions of (un)trustworthiness. Combatting epistemically toxic online behavior will require more radical, environmental solutions that address skewed distributions of trustworthiness.

---

# A B S T R A C T S

November 3rd

---

## Cults, Conspiracies and Fantasies of Knowledge

**Daniel Munro (University of Toronto)**

There's often pleasure in possessing knowledge, especially in possessing secret knowledge to which outsiders don't have access. As with most sources of pleasure, there's also pleasure in merely fantasizing about possessing secret knowledge: it can be fun to imagine you were first to make some monumental scientific discovery, that you've uncovered a government conspiracy to cover up UFOs, or that you've been chosen to receive a prophesy directly from God. Such fantasies are typically harmless. However, I argue that the allure of fantasizing about secret knowledge helps to explain the dynamics of certain epistemically problematic online communities. Specifically, it can help to explain why people join online communities which appear from the outside to have lost touch with reality.

Recent public discourse is riddled with examples of internet subcultures trafficking in conspiracy theories and other forms of misinformation: groups such as QAnon, flat-earthers, radical anti-vaxxers, and "incels" often congregate on fringe online message boards. Many such groups seem to exhibit the epistemic structure of what Nguyen (2020) calls "echo chambers": they profess a commitment to beliefs which allow them to predict and explain away outside sources of counterevidence, beliefs which are thereby insulated from evidential undermining. In this sense, these groups resemble religious cults, who often profess a commitment to bizarre conspiracy theories (Dyrendal 2013; Ryutaro 2018).

It can appear that members of such groups have been brainwashed into accepting delusional beliefs. However, empirical research complicates this. It suggests that, for example, whether members genuinely believe the claims they seem to profess is less important than the pleasure and entertainment they get out of participating in a community that claims to have access to secret knowledge unavailable to outsiders (Dyrendal 2016; Rosenblum and Muirhead 2019; Mercier 2020). In light of data like this, I argue that members of such communities often don't genuinely believe in their community's ideologies but are instead absorbed in fantasies of possessing secret knowledge. I identify certain features of online environments that make them especially susceptible to fostering communities built around such fantasies. And I argue that this again makes them similar to insulated religious cults, where many of the same forces seem to be at work. This conclusion might seem to some extent to neutralize the epistemic threat of online communities like these. However, I furthermore argue that the dangers of getting too caught up in group fantasies of knowledge is that they can become mistaken for reality, with mere fantasies of knowledge becoming illusions of knowledge. This explains why only a select few members of such communities evince a much deeper commitment to their group's ideology, sometimes with violent and tragic results. This again instantiates a broader kind of social phenomenon also seen in religious cults, which adds to the empirical plausibility of my account.

---

# A B S T R A C T S

November 3rd

---

## **Towards an Epistemic Compass for the Internet Abraham Tobi (University of Johannesburg)**

The internet is a source of information, widely available to all as consumers of or contributors to this bank of information. The ease of access to the internet means that the information we get from the internet does not always go through the same rigorous verification and validation process that other sources of information – like an academic journal, for instance – go through. This is a growing problem, leading to the spread of misinformation and disinformation (Miller & Record 2013, Worden 2019). A recognition of the pernicious effects of the spread of ignorance through the means of the internet has prompted major social media outlets like Facebook, Twitter, and Instagram to employ a system of independent fact-checking that flags or blocks information shared on their platforms that are factually incorrect or violates community standards, amongst other considerations. This move is a step in the right direction. Still, it is not all-encompassing as the epistemic downside is that epistemic authority is now vested on companies with a clear interest in manipulating information (for advertising purposes, for instance (Trombley & Flanagan 2001)).

The problem of misinformation and disinformation through the internet remains, and we are saddled with a solution that has its epistemic downsides. Hence, what sort of online epistemic standards are required to attend to these issues? I argue in this talk that a plausible solution to these issues should, at least, prioritize two factors.

