



**University of  
Zurich**<sup>UZH</sup>

**Zurich Open Repository and  
Archive**

University of Zurich  
University Library  
Strickhofstrasse 39  
CH-8057 Zurich  
[www.zora.uzh.ch](http://www.zora.uzh.ch)

---

Year: 2023

---

## **Fake news and epistemic flooding**

Anderau, Glenn

**Abstract:** The advance of the internet and social media has had a drastic impact on our epistemic environment. This paper will focus on two different risks epistemic agents face online: being exposed to fake news and epistemic flooding. While the first risk is widely known and has been extensively discussed in the philosophical literature, the notion of ‘epistemic flooding’ is a novel concept introduced in this paper. Epistemic flooding occurs when epistemic agents find themselves in epistemic environments in which they are routinely confronted with more information and evidence than they can diligently process. Epistemic flooding is one of the most significant risks epistemic agents run while using social media and one of the reasons why the first risk (being exposed to fake news) is especially pernicious. It is particularly harmful if one ascribes to the Spinozan model of belief acquisition in which belief comes easy. Mitigating the combined threat of fake news and epistemic flooding requires us to think normatively about our epistemic environments and how to better them. This paper turns to the normative framework of epistemic environmentalism in order to accomplish this.

DOI: <https://doi.org/10.1007/s11229-023-04336-7>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-237567>

Journal Article

Published Version



The following work is licensed under a Creative Commons: Attribution 4.0 International (CC BY 4.0) License.

Originally published at:

Anderau, Glenn (2023). Fake news and epistemic flooding. *Synthese*, 202(4):106.

DOI: <https://doi.org/10.1007/s11229-023-04336-7>



# Fake news and epistemic flooding

Glenn Anderau<sup>1</sup> 

Received: 14 July 2022 / Accepted: 5 September 2023  
© The Author(s) 2023

## Abstract

The advance of the internet and social media has had a drastic impact on our epistemic environment. This paper will focus on two different risks epistemic agents face online: being exposed to fake news and epistemic flooding. While the first risk is widely known and has been extensively discussed in the philosophical literature, the notion of ‘epistemic flooding’ is a novel concept introduced in this paper. *Epistemic flooding* occurs when epistemic agents find themselves in epistemic environments in which they are routinely confronted with more information and evidence than they can diligently process. Epistemic flooding is one of the most significant risks epistemic agents run while using social media and one of the reasons why the first risk (being exposed to fake news) is especially pernicious. It is particularly harmful if one ascribes to the Spinozan model of belief acquisition in which belief comes easy. Mitigating the combined threat of fake news and epistemic flooding requires us to think normatively about our epistemic environments and how to better them. This paper turns to the normative framework of *epistemic environmentalism* in order to accomplish this.

**Keywords** Social epistemology · Fake news · Epistemic environmentalism · Spinozan belief acquisition · Epistemology of the internet

## 1 Introduction

Being an epistemic agent online can be a daunting prospect. Even though the internet and social media hold a lot of potential as easily accessible sources of knowledge, one is also exposed to a series of risks such as the heightened threat of being exposed to disinformation.<sup>1</sup> While fake news is not new, it has become much more widespread online, especially because the technology of the internet and social media allows for

---

<sup>1</sup> The epistemic risks one can face online of course transcend merely being exposed to disinformation. For example, one could also be in greater danger of entering epistemic bubbles or echo chambers (Nguyen, 2020).

---

✉ Glenn Anderau  
glenn.anderau@philos.uzh.ch

<sup>1</sup> Philosophy Department, University of Zürich, Zurich, Switzerland

much easier dissemination of it (Jaster & Lanus, 2019; McIntyre, 2018).<sup>2</sup> Fake news poses a risk because it can lead to the adoption of false beliefs and it is this threat which most analyses of fake news are primarily concerned with. But the impact of fake news runs deeper than that: As Blake-Turner (2020) highlights, fake news can also cause serious problems by degrading our epistemic environments.<sup>3</sup> Expanding our understanding of the threat posed by fake news is important in order to fully grasp its impact.

This paper will address an additional threat epistemic agents face online, *epistemic flooding*. Epistemic flooding is meant to describe epistemic environments in which epistemic agents are continuously inundated with more information than they can diligently process. Fake news is an especially potent threat when combined with epistemic flooding. In order to fully grasp the impact of fake news online, we need to consider the dual threat of fake news and epistemic flooding.

This paper will proceed in the following way: First, it will outline the concept of epistemic flooding by providing a definition and considering an example. Next, the focus will be on the connection between fake news and epistemic flooding, and the combined threat they pose. I will also argue that both phenomena are increasingly widespread online. Given the newfound ubiquity of epistemic flooding, tackling its combined threat with fake news is a tricky task. Finding solutions to it should be seen as an occasion for epistemologists to think about the nature of epistemic normativity. Finally, this paper will consider the notion of epistemic environmentalism (Ryan, 2018, 2021) in order to stymie the combined threat of fake news and epistemic flooding.

## 2 Defining epistemic flooding

In order to provide a functioning definition of epistemic flooding, I first need to explain how I understand the terms ‘epistemic environment’ and ‘epistemic agent’, as both are central to my definition. I use ‘epistemic environment’ as defined by Shane Ryan, who posits that epistemic environments are “constituted by facts, whether social or physical, that bear on epistemic attainment (gaining knowledge, understanding, and so on) in that environment” (2021, p. 2). Epistemic environments by this definition can come in many forms: a classroom can constitute an epistemic environment but so could an airport or social media. The latter is especially important, given that I will argue that the technology of the internet and social media has made epistemic flooding commonplace.<sup>4</sup>

I use ‘epistemic agent’ rather than doxastic agent to include actions agents take which can bear on epistemic attainment but which are not directly tied to forming

---

<sup>2</sup> I adopt a specific definition of fake news on page 14 of this paper.

<sup>3</sup> There are of course also a host of political concerns brought about by fake news. I mainly focus on the epistemic perspective in this paper, however. A concern which could be both politically and epistemically worrisome is the erosion of trust in genuine news sources which the consumption of fake news could elicit.

<sup>4</sup> I will not argue that epistemic flooding did not exist in the pre-internet age. I merely argue that it has become more prevalent through the internet and social media. Given how ubiquitous the internet and social media usage have become, this means that epistemic flooding is now a phenomenon epistemologists can no longer ignore (at the very least, if they have some interest in speaking about non-ideal epistemology).

beliefs. Importantly, actions which regulate one's epistemic environment would fall under the realm of epistemic agency even if said actions do not directly lead to the formation of a belief or to the acquisition of knowledge. We can call this a broadscale view of epistemic agency and it will be of essence to me in the latter part of this paper. The actual definition of epistemic flooding itself would not be affected by a more narrow definition of 'epistemic agent'.