The first is a recognition of the democratic nature of information shared on the internet. Damien Smith Pfister (2011: 220) describes the ease of access to the internet as having the potential for the 'democratization of knowledge'. The democratization of knowledge here can simply be understood as the process of 'bringing a wider range of people into the exchange of ideas, or as introducing new processes of information dissemination' (Mößner and Kitcher 2016: 1).

The second, which builds on the first, is the need to minimize the epistemic downside of placing epistemic authority on companies with non-epistemic motives. Taking the democratic nature of knowledge on the internet into account, the verification and validation process of this knowledge must involve the agents that are both responsible for it and its consumers.

---

# ABSTRACTS

November 3rd

---

Resistance strategies. The (Possible) Bright Side of Online Epistemic Bubbles.

**Daniel Barbarrusa and Lola M. Vizuete (University of Seville)**

Recent discussions in the epistemology of the Internet have shown serious concern about the isolation that the Internet can produce over communities through the creation of epistemic bubbles and echo chambers. Bubbles have been considered as especially problematic when they result from personalisation technologies, ubiquitous along the Web (Pariser, 2011). Moreover, epistemic bubbles and echo chambers have been widely regarded as harmful for both their members and the public sphere as a whole, as they "lead their members astray" and reinforce ideological separation (Nguyen 2020, 141).

We aim to question this generalisation. Under proper conditions, online epistemic bubbles can be useful for unfairly disadvantaged communities to resist epistemic injustices, constituting a form of epistemic activism (Medina, 2021). The crucial condition would be that their distinctive selective exposure (Nguyen, 2020) remains under the agent's control. This means that although temporarily shielded from voices—even relevant ones—beyond the community, members are still aware that these voices do exist and retain enough agency to withdraw from the bubble when needed. This way, the bubble could be beneficial for the community itself, "by augmenting and protecting their epistemic group agency" (Medina, 2021, 185), but also for the wider society, as it allows for silenced voices and discredited knowledge to flow through society.

Two considerations must be observed for our purposes.

First, we draw from Nguyen's (2020) demarcation on epistemic bubbles and echo chambers. While the former lack exposure to outsiders' voices and arguments, this is not the case in echo chambers, whose distinctive feature is the massive discredit towards outsiders. In our proposal, we will deal only with epistemic bubbles (which Nguyen takes as less resilient than echo chambers), leaving echo chambers as more problematic and out of our scope.

Second, we underline the difference between bubbles generated on purpose by users, and those invisible (Pariser 2011) bubbles resulting from personalisation technologies. The former are arguably easier to control, and thus, to be subordinated to epistemic activism objectives; while the latter will probably remain in the domain of secret algorithms—which could even absorb users' agency. Since these filter bubbles represent a greater challenge to the community's agency, we will focus on epistemic bubbles created on purpose.

How can they be helpful? As stated above, the key is to retain and develop the community's agency. In this sense, epistemic bubbles can comply with the features that Medina (2021) attributes to epistemic activism. Indeed, they can (i) augment the community's epistemic agency, (ii) amplify their voices, and (iii) facilitate the development and exercise of their epistemic capacities.

While acknowledging the risk of data surveillance and filter bubbles, we emphasize the Internet still offers unique chances for marginalised communities to fulfill these objectives. For example, easily controlling access to the community can boost the group's agency. Moreover, the Internet allows communicating the community's experiences and requests to larger auditoriums. Finally, asynchronous communication can promote a reflective discussion in which participants share in detail their own experiences, thus developing and exercising further their epistemic capacities in justice.

---

# A B S T R A C T S

November 3rd

---

How Twitter Gamifies Communication.  
**Thi Nguyen (University of Utah)**

Twitter makes conversation into something like a game. It scores our communication, giving us vivid and quantified feedback, via Likes, Retweets, and Follower counts. But this gamification doesn't just increase our motivation to communicate; it changes the very nature of the activity. Games are more satisfying than ordinary life precisely because game-goals are simpler, cleaner, and easier to apply. Twitter is thrilling precisely because its goals have been artificially clarified and narrowed. When we buy into Twitter's gamification, then our values shift from the complex and pluralistic values of communication, to the narrower quest for popularity and virality. Twitter's gamification bears some resemblance with the phenomena of echo chambers and moral outrage porn. In all these phenomena, we are instrumentalizing our ends for hedonistic reasons. We have shifted our aims in an activity, not because the new aims are more valuable, but in exchange for extra pleasure.

---

# ABSTRACTS

November 4th

---

## Anonymity and Asynchronicity as Key Design Dimensions for the Reciprocity of Online Democratic Deliberation

**Daniel Barbarrusa and Lola M. Vizquete (University of La Frontera and University of Chile)**

The online world and its impact for democracy were early seen by some in a positive light (e.g. Rheingold, 1993; Price 2009), while others saw it in a negative light (e.g. Sunstein, 2001; Curran et al., 2012). The optimists thought that this new world would strengthen our democracies and solve many of its offline problems. For example, deliberative democracy is sometimes thought to require a State-wide democratic deliberation (cf. Tanasoca, 2020; see also Young, 2000: 44-45). However, a State-wide, face-to-face deliberation is rightly thought to be unfeasible (cf. Ackerman and Fishkin, 2004). Given this, with the advent of the online world, some thought that online deliberation could help solve this alleged scaling-up problem facing deliberative democracy, given that it does not require co-presence and co-temporality.

Whether online deliberation can help with such an issue (cf. Ito et al., 2017; Klein, 2015), we think it is worth exploring other possibilities of the online world for democracy. In this talk we are interested in examining the sort of design features that it should have to promote democratic deliberation. In particular, we consider one crucial aspect of deliberative argumentation: namely, its reciprocity, which puts interaction centre stage to capture the back-and-forth of reasons. More precisely, we focus on two essential features of the deliberative interaction: namely, its listening widely and listening carefully. Given that, we study different combinations of two of the most innovative variables that the online world brings to the design of online deliberation, namely synchronicity/asynchronicity and identification/anonymity. We conclude that one sort of online deliberation that combines the two design features of anonymity and asynchronicity is likely to better promote the reciprocity required for democratic deliberation than both natural and designed offline deliberations (such as the designed deliberation in Deliberative Polling) and online simulations of them.

The talk proceeds as follows. Firstly, we introduce the deliberative model of democracy and the general notion of democratic deliberation that it often embraces. In particular, we introduce the reciprocity that is uncontroversially accepted as an essential feature of democratic deliberation. Secondly, we clarify the notion of online deliberation and present some design features, focusing on asynchronicity and anonymity. Finally, we assess four different alternative designs given the two mentioned variables and offer some brief concluding remarks.

---



# A B S T R A C T S

November 4th

---

## Epistemic Injustice in Data Mining **Damon Mackett (University of Johannesburg)**

The widespread use of the internet has led to an increase in the collection of personal information or data. The collection of personal data is facilitated by agreements set out by publishers. How the general public (general public, in this case, referring to those who are not aware of data mining or the techniques used) come to understand End-User Licence Agreement (EULA), Click-Through Agreement (CTA) or any variation of the above and what these agreements represent are fundamentally different (if any understanding takes place, to begin with). I argue that this misunderstanding comes from how information in these agreements is conveyed, making this primarily epistemic in nature. The information presented is accurate but it is not conveyed in a sincere way. The lack of effort in conveying information in these agreements sincerely in addition to Dotson's theory of contributive injustice results in epistemic hoarding of specialized hermeneutical resources.

I argue that the above phenomenon being a new distinct form of epistemic injustice. A form of epistemic injustice that results in economic harm in economic exploitation. The argument rests on the idea that the amount of data collected is valuable, in the sense that companies who employ data mining techniques can profit from them. The users are unaware of data mining techniques, and therefore do not know how valuable the amount of data might be. They do not know the value the amount of data holds because there is little effort put in conveying that in the agreement.

I do not argue that publishers need to make the public aware of how much economic worth the data holds. Rather I argue that users do not know how much data is collected and harvested due to the lack of sincerity. Only those aware of data mining techniques can begin to understand the amount of data one can collect and harvest for knowledge. Only if a person who has an interest in or has been trained in data mining would be able to understand how data can be valuable due to specialized hermeneutical resources.