## 2.1 Definition

*Epistemic Flooding:* Epistemic flooding occurs when epistemic agents find themselves in epistemic environments in which they are routinely confronted with more information and evidence than they can diligently process.

Several further clarifications are important to add here: It is only correct to speak of epistemic flooding if being exposed to a deluge of information we are unable to adequately process is a recurring problem we face over a longer period of time in the epistemic environment we inhabit. Single instances of 'flooding' pertaining to singular, self-contained situations and inquiries are not enough to qualify as epistemic flooding in this sense.<sup>5</sup> So an isolated situation in which we are flooded with more information than we can diligently process for a limited amount of time does not qualify as 'epistemic flooding' in this sense. We could still call such situations instances of epistemic flooding in theory, but I am only interested in scenarios in which we inhabit an epistemic environment which is habitually flooded over a longer period of time in this paper.

Another is that I speak of epistemic flooding in the context of the epistemic environment we inhabit. It is epistemic environments which are flooded which then impacts the agents inhabiting them. This is important because it means the information we are flooded with impacts our epistemic surroundings in general and affects our everyday epistemic practices. It also means that the effects of epistemic flooding do not solely pertain to individuals but epistemic communities as a whole.

There are two reasons why epistemic flooding can be considered worrisome: The first is that epistemic flooding bears negatively on our epistemic attainment.<sup>6</sup> This worry is multifaceted. The most simple worry is that epistemic flooding leads us to adopt false beliefs: If we are less diligent in processing new information, it stands to reason that we are more likely to accept false beliefs. However, it is debatable whether epistemic flooding really makes us more susceptible to adopting false beliefs. The worry becomes more poignant if we adopt a Spinozan model of belief acquisition,

---

<sup>5</sup> Being a neologism, one could of course conceive of 'epistemic flooding' in this way. However, I believe that singular cases of 'epistemic flooding' are a smaller epistemic risk, one which epistemic agents need to learn to live with. Singular instances of 'epistemic flooding' could also be equivocated with the psychological term *information overload* (Bermes, 2021). It is only once epistemic flooding takes place over a prolonged period of time that the concept becomes a more significant issue and one deserving of more theoretical attention by epistemologists.

<sup>6</sup> I use 'epistemic attainment' here in order to remain neutral about different goals epistemic agents might pursue (such as knowledge, truth, forming true beliefs, acting in accordance to epistemic virtues etc.). Bearing negatively on epistemic attainment is meant to be understood in a broad sense here: It is enough for epistemic flooding to indirectly impact our epistemic goals negatively for the label to apply.

in which belief comes easy. I will discuss this in more detail in Sect. 3 of the paper. However, we need not adopt the Spinozan model or even believe that epistemic flooding directly leads to the adoption of false beliefs in order to accept that epistemic flooding bears negatively on epistemic attainment. Epistemic flooding could be harmful if it leads us to accept misinformation in the form of ‘stored information’ (even if this does not amount to the formation of a belief). We could also be flooded with information which is useless to us as epistemic agents (even if it is not wrong) and therefore prevents us from engaging with pertinent information and from forming true beliefs about issues we care about. Finally, epistemic flooding might lead us to adopt epistemically harmful sub-doxastic states such as aliefs which bear negatively on our epistemic attainment. I discuss how epistemic flooding can lead to the adoption of beliefs in the Spinozan sense or aliefs more extensively in Sect. 3 of this paper.

The second risk is that epistemic flooding limits our epistemic agency. Since epistemic flooding impedes our ability to diligently process new information, it impacts our ability to act freely as epistemic agents. It is important to note that we mostly do not diligently screen or fact-check new information we receive even in non-epistemically flooded environments. The difference however, is that in non-epistemically flooded environments we are not impacted in our ability to do so. If a claim we encounter strikes us as odd, we are completely free to inquire further into the matter. In epistemically flooded environments our ability to do so is denied by the fact that we have too much information to process. It is important to note that epistemically flooded environments do not fully prevent us from acting: they merely reduce the amount of freedom we have over epistemic activities such as fact-checking etc. We are not completely unable to perform these actions in epistemically flooded environments.

Nevertheless, I consider the restriction of our epistemic agency a considerable worry, one which should not be neglected. This is especially true if we think of our epistemic agency as a whole rather than of isolated cases of inquiry. The loss of our epistemic agency remains even if we were to find ourselves in an epistemically flooded environment in which we encounter only correct information. Such a scenario, though relatively unlikely if we keep in mind that I am only considering cases in which epistemic flooding takes place over a prolonged period of time, would at least partially rule out the first worry of bearing negatively on epistemic attainment. Forming false beliefs for instance, would not be an issue we worry about in such a scenario. Of course, even an epistemic environment in which we encounter only correct information could still have a negative impact on our epistemic attainment in other ways. For instance, seeing certain beliefs overrepresented in the information I encounter could lead me to form false higher-order evidence about how many people hold this belief. This could be true even if the initial belief is correct. So epistemic flooding could be worrisome even in situations in which we encounter only accurate information. However, we could imagine a scenario in which we are flooded not only with accurate information, but also accurate information which does not have any negative impact on our epistemic attainment at all. While such a scenario is even more unlikely, it is at least theoretically possible to imagine epistemically flooded epistemic environments which completely rule out the first problem of bearing negatively on our epistemic attainment.

Nevertheless, the manner in which we form beliefs and in which we take decisions as epistemic agents (such as diligently screening and fact-checking a new piece

of information) would be impacted even if we are in such a scenario. What I will grant is that there could be both practical and epistemic reasons which would lead us to prefer yielding some of our epistemic agency in order to be in an epistemically flooded environment which contains correct information which is beneficial to us. But such scenarios are hypothetical and rare and even in these scenarios the loss of epistemic agency we always experience in epistemically flooded environments needs to be outweighed by some practical or epistemic reasons.

One could of course be skeptical about the notion that we lose epistemic agency in epistemically flooded environments. The argument could go something like this: Even though I am not able to diligently process every piece of new information, I still have the same amount of diligence regarding any given *single* piece of information. Therefore, my epistemic agency is not impacted even in epistemically flooded environments. We could compare this to entering a cake shop in which our practical agency is not impacted by the fact that we cannot consume every single cake in the shop (but presumably, we can still pick the cake we want).<sup>7</sup>

However, I believe that epistemically flooded environments and the cake shop scenario are disanalogous in important ways. Even if we ignore the differences between epistemic and practical agency, the case of epistemic flooding has a bigger impact on us since being exposed to information passively has an impact on our epistemic attainment. I argue for this because being exposed to information without carefully engaging with it can lead us to adopt false beliefs (if we buy into a Spinozan Model in which belief comes easy) or at least leave us with sub-doxastic states such as aliefs which can indirectly affect our belief formation process. Our epistemic attainment is impacted by the passive exposure to information and so the fact that we cannot engage with all of it does leave us more vulnerable. This differs from the cake shop analogy: We do not partially ingest a cake simply by looking at it and being exposed to many cakes in a room has no effect on our caloric intake or how we cover our nutritional needs. I explain why I believe that being exposed to information in epistemically flooded environments does leave us with beliefs (in a Spinozan sense) or sub-doxastic states (such as aliefs) in more detail in Sect. 3.

Another worry might be that epistemic flooding in its current definition is too ubiquitous to truly be a worry and that our epistemic environments are flooded even in the most mundane situations. For example, if I am in a café and there are several conversations around me, I might already be in an epistemically flooded environment since I cannot possibly follow all of them.<sup>8</sup> Even the amount of sensory input which is available to us might be considered a form of epistemic flooding. But since such situations are extremely commonplace, we should not consider them instances of epistemic flooding because in that case epistemic flooding is not worrisome, seeing that it includes mundane examples of this kind.

I agree that trivial situations such as being in a coffeeshop and overhearing several conversations should not count as instances of epistemic flooding, at least not in the way I am describing the term in this paper. Being in a crowded coffeeshop is not an instance of epistemic flooding since in this scenario, the flooding occurs in an

<sup>7</sup> I thank an anonymous reviewer for raising this worry and this example.

<sup>8</sup> I thank an anonymous reviewer for raising this worry and providing this example.

isolated instance and not over a longer period of time. However, there is another aspect of epistemic flooding which sets it apart. My definition argues that in epistemically flooded environments we are ‘confronted’ with more evidence than we can diligently process. It is this term which does some of the heavy lifting here to explain the difference between genuine cases of epistemic flooding in my understanding and more mundane examples such as overhearing conversations in the coffeeshop. Being confronted with information is having this information made accessible to us in a very tangible way, it is not simply the totality of possible evidence in our surroundings. In the coffeeshop example, there is a difference between several conversations taking place in the background (in which we are not directly confronted with this information) and several conversations being so loud we have to actively try not hear them or being so close to us that they draw our attention. The latter would count as being confronted, although it still seems like an isolated instance of flooding. In the unlikely event that we are confronted by new information in such a way in a coffeeshop over a prolonged period of time, I would grant that this counts as an instance of epistemic flooding. However, I would not grant that such examples are common occurrences.

The question of what precisely counts as epistemic flooding according to my definition is a good one and one which I will not be able to answer in more detail here. This is because I want to focus on a particular kind of epistemic flooding in the rest of this paper, namely Online Epistemic Flooding (OEF) which takes place on social media. While I do not claim that social media is the only epistemic environment which is permanently flooded, I believe that it constitutes both a very clear example of epistemic flooding as well as having some unique characteristics which make it worth thinking about OEF as its own subcategory. I maintain that epistemic flooding occurs ‘offline’ as well and that such instances of epistemic flooding pose a worry to epistemic agents. This can be true even in epistemic environments which are not necessarily flooded with epistemically harmful content. However, I believe that it is worth exploring this point in more detail in future work. I will focus in this paper exclusively on OEF, which I believe should also be more resistant to skepticism about whether such environments are truly flooded. I discuss the specific characteristics of OEF in more detail in Sect. 2.3 of this paper.

## 2.2 Example of epistemic flooding

Epistemic flooding has become more commonplace in the epistemic environment of social media. In fact, facing it has become the *default* state for most epistemic agents online. Consider the following example:

*Leila.* Imagine Leila is an epistemic agent active on three social media platforms. Her everyday routine includes scrolling through her social media feeds for half an hour each when she wakes up and before she goes to bed. While scrolling through her feeds she encounters content and information from her social media friends and accounts she follows. Often she finds herself aimlessly scrolling through this content, taking it in without aiming at or inquiring into a specific question. Many of her friends post news articles which contain (accurate) information about the winner of a presidential election. Leila is not interested in the presidential election

and does not pay a lot of attention to this news. Nevertheless, after it is shared several times by her friends, she adopts a correct belief about the winner of the presidential election.

Leila's consumption of social media does not seem excessive and it does not seem as if she is doing anything wrong in this scenario. In many ways, her use of social media offers many benefits to her, not least from an epistemic perspective. She is able to keep up with the social life of her friends and she can garner insight into news or other content which her online social circle has posted. Leila's situation is by far not the most extreme example of epistemic flooding: We could imagine an epistemic agent who is completely addicted to social media and spends much longer hours obsessing over their social media accounts and soaking up more information.

But the reason Leila's example is interesting is that even in her more moderate scenario, Leila inhabits an epistemically flooded environment.<sup>9</sup> The amount of content Leila encounters in her daily hour of private social media usage (notwithstanding a potential use of social media in her work life) is still large enough to overwhelm Leila epistemically. Diligently screening, fact-checking, and deliberating upon every bit of content she consumes in her time online would require considerable time and effort, at the very least much longer than the hour she spends to passively consume the content and quite likely more time and effort than she has in a day. Leila's epistemic environment is epistemically flooded because it is routinely filled with much more information than she can diligently process even though she can passively consume it in much less time.

We have good reason to assume that Leila is not alone in her situation: Her circle of friends consume social media in a similar way and often they are directly sharing content amongst themselves. Nor is her social circle largely different from the society she inhabits or even the global online community. It is estimated that roughly four billion people own a social media account, a number which is steadily increasing (Dean, 2021). While the manner in which social media is used varies among individual users, it is important to note that Leila is hardly on the extreme end of the spectrum given her more moderate use of social media. It is quite likely that many and maybe even most social media users' epistemic environment is even more flooded.

In this specific example, there is no negative impact on Leila's epistemic attainment despite inhabiting an epistemically flooded environment. However, she is still prone to adopting beliefs in a way which bypasses her epistemic agency. Leila has neither inquired into nor actively engaged with sources about the winner of the presidential election. Nevertheless, when questioned about it, she would formulate the correct belief about the winner. This is a belief which she has acquired passively by inhabiting an epistemically flooded environment. Even though she is not actively inquiring into the question, she is quite directly confronted with this information when scrolling through

---

<sup>9</sup> One might wonder what the quantitative or qualitative threshold for an epistemically flooded environment is. While I cannot offer a quantitative answer, I will propose the following heuristic: Cases of epistemic flooding should be clear-cut. The amount of information we are confronted with (and the necessary period of time we are flooded) needs to vastly outnumber our ability to diligently process it. In cases in which we are uncertain whether this is the case, it is better to not accept them as genuine instances of epistemic flooding. Crucially, I believe that social media feeds are flooded with so much information by design that Leila's example easily clears this hurdle, even if her social media use is not particularly excessive.



her feed. Although Leila's epistemic agency is undermined in this scenario, it seems as if she does nothing wrong in this case. However, there might be some objections to this reading of the scenario.