This can be seen in a specialized language that a field develops that removes its discourse from the public discourse. Doctors and lawyers are easy examples of specialized hermeneutical resource hoarding. I argue that data mining employs the same technique. This technique allows for the creation of exploitation, users do not know that the value of their data because they are not aware of how much data is collected, as well as its continuation. It is difficult to identify something you are unaware even exists in the first place. The general public may have a more accurate picture of the economic worth their value holds if they first understand how much data is collected in the first place. They would have a better chance of understanding this if the information in the agreements were conveyed more sincerely.

---



# ABSTRACTS

November 4th

---

## Epistemic Corruption in Online Environments Daniella Meehan (University of Glasgow)

The general meaning of corruption is understood as articulating damage or forms of degeneration, and it is usually confined to both moral and political philosophy (Miller 2018, Ceva & Ferretti 2017, 2021). However, despite the array of literature on this topic, there has been significantly less focus on epistemic corruption and the harms that it can cause to individuals and their wider epistemic communities. I argue that one such example of these corruptive harms is those that are inflicted on our intellectual character, which in turn gives way to a variety of epistemic vices.

The plan for this paper is as follows. Firstly, I will introduce the notion of epistemic corruption generally, before focusing on how our online environments can be epistemically corrupting through the increasing presence of information disorder - understood as the numerous ways that online information can distort the truth and our trust in the media (Wardle 2019). Crucially, information disorder can provide optimal conditions for the flourishing of civic and epistemic vices and the suppression of corresponding virtues. In particular, I focus on three distinct epistemic vices - prejudice, conspiratorial thinking and epistemic capitulation. I outline how these vices present themselves in these corruptive online environments in accordance with the 'five modes of epistemic corruption' - understood as the various (non-exhaustive) ways that a corrupting system can install epistemic vices in its corruptees (Kidd 2021).

With regards to the former vice, I argue that prejudice can be intensified in corruptive online environments with reference to cases of false stereotyping and testimonial injustice (Fricker, 2017). Next, I argue that corruptive online environments can increase the scope of the vice of conspiracy thinking (Cassam, 2020) through the presence of misinformation. Finally, I argue that corruptive online environments can enable the acquisition of epistemic capitulation (Battaly, 2017) through the overwhelming excess of information that these environments create.

I then conclude by examining the various ways systems and environments can overcome their corruptive state, by assessing both individualistic and structural ameliorative solutions. These solutions concern educating for intellectual virtue, and the wider approach of fostering civic epistemic virtues in ourselves and wider communities (Baehr (2011, 2015, Pritchard 2013, 2014, Sullivan and Alfano 2021).

---

# A B S T R A C T S

November 4th

---

## Computer Software Agency and the Bubble Effect: A Case for Epistemic Injustice

**Peter Ikhane (University of Ibadan)**

In this paper, I argue that personalized search resulting from how internet search engines tailor search results according to derived user preferences based on irreproducible criteria incarnates epistemic injustice. I do this by showing that the exclusion of certain data in search engine results on the basis of user search history and location amounts to epistemic injustice.

The growth of Artificial intelligence technology has seen to the possibility of computer softwares, such as internet search engines as Google, being able, like humans, to explore complex issues, resulting in their providing machine-generated explanation, as well as giving large sizes of information on request (Powers, 2017). Although internet search engines as software agents seem immensely useful, they produce multiple instances of biases. The nature of these biases has been examined in the literature (Curkovic & Kosec, 2018). From an epistemic perspective, while there have been discussions that focused on the nature of justification of beliefs formed from data from the internet (Miller and Record, 2013), a virtue-based epistemological consideration of how to use the internet in seeking to improve our information-seeking behaviour (Heersmink 2013), the question of whether personalized search with regards to the use of internet search engines is a boon or burden (Smart and Shadbolt, 2019), the question of epistemic trust with respect to search engines (Frost-Arnold, 2019; Gunn and Lynch, 2019), there is little attention to the nexus of personalized search results and epistemic injustice.

Eli Pariser, in *The Filter Bubble: What the Internet is Hiding from You*, 2011, introduced the term "filter bubble" as a personal selection bias inherent to internet search engines. He argued that internet search engines use personalized algorithms in order to evaluate and stratify results by adjusting to users' beliefs and attitudes, leading to the creation of personalized search bubble. In the light of this, I argue in the paper that the bubble effect generates an instance of epistemic injustice, where internet search engine algorithm systematically excludes certain information on the basis of personalized search results. This is such that if we choose to encompass some breath of knowledge through internet search engines, there is the good chance that the search bubble effect will lead us to results that have already been chosen for us.

---

# A B S T R A C T S

November 4th

---

Ways of Worldfaking

**Boaz Miller (Zefat Academic College)**

Deepfakes, namely, algorithmically created realistic images and videos that make it appear as if people did something they didn't, undermine our fundamental epistemic standards and practices. Yet the nature of the epistemic threat they pose remains elusive. After all, fictional or distorted representations of reality are as old as cinema. Existing accounts of technology as extending the senses (Humphreys 2004), mediating between subjects and the world (Verbeek 2011), or translating between actants (Latour 2005) cannot characterize this threat. Existing concrete accounts of the threat of deepfakes by social epistemologists such as Regina Rini (2020) and Don Fallis (2020) fall short of their target.

Employing the notions of artifact affordance and technological possibility (Record 2013; Davis 2020), I argue that the epistemic threat of deepfakes (and CGI more generally) is that for the first time they afford ordinary computer users the practicable possibility to fairly cheaply and effortlessly make fictional worlds indistinguishable from the real world. Normatively, a deepfake is epistemically malignant when (1) a reasonable person is misled to believe that the fictional world is the actual world; (2) she forms beliefs about the actual world about issues that are morally or epistemically important. For example, a satirical deepfake of Queen Elizabeth dancing to hip-hop song is benign because a reasonable person understands this is fiction. But a deepfake of a misogynic speech by Obama is malignant because it misleads a reasonable person about Obama's views of women. I illustrate how this analysis generalizes to other case studies, such as a Photoshop makeover, or a QAnon discussion group.

---

# ABSTRACTS

November 5th

---

## Starting From Vice: Virtue Responsibilism Online Lillianne John and Urna Chakrabarty (Duquesne University and University of Delhi)

Philosophising about the internet is a tricky prospect, given how quickly the complex industry of the web advances every few years. We focus on how computer-mediated-communication (CMC) via social networking services (SNS) provides scant access to that variety of personal identifiers that face-to-face interaction yields (Parcell 2008). Individual difference is displaced online by the exaggeration of social identities. As a consequence, members of SNS collectives commonly perceive themselves (and, we add, those from outgroups) as representatives of a group, rather than as individuals. Crucially, those members who post information are taken to be "idealised community members", in terms of their commitment to community principles. This is the Social Identity Model of Deindividuation Effects (SIDE) (Postmes et al. 1998); it has informed much of the discourse around the dangers of group polarisation in the digital age.

We marshal the SIDE model to argue that there is a need to formulate strategies of communication—especially between or among different communities online—that proceed from the assumption that we operate in these fragmented (as 'smeared selves' Ess 2010), epistemically vicious (Heersmink 2018) ways on SNS: we are not self-contained, atomic nodes that operate uniformly or perfectly when online. Extending a version of the principle of charity to one's perceived SNS opponent (a "believe the best of the other person" approach), though ethically desirable, appears to be unworkable in practice, given (1) the inherent limitations of mass interaction online, and (2) the propensity of SNS algorithms to encourage aggressive interactions. Such an approach, covering the majority of 'average' SNS communications—that are neither 'troll behaviour' nor explicit proselytisation/propagandism—would lead to more fruitful interactive experiences online.

We will be drawing our insights from 'ship wars' (disagreements between the supporters of two different romantic character pairings regarding the relative superiority of their pairings, or ships) within and among Western fan collectives on Twitter and Tumblr that are clustered around different Chinese danmei (also called Boys' Love; refers to fictional media romanticizing male-male relationships, largely for and by straight women in contemporary China) novels. The ship wars we study involve fandom purity culture, which refers to a contemporary practice of critiquing fanworks on the basis of their being morally or politically "problematic", usually in harmfully reductive ways that mirror the tactics of cancel culture.