One, we might dispute that Leila forms a belief in the first place. I think this is a pertinent worry and one which can only be answered if we have a certain understanding of what constitutes a belief. If we adopt a Spinozan Model of belief acquisition, then it seems correct to say Leila has formed a belief about the winner of the presidential election. I will discuss this model in more detail in Sect. 3 of the paper. However, even if we do not grant that Leila forms a belief in this scenario, we should still assume that the information she encounters has some kind of impact on her. This could amount to stored information or a sub-doxastic state about the winner of the presidential election. Both would make her predisposed towards taking a certain stance regarding this question in the future. What should be noted is that this process is commonplace in epistemically flooded environments: Leila is prone to adopting many beliefs (or at least sub-doxastic states) in this manner.

The second worry is that Leila is not justified in forming a belief in this scenario regardless of whether she actually does so. Given the fact that she did not properly inquire into the question and has only sparingly engaged with the evidence, it would be better for her to suspend. To a certain extent, this worry seems justified, however, I think there are two important mitigating factors which would excuse Leila in this case. The first is a question of control: Given Leila's passive uptake of information, it is unclear whether she can actually help forming this belief. And even if we think that she does not form a belief, but rather a sub-doxastic state such as an alief in this scenario, it is unclear how much control she has in doing so. The second mitigating factor is that Leila's social media feed features posts from her circle of friends and peers. It is reasonable of her to be less skeptical towards content posted by people she trusts. Levy (2022) makes a strong case for the importance of placing trust in our epistemic peers. According to him even encountering false information through testimony should not lead us to abandon the trust we place in our friends. This is because trust is intrinsically valuable from an epistemic point of view according to Levy (2022). Even accounts which worry about the erosion of trust in the face of fakes or falsehood such as Harris (2022) maintain that trust is of vital importance for epistemic attainment. Therefore, it seems reasonable for Leila to place trust in her friends' social media testimony. From this perspective, Leila is entitled to her behavior in the scenario described above.

A final question we might be interested in is whether Leila is in anyway engaging in an inquiry by 'aimlessly scrolling' through her social media feed. While she is not directly interested in asking or answering a question, her aimless scrolling might still be interpreted as a minimal amount of inquiry. Friedman considers the possibility of a "more passive or minimal notion of inquiry" (2020, p. 13). However, even on a more broadscale view of inquiry, Leila's example would not qualify. This is because some minimal level of interrogative attitude or mild curiosity needs to be present in all inquiries, which is not the case for Leila (Friedman, 2020, p. 9). Inquiries on the other hand need to be intentional according to Friedman (2020, 2023).

### 2.3 Online epistemic flooding

While I believe that social media or online environments are not the only instances of epistemic flooding, I do believe that it is reasonable to pay special attention to online epistemically flooded environments for two reasons: 1. The internet and social media in particular have made epistemically flooded environments much more common. 2. Epistemically flooded social media environments feature certain unique characteristics which make the worries we have about epistemic flooding more poignant. While I take the first point to be uncontroversial as an empirical claim, the unique characteristics of epistemically flooded social media environments deserve more attention.

First, epistemic environments on social media are flooded *by design*. Social media feeds thrive off of engagement and the prolific sharing of content. If you find yourself scrolling through an empty social media feed, you would assume it is faulty.<sup>10</sup> Second, the kind of information we encounter on social media differs from other epistemically flooded environments. We are more likely to encounter news content or engage with testimonial sources online. This also makes worries about epistemic flooding in conjunction with fake news and misinformation (which I will discuss in the next section) more poignant. Third, the content we are shown on social media is not neutral but is selected by algorithms which for the most part aim at showing us content we find engaging. Given the sophistication of algorithms on large social media platforms such as Facebook or Twitter, we can assume that content encountered on social media is quite carefully curated in order to be engaging for the user. These algorithms are also likely to exploit cognitive biases human agents are vulnerable to, such as the repetition effect or confirmation bias (Fazio, 2020). Although the algorithms may not *solely* aim at driving engagement with their platforms, social media platforms do rely on keeping their users engaged. As such, there is a good argument to be made that social media algorithms are likely to select personalized content which we find engaging in most cases, rather than catering to our best epistemic interests.

Given these unique characteristics, it is worth considering Online Epistemic Flooding (OEF) as a special kind of epistemic flooding. While I am interested in introducing the concept of epistemic flooding in a more general sense, the rest of this paper will focus on OEF in particular. This also makes sense with regard to the discussion of epistemic flooding and fake news in the following section. The first two characteristics of OEF highlighted above regard the frequency of flooding as well as the kind of information we are most likely to encounter. Both factors make OEF particularly worrisome with regard to the threat to our epistemic attainment induced by (O)EF. The worries about social media algorithms and the exploitation of cognitive biases

---

<sup>10</sup> It might be argued that social media platforms are a bit more heterogenous than described here. I agree with this assessment: There could definitely be social media platforms which are not epistemically flooded (or at least ways to interact with platforms in order to avoid epistemic flooding). This might especially be true as a historical claim about older social media platforms from the early days of the internet. I am oversimplifying this claim slightly because I believe that the majority of social media platforms now are permanently epistemically flooded (barring fringe cases in which one is new to a platform and has not formed a following yet). This is particularly true for large social media platforms such as Facebook or Twitter. Not only are these platforms flooded with content but these platforms become more desirable the larger their user base is. It is therefore fair to say that modern day social media platforms are permanently flooded by design (notwithstanding a few exceptions).

on the other hand might also bear negatively on our epistemic attainment but more poignantly, they underline the loss of epistemic agency induced by epistemic flooding. Being shown content which an algorithm has pre-selected based on data it has about what we are likely to find engaging makes us more vulnerable to being manipulated by the content we encounter in OEF environments. More worryingly, we may not even be aware of this additional loss of agency in OEF. These are all good reasons to think about OEF environments in particular, especially if we are worried about their role in the spread of fake news.

### 3 The combined threat of fake news and epistemic flooding

#### 3.1 Epistemic flooding and spinozan belief acquisition

Epistemic flooding has become commonplace because the technological advancements of the internet have facilitated spreading and accessing content online. This change means that information is much more readily available to us than ever before. From an epistemic point of view this means that for most inquiries we have, evidence is extremely easily available. But our evidential situation has not only changed if we are acting as active inquirers. The nature of social media means we are constantly inundated with information on topics and questions we are not actively inquiring into. Scrolling through a social media feed exposes epistemic agents to an avalanche of information. This information however, would be incredibly time-consuming to process diligently and for the most part is not actively engaged with.