---

# A B S T R A C T S

November 5th

---

## Epistemic Justice and Empowerment for Autistic Women: The Role of the Internet

**Amandine Catala (Université du Québec à Montréal)**

Because diagnostic criteria for autism were originally formulated based on autistic boys, and because autistic traits typically manifest differently in girls due largely to gender norms, autistic women tend to be underdiagnosed relative to their male peers. This male bias in both the clinical and social understandings of autism generates hermeneutical injustice for undiagnosed autistic women, by obscuring from understanding a part of their experience “which it is strongly in [their] interest to understand” (Fricker 2007, 149).

In her analysis of hermeneutical injustice, Miranda Fricker (2007, 149) refers to the situation of women who had been experiencing postpartum depression without knowing that it was postpartum depression, and who, upon learning about postpartum depression, suddenly realized this was exactly what they had been experiencing. Fricker refers to this type of realization as a “hermeneutical breakthrough” (Fricker 2007, 149), and characterizes it as a “revelation,” “a life-changing flash of enlightenment” (Fricker 2007, 153).

In this paper, I look at the ways in which the internet can enhance epistemic agency by enabling greater epistemic justice and empowerment for autistic women. I proceed in two main steps. First, I argue that the internet can play an important role in helping previously undiagnosed autistic women reach a hermeneutical breakthrough, which I argue involves a hermeneutical tool, hermeneutical representativeness, and improved self-understanding. From blogs to webcomics to YouTube channels, the internet allows autistic women to widely share their stories and spread new, intersectional understandings of autism. In providing a channel for these various epistemic resources, the internet makes it possible for undiagnosed autistic women to recognize themselves in first-person testimonies and more accurate representations of autism stemming from autistic women. In this way, the internet can shed new and crucial hermeneutical light on so many aspects of the everyday experience and entire life of many autistic women who have made it to adulthood without a diagnosis and who can now understand themselves better and navigate the neurotypical world accordingly.

Second, I argue that certain characteristics of the internet make it conducive to epistemic empowerment for autistic women. In particular, the possibility of asynchronous, anonymous or pseudonymous, not face-to-face epistemic exchanges and communities through various online platforms alleviates many of the sensory, communicative, and social difficulties that often arise in “real-world” environments. In this sense, the internet can provide an optimal space for autistic women not only to acquire, share, and exchange knowledge regarding their experiences, but also to build their own community, support practices, and collective projects.

In arguing for these claims, the paper not only brings to the fore the philosophically neglected intersection of gender, neuronormativity, and neurodiversity, but it also contributes to a more complete understanding of hermeneutical breakthroughs, epistemic empowerment, as well as some of the features that can make online environments epistemically fruitful and conducive to greater epistemic justice and empowerment.

---

# A B S T R A C T S

November 5th

---

## The Knowledge Machine: Googling and Testimony Caitlin Rybko (University of Johannesburg)

Since its inception, Google has fundamentally altered how we parse and peruse information on the internet. With 5.6 billion searches per day, it has cemented itself into the daily lives of users. And in doing so, has become a vital epistemic resource. But we ought to recognise that it has also fundamentally changed the way we access and consume information.

If our technological infrastructures are embedded within our knowledge practices, then how we are able to acquire knowledge is dependent on the resources that are available to us. As our technology has evolved, so too have our standards for epistemic responsibility (Miller and Record 2017). As such, our engagement with these new epistemic practices should evolve as well. The internet presents us with a curious epistemic tangle because it operates differently from traditional sources of information and knowledge. And Google, as a distributor of this information has come to be seen as both a source of information and a knowledge provider. This paper then, is an effort to address and engage with what I take to be a new problem for epistemology that has emerged from Google's status as a source of information.