Even such ‘passive’ exposure to vast amounts of information could massively impact us as epistemic agents. Indeed, some of the most pernicious effects of epistemic flooding occur precisely when epistemic agents are exposed to content passively. This is especially true if we buy into the Spinozan belief acquisition model. The Spinozan model is usually juxtaposed with the Cartesian model. The Cartesian view holds that we can understand a proposition without believing it. In this model, understanding precedes belief. Ideally, rational agents first understand a proposition and then, after rationally deliberating its merits, they can decide whether to believe it or not. The Spinozan Model reverses this order: In the Spinozan Model, when we first encounter a proposition, we initially believe it. It requires additional mental effort to subsequently understand it. Only once we have undergone this additional work can we decide to disbelieve a proposition. On this view, belief comes easy and understanding and potential disbelief require more effort.

While the Spinozan Model was not exceedingly popular, Dan Gilbert revived interest in this model and provided empirical evidence from psychological studies to support it (1991, 1993). Mandelbaum (2014) is a philosopher who offers a thorough defense of the model.<sup>11</sup> It is apparent why the Spinozan Model makes the threat of epistemic flooding even more poignant: If we find ourselves in epistemic environments which are routinely flooded with more information than we can diligently process and which we often consume only passively, it seems likely that we would form many

<sup>11</sup> For a paper arguing against Mandelbaum and the Spinozan Model, see Peters (2017).

beliefs without properly vetting them. And since it requires more effort to disbelieve, we might not have the time to reevaluate most of the new beliefs we form.

It is important to note that this should concern us even if we do not fully buy into the Spinozan Model. Many epistemologists feel that the kind of ‘beliefs’ described in the Spinozan Model are not truly beliefs at all, at least not in the sense in which we are interested in the term in epistemology. It is possible that ‘beliefs’ in the Spinozan Model are merely “stored information” and that this alone is not a belief (Whittle, 2012). But while this may be convincing on a semantic level and negate some of the impact of the Spinozan Model, it does not fully alleviate the concerns it raises with regard to epistemic flooding. Because even if most of the ‘beliefs’ which are formed in epistemically flooded environments without due diligence are not beliefs proper at all but rather stored information, they still pose an epistemic threat. If false ‘beliefs’ or stored (mis-)information are taken up without being re-evaluated, they can still impact our epistemic behavior negatively or indirectly lead to the formation of actual false beliefs (or prevent us from forming true beliefs which conflict with them).

Even if we do not form beliefs in epistemically flooded environments, it is possible that we adopt sub-doxastic states instead. While it is possible that we could adopt a host of different sub-doxastic states, I will focus on one in particular, namely *aliefs*. Gendler describes aliefs as “a mental state which is [...] associative, automatic and arational” (2008, original emphasis, p. 641). While this state does not amount to belief, it has a huge impact on our behavior and can tamper with our epistemic attainment indirectly. Aliefs for instance can explain implicit biases, which could negatively impact our epistemic goals. Even if we do not form beliefs directly by being exposed to information in epistemically flooded environments, we are still prone to adopting sub-doxastic states. It is unclear to what an extent we have control over sub-doxastic states such as aliefs. At the same time, they can lead us to biases which stymie our epistemic attainment. As such, we should be worried about the amounts of information we are exposed to even if we do not buy into the Spinozan model of belief acquisition. Accepting that we form aliefs and that aliefs can bear negatively on our epistemic attainment might be less controversial to accept for some than the Spinozan model.

I will leave the question of whether the Spinozan model is right and how strictly we should define ‘belief’ open here. However, I maintain that the Spinozan model does highlight one of the main difficulties of epistemic flooding, one which remains troublesome even if we do not agree with the Spinozan definition of ‘belief’. The worries raised by the Spinozan model seem particularly poignant with regards to the first worry about epistemic flooding, which is that it bears negatively on our epistemic attainment. We should be most concerned by beliefs or even sub-doxastic states coming to us ‘easy’ if the information we are confronted with is false or epistemically harmful in other ways. However, they address the second worry about epistemic agency as well, which seems to be bypassed in this belief formation process. Of course, the undermining of epistemic agency seems to be a problem within the Spinozan model of belief formation (or even the adoption of aliefs) outside of epistemically flooded environments as well. But epistemic flooding would at the very least increase the frequency with which our epistemic agency is undermined significantly. As such, the Spinozan model of belief acquisition raises problems independent of the content we encounter in an epistemically flooded environment.

### 3.2 Fake news

If epistemic flooding leaves us vulnerable to believing misinformation because it lowers our ability to diligently screen new information, it seems clear that fake news and mis- and disinformation in a broader sense are especially harmful in epistemically flooded environments. Since fake news has a greater chance of being accepted and not diligently screened in epistemically flooded environments, fake news and epistemic flooding pose a particularly strong combined threat. This threat should not be neglected, especially since both epistemic flooding and fake news are already worrisome on their own. I offer the following definition of fake news in a previous paper:

*Fake news:* Fake news is misleading information intentionally published and presented as news which has the function of deliberately misleading its recipients about its status as news. (Anderau, 2021, p. 214)

This definition places a lot of emphasis on the intention underlying fake news and marks out fake news as a form of disinformation. However, for the purposes of this paper, most other philosophical definitions of fake news would be acceptable as well (such as Jaster & Lanus, 2021, Croce & Piazza, 2021, Fallis & Mathiesen, 2019, Gelfert, 2018, Pepp et al., 2019, Rini, 2017). Epistemic flooding heightens the chance we take up false information without diligently processing it and as such would exacerbate not only fake news but also disinformation and misinformation in a broad sense. Therefore, the precise definition of fake news should not be a pressing concern here.<sup>12</sup> It is worth highlighting that Pepp et al.'s (2019) view posits that the main threat of fake news is not that people necessarily believe it but rather that they treat it as if it was produced by "standard journalistic standards" when it has not (2019, p. 69).

It is widely accepted in the philosophical literature that while fake news is not new, it has become increasingly widespread on social media because of the ease with which it can be disseminated there (McIntyre, 2018; Jaster & Lanus, 2019, Pepp et al., 2019). This is a similarity shared by fake news and epistemic flooding. Since both are more prevalent online, their combined threat seems especially potent in the epistemic environment of the internet. Let us revisit the example of Leila from above to illustrate how the two concepts work in unison:

*Leila II.* Imagine Leila is an epistemic agent active on three social media platforms. Her everyday routine includes scrolling through her social media feeds for half an hour each when she wakes up and before she goes to bed. Often she finds herself aimlessly scrolling through this content, taking it in without aiming at or inquiring into a specific question. Many of Leila's friends on social media post content spreading fake news promoting anti-vax sentiments. Leila does not pay active attention to these posts and mostly just sees their headlines. The topic of the Covid-19 vaccination does not initially interest her and neither do politics.

<sup>12</sup> For the purpose of brevity I ignore voices who believe 'fake news' as a term is not useful and should be abandoned such as Habgood-Coote (2019, 2020). This paper does pre-suppose the default assumption that fake news is real and worrisome from an epistemic point of view. However, readers who would balk at such an assumption could still consider the combination of epistemic flooding and misinformation in general as a problem instead.

Nevertheless, after being exposed passively to many of her friends' posts, she starts to believe that she should not take the Covid-19 vaccine.

Again, this example does not seem too far-fetched nor do Leila's initial actions seem extreme. Nevertheless or maybe even because of this, it is a useful way to demonstrate how epistemic flooding can leave an epistemic agent vulnerable to fake news. If we grant that Leila did nothing wrong epistemically in the previous example, then we need to grant that she does nothing wrong here either. Regardless, she still ends up believing in fake news. What is striking about this example is that Leila might very well have been able to ward off fake news if she had encountered it in (a) an active inquiry on her part and (b) a less flooded epistemic environment.

Of course, there might be skepticism about whether Leila has formed a belief here at all and if she has, whether she has been justified in doing so. Regarding the first question, I think this problem can be sidestepped because we can expect similar outcomes even if we do not grant that her attitude should be called a belief. Even if all she has developed in this scenario is a sub-doxastic state such as an alief (consisting of a negative reaction to the Covid-19 vaccine), this may still prevent her from taking the vaccine. In fact, this is particularly likely in the case of an alief since they are action-oriented attitudes (Gendler, 2008).<sup>13</sup> But we are still left with the question of whether Leila is blameworthy for her belief (in a Spinozan sense) or alief.

At the very least, it is hard to see how much doxastic control Leila has in this situation. It is likely that Leila has no control over forming a belief nor an alief in this context. Avoiding this attitude is something which she might only be responsible for indirectly, say by pre-emptively regulating her epistemic environment. I will explore to what an extent it is reasonable to hold Leila responsible for governing her own epistemic environment in Sect. 4. In addition, Leila has a special reason to place trust in her friends: Since epistemic trust is vital for epistemic attainment, Leila is particularly justified in taking the online testimony of her friends at face value (Levy, 2022).

Such trust could also be deliberately exploited: The Russian Firehose of Falsehood model for instance can be seen as a deliberate attempt to exploit our vulnerability in epistemically flooded environments (Paul & Matthews, 2016). One of the strategies of this model of propaganda, which is usually disseminated online, is to overwhelm victims through the sheer quantity of propaganda even in cases in which the original propaganda is not particularly convincing (Paul & Matthews, 2016). Pepp et al. (2022) also highlight how public conversations themselves can be made knowledge-resistant by deliberately flooding the space in which they take place with less relevant news. The success and existence of such tactics suggest that epistemic flooding is a real concern and one which can be exploited by unscrupulous actors. However, while such targeting is especially worrisome, epistemic flooding and fake news can pose problems even if

---

<sup>13</sup> Levy (2017) also highlights the dangers of fake news even if it does not lead to the adoption of (false) beliefs.

they are not paired as a deliberate strategy. In fact, epistemic flooding can cause issues even in the absence of fake news.<sup>14</sup>

One question to keep in mind is whether it would make a difference if we subscribe to Pepp et al.'s (2019) view that the main issue with fake news is not that we believe its content, but rather that we treat it as if it has been produced by standard journalistic practices when it has not. While I think this view can be appealing when trying to define fake news, I would argue that adopting such a view would still make the combined threat of fake news and epistemic flooding more worrisome.

Firstly, because treating fake news as if it has been produced by standard journalistic practices when it has not is worrisome and we would still be more likely to misjudge fake news this way in epistemically flooded environments.<sup>15</sup> Secondly, while it is not a necessary condition that the content of fake news is false on this definition, fake news still usually features false information even according to Pepp et al. (2019). Since epistemic flooding raises a worry about the volume of misinformation encountered, this concern will persist even if fake news features wrong or misleading content only most of the time.

And finally, Pepp et al.'s (2019) view still presupposes that readers of fake news are directly and actively consuming it. In the very passive encounter of fake news outlined in Leila II, it still seems plausible that consumers would develop either beliefs (in the Spinozan sense) or sub-doxastic states about the content of the fake news they are exposed to. At the very least, it seems at least equally plausible that they might passively adopt some sort of attitude towards the content of fake news than they would towards the manner in which it has been produced. Therefore, epistemic flooding will still remain especially pernicious in combination with fake news even if we adopt Pepp et al.'s (2019) view.

There is empirical evidence that, contrary to popular belief, people are remarkably good at identifying fake news stories (Altay et al., 2020). The problem of this study is similar to that of many of the philosophical examinations of fake news: Both assume that epistemic agents are actively engaging with fake news sources and following standard norms of inquiry. However, this is often not possible in flooded epistemic environments.<sup>16</sup> And it is exactly in such surroundings that fake news or the Russian Firehose of Propaganda model thrives. We are vulnerable in cases of epistemic flooding and cede some of our epistemic agency. This is an issue in epistemically flooded environments in general. However, these worries become exacerbated in environments in which we are flooded with fake news or other epistemically harmful content. While epistemic flooding poses a minimal threat independent of the quality of information we

<sup>14</sup> In the cases Pepp et al. (2022) have in mind, the strategy does not necessarily rely on *fake* news being spread but rather just news irrelevant to certain public discussions. While this would not rely on fake news per se, it is an example of how epistemically flooded environments can be deliberately exploited.

<sup>15</sup> Pepp et al. (2019) explicitly mention that we are more susceptible to fake news when it is shared on social media. OEF in particular would be a strong concern even when adopting this view of fake news.

<sup>16</sup> If we think of the example of Leila, it is debatable whether her 'aimless scrolling' through social media is a form of inquiry at all. If it is, we might think of this as a kind of 'bent inquiry' (in a nod to Rini's 2017 description of social media testimony as 'bent' testimony) which has become commonplace on social media. Even when we begin by inquiring into a specific question online, we run the risk of losing ourselves in the act of 'aimlessly scrolling' eventually, falling prey to what we might call a zetetic form of the Gruen Transfer.

encounter in the flooded environment, being exposed to epistemically harmful content becomes an even bigger risk in combination with epistemic flooding. This effect can but need not be deliberately exploited as it is within the firehose of falsehood propaganda model. This poses a dilemma for epistemologists: Traditional tools of epistemology such as evidential norms or the norms of inquiry are of no or at least less use in cases of epistemic flooding. As such, we need to consider different normative avenues to face these problems. Although this paper is very much located in non-ideal epistemology, the problem of epistemic flooding might be taken as a catalyst to reconsider epistemic normativity from an ideal point of view as well.

## 4 Possible solutions

### 4.1 Adapting our everyday epistemic practices

One of the most straightforward solutions is that we adapt our epistemic environment preemptively to prevent epistemic flooding. The simplest solution is to abstain from social media use or to at least heavily modify it. The latter option would include a more stringent and conscious decision of who we include in our social media circles and whose content we make visible to us. It would also mean trying to engage with content only actively and spending a lot of time fact-checking and deliberating on it. While a radical modification of (or complete abstinence from) social media use is theoretically possible, I do not believe it is viable as a practical solution for two reasons: 1. It comes at too high of a cost 2. the approach considered is too individualistic. I will consider the second point in the following subsection.

Regarding the first point, I believe that abstaining from or minimizing and micro-managing our social media use places a strong burden upon epistemic agents. Abstaining from social media can come at high social or practical costs which seem unfair to ask of individual epistemic agents. For some, abstaining from social media might simply not be possible because they rely on it for their work. This also highlights why epistemic flooding affects epistemic communities as a whole and not just individuals. The social repercussions of leaving social media in an environment in which everybody uses it are more severe than they ever could be in an individual example of epistemic flooding. It should be noted that despite the epistemic risks associated with social media use, abandoning social media might have epistemic repercussions as well: We are increasingly reliant on information we glean from social media, whether it be important personal information within our social circle or news about the world at large. Being deprived of the tool of social media while everybody else uses it puts us at a disadvantage epistemically, one which might be hard to mitigate even if the upshot is that we are less vulnerable to misinformation in an epistemically flooded environment.



## 4.2 Epistemic environmentalism

The second problem is that the approach is too individualistic even though it aims at solving an issue which ails an epistemic community as a whole. Leaving the sole responsibility of improving our epistemic environment with individuals seems both unfair since epistemic agents are only partially responsible for their epistemic environment and inefficient since they are limited in their ability to improve it (and where they are, face harsh social or practical sanctions for certain epistemically beneficial interventions).

This is not to say that adapting our everyday epistemic practices is completely pointless or should be discouraged. It is absolutely reasonable for epistemic agents to be more diligent in the choices which are open to them in regulating their online epistemic environment. Individuals should be encouraged to block or confront people in their social media circle spreading misinformation and be more diligent in the way they process information online. Another possible solution is to reconsider our zetetic norms to help our epistemic behavior online become more focused and attuned to our epistemic attainment. This might encourage us to think along the lines of Friedman's (2020, 2023) proposed zetetic turn in epistemology. While I am sympathetic to this notion for the most part, our zetetic norms will not be sufficient to address many instances of epistemic flooding, including the examples of Leila discussed above. This is because Leila is not inquiring at all in these examples, not even in a more 'minimal' understanding of inquiry proposed by Friedman (2020). Therefore, zetetic norms cannot give us a full answer to the issue. Pointing to zetetic norms also leaves the responsibility with the individual inquirer or social media user, which might again place too strenuous of a burden on them.

Similar to how individual diligence is encouraged but not sufficient in order to tackle problems like climate change in order to improve our environment, we need a normative framework which allows us to address epistemic environments from a collective perspective. Luckily, there is a normative framework in the literature which allows us to speak normatively (and collectively) about epistemic environments and how to improve them. This is Ryan's notion of *epistemic environmentalism*. Ryan himself references "geo-environmentalism" as an analogy for his concept, which he uses in order to hold accountable what he calls "epistemic pollutants" (2018, p. 99).<sup>17</sup> Epistemic environmentalism takes the social dimension of epistemic environments into account and holds both individuals and institutions responsible for bettering the epistemic environment they inhabit (Ryan, 2018, p. 101).

Apart from being able to hold those spreading fake news accountable, this also allows us to criticize those who indirectly enable its spread.<sup>18</sup> These are

---

<sup>17</sup> We should be careful about the comparison to geo-environmentalism since it can be slightly disanalogous in some aspects. One important difference is that epistemic flooding is an issue independent of content, whereas we might be extremely selective about what is regarded as a pollutant in geo-environmentalism. Nevertheless, what is important is that epistemic environmentalism gives us a framework which allows for normative claims about bettering the epistemic environment as a whole. Levy (2022) and Blake-Turner (2020) offer other views which promote bettering the epistemic environment.

<sup>18</sup> Another normative framework which might be useful in this regard is Kawall's (2020) notion of *epistemic credit*, provided that we extend it in order to include 'epistemic discredit' as well.

social media platforms such as Twitter and Facebook themselves. Arguably, these platforms are best placed to prevent the spread of fake news on social media and can be considered epistemic pollutants.<sup>19</sup> Ensuring that social media platforms actually assume responsibility is a task which extends beyond epistemology. Arguably, it would require an intervention on an institutional level, which could be considered worrisome for ethical or political reasons, such as concerns about freedom of speech.<sup>20</sup> Weighing these arguments goes beyond the scope of this paper. However, the combined threat of epistemic flooding and fake news described here might at least offer an epistemological motivation to consider such actions.

## 5 Conclusion

This paper has outlined the novel concept of epistemic flooding, an underexamined threat which is increasingly faced by epistemic agents in online epistemic environments. This threat makes the danger of fake news online especially poignant, since both fake news and epistemic flooding grow more dangerous in combination. Addressing their dual threat requires us to rethink how we conceive of epistemic normativity. The normative tools of traditional epistemology may increasingly be found too blunt to combat it. It is unclear if traditional evidential norms can help us if the problem is that we are routinely inundated with more evidence than we can diligently process.

Addressing epistemic flooding requires us to think in normative terms about our epistemic environments and how to better them. It also forces us to consider epistemic normativity from a communal, rather than solely an individual perspective. It is possible that an effective solution hinges upon an intervention into our online epistemic environment on an institutional level. While enacting such a political intervention goes beyond the field of epistemology, re-thinking epistemic normativity might still be necessary to address epistemic flooding. It will force us to consider actions epistemic agents take outside of directly reacting to evidence or acting as active inquirers from a normative perspective. Epistemic flooding forces us to consider our passive and indirect responsibilities as epistemic agents.<sup>21</sup> This indirect form of epistemic normativity deserves more attention in the literature. Studying it can hand us more and sharper tools to address new threats such as epistemic flooding which develop in the epistemic environment of social media.

**Acknowledgements** I am a Research Associate at ACEPS (African Centre for Epistemology and Philosophy of Science, University of Johannesburg). I would like to thank the following people for their insightful comments and support which greatly helped me improve this paper: Anne Meylan, Kevin Reuter, Axel Gelfert, Melanie Sarzano, Pablo Hubacher Haerle, Karen Frost-Arnold, Fenner Tanswell, Finnur Dellsén,

<sup>19</sup> For empirical evidence that social media companies such as Facebook enable the spread of fake news and misinformation and hold the power to prevent it, see Rosza (2021).

<sup>20</sup> According to Ryan (2018), epistemic environmentalism itself is split over state intervention in order to better epistemic environments. However, given the state's role in education and the epistemic environments in which it takes place, it does seem possible that state intervention in epistemic environmentalism is an option. State regulation of social media epistemic environments might still remain controversial, however.

<sup>21</sup> This indirect form of epistemic normativity may be analogous to considering a football player's responsibilities beyond directly kicking or playing the ball, such as their positioning or general off-the-ball play.

Endre Begby, Benoit Guilielmo, Neri Marsili, Peter Graham, Daniel Barbarussa & Lola Medina Vizuete, the anonymous reviewers, the whole ZEGRa team, and Caroline Nganga.

**Funding** Open access funding provided by University of Zurich.

## Declarations

**Conflict of interest** The author has no competing interests to declare that are relevant to the content of this article and has no relevant financial or non-financial interests to disclose.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Altay, S., Hacquin, A.-S., & Mercier, H. (2020). Why do so few people share fake news? It hurts their reputation. *New Media & Society*. <https://doi.org/10.1177/1461444820969893>
- Anderau, G. (2021). Defining fake news. *KRITERION—Journal of Philosophy*, 35(3), 197–215. <https://doi.org/10.1515/krt-2021-0019>
- Bermes, A. (2021). Information overload and fake news sharing: A transactional stress perspective exploring the mitigating role of consumers' resilience during COVID-19. *Journal of Retailing and Consumer Services*, 61, 1–10. <https://doi.org/10.1016/j.jretconser.2021.102555>
- Blake-Turner, C. (2020). Fake news, relevant alternatives, and the degradation of our epistemic environment. *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1725623>
- Croce, M., & Piazza, T. (2021). Misinformation and intentional deception: A novel account of fake news. In N. Snow & M. S. Vaccarezza (Eds.), *Virtues, democracy, and online media: Ethical and epistemic issues* (pp. 49–63). Routledge.
- Dean, B. (2021). Social network usage & growth statistics: How many people use social media in 2021? *Backlinko*. Retrieved July 7, 2022, from <https://backlinko.com/social-media-users>.
- Fallis, D., & Mathiesen, K. (2019). Fake news is counterfeit news. *Inquiry*. <https://doi.org/10.1080/0020174X.2019.1688179>
- Fazio, L. K. (2020). Repetition increases perceived truth even for known falsehoods. *Collabra: Psychology*, 6(1), 38. <https://doi.org/10.1525/collabra.347>
- Friedman, J. (2020). Zetetic Epistemology. In B. Reed & A. K. Flowerree (Eds.), *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge.
- Friedman, J. (2023). The aim of inquiry. *Philosophy and Phenomenological Research*. <https://doi.org/10.1111/phpr.12982>
- Gelfert, A. (2018). Fake news: A definition. *Informal Logic*, 89(1), 84–117. <https://doi.org/10.22329/il.v38i1.5068>
- Gendler, T. S. (2008). Alief in action (and reaction). *Mind & Language*, 23, 552–585. <https://doi.org/10.1111/j.1468-0017.2008.00352.x>
- Gilbert, D., et al. (1993). You can't not believe everything you read. *Journal of Personality and Social Psychology*, 65(2), 221–233. <https://doi.org/10.1037/0022-3514.65.2.221>
- Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, 46(2), 107–119. <https://doi.org/10.1037/0003-066X.46.2.107>
- Habgood-Cooté, J. (2019). Stop talking about fake news! *Inquiry*, 62(9–10), 1033–1065. <https://doi.org/10.1080/0020174X.2018.1508363>

- Habgood-Coote, J. (2020). Fake news, conceptual engineering, and linguistic resistance: Reply to Pepp, Michaelson and Sterken, and Brown. *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1758770>
- Harris, K. R. (2022). Real fakes: The epistemology of online misinformation. *Philosophy & Technology*, 35(83), 1–24. <https://doi.org/10.1007/s13347-022-00581-9>
- Jaster, R., & Lanius, D. (2019). *Die Wahrheit schafft sich ab: Wie Fake News Politik machen*. Reclam.
- Jaster, R., & Lanius, D. (2021). Speaking of fake news: Definitions and dimensions. In S. Bernecker, A. K. Flowerree, & T. Grundman (Eds.), *The epistemology of fake news*. Oxford University Press.
- Kawall, J. (2020). Testimony, epistemic egoism, and epistemic credit. *European Journal of Philosophy*, 28, 463–477. <https://doi.org/10.1111/ejop.12447>
- Levy, N. (2017). The bad news about fake news. *Social Epistemology Review and Reply Collective*, 6(8), 20–36.
- Levy, N. (2022). In trust we trust: Epistemic vigilance and responsibility. *Social Epistemology*, 36(3), 283–298. <https://doi.org/10.1080/02691728.2022.2042420>
- Mandelbaum, E. (2014). Thinking is believing. *Inquiry*, 57(1), 55–96. <https://doi.org/10.1080/0020174X.2014.858417>
- McIntyre, L. (2018). *Post-truth*. MIT Press.
- Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme*, 17(2), 141–161. <https://doi.org/10.1017/epi.2018.32>
- Paul, C., & Matthews, M. (2016). The Russian “firehose of falsehood” propaganda model. *Rand Corporation*, 2(7), 1–10.
- Pepp, J., Michaelson, E., & Sterken, R. K. (2019). What’s new about fake news? *Journal of Ethics and Social Philosophy*, 16(2), 67–94. <https://doi.org/10.26556/jesp.v16i2.629>
- Pepp, J., Michaelson, E., & Sterken, R. (2022). Relevance-based knowledge resistance in public conversations. In *Knowledge resistance in high-choice information environments* (pp. 106–127). Routledge.
- Peters, U. (2017). On the Automaticity and Ethics of Belief. *Teoria. Rivista Di Filosofia*, 37(2), 99–114. <https://doi.org/10.4454/teoria.v37i2.20>
- Rini, R. (2017). Fake news and partisan epistemology. *Kennedy Institute of Ethics Journal*, 27(2), E43–E64. <https://doi.org/10.1353/ken.2017.0025>
- Rosza, M. (2021). Facebook could have stopped 10 billion impressions from “repeat misinformers”, but didn’t: Report”. *Salon*. Retrieved July 7, 2022, from <https://www.salon.com/2021/04/12/facebook-could-have-stopped-10-billion-impressions