In 2012 Google launched what it terms a 'knowledge panel', and in 2013 announced that the goal of their service is to "know what you want, before you do" (Adams 2013). The knowledge panel is a fundamental tool for achieving this goal. It is a rectangular module that provides direct answers to an inquiry placed either at the top, or on the right of search results. In providing information directly rather than merely linking users to possible sources of information, Google has moved from being something akin to a directory that points users towards informants, to becoming an informant itself. But, Google differs substantially from traditional examples of testifiers in the epistemic literature and so an investigation is in order. This paper then, is split into two parts. In the first, I will argue that there are three prima facie reasons why we should take the information contained in knowledge panels to be testimony. And in the second I will show that not taking knowledge panels as testimony leads to negative epistemic consequences.

---

# A B S T R A C T S

November 5th

---

## Is a Retweet not an Endorsement? Understanding and Evaluating Novel Speech Acts on Social Media

**Glenn Anderau (University of Zürich)**

The advance of the internet and social media has led to changes in the way we communicate. Nevertheless the epistemic environment of social media remains largely unexamined. In my talk, I want to focus on analysing social media testimony. I will focus on two novel speech acts social media has made possible, namely the sharing and liking of content. It is unclear how these acts are meant to be understood. In particular,

I will examine whether the sharing or liking of content on social media constitutes an endorsement of said content. This question has only sporadically been discussed in philosophy. There is no clear consensus on its answer in everyday language either, which means that the meaning of these speech acts remains up in the air. The lack of clear norms regarding social media testimony can cause confusion. In my talk, I will defend the endorsement view of social media testimony which holds that sharing or liking content on social media should be seen as an endorsement by default unless the content is explicitly disavowed in the sharing process itself. This view should be understood as a constitutive norm: It does not mean that scenarios in which content is shared or liked on social media but not meant to be endorsed cannot exist.

The motivation behind supporting the view that the sharing and liking of content on social media should by default be seen as an endorsement is twofold: Firstly, an endorsement view aligns most closely to the spirit with which social media platforms allow us to share and like content. In particular, I will argue that the endorsement view is superior to the quotation view, which holds that the speech acts of liking and sharing content online are best understood as quotations. This is not to say that the liking and sharing of content on social media are not complex speech acts, which could be used to express a host of different sentiments. Rather, the argument is that *ceteris paribus*, the endorsement view holds more promise than the quotation view as a default position. This has to do with the way social media platforms are structured and the emphasis they place on the person sharing or liking the content while they perform these speech acts.

Secondly, I argue for the endorsement view as a norm because it is better at addressing the harmful epistemic effects of social media testimony than rival views. Social media has proven a fertile breeding ground for fake news and other forms of disinformation. An endorsement view places the responsibility for social media testimony at the feet of the agents sharing or liking content on social media rather than its recipients. This is a virtue of the endorsement view, at least if we are interested in using it to combat the spread of online disinformation.

---



# ABSTRACTS

November 5th

---

## Echo Chambers, Trust, and Epistemic Agency under Oppression Karen Frost-Arnold (Hobart and William Smith Colleges)

I argue that much of the discourse about echo chambers and filter bubbles ignores the agency of marginalized people, and I suggest some ways that epistemologists can center oppressed people in their analysis. I begin by diagnosing a common pattern in internet epistemology. Step 1: concern grows about bad actors doing epistemic harm. Step 2: a structural feature of the internet is identified that is being exploited by the bad actors to do epistemic harm. Step 3: arguments are made that the structural feature is an epistemic threat and/or should be eliminated. I show how this pattern appears in a recent account of the harms of online personalization (Gunn and Lynch 2021). I argue that this account ignores how marginalized people use the internet to protect their epistemic agency. In hostile epistemic environments, oppressed people often need to protect their epistemic agency by shielding their time, maintaining their self-trust, and sustaining their connection to epistemic communities. By filtering out hostile voices and cultivating distrust in those who disrespect them as knowers, marginalized people use online personalization for epistemic good. I conclude by arguing that adopting standpoint methodology and centering the concept of power in internet epistemology will provide more fruitful accounts of the promise and pitfalls of online personalization and other features of the internet.

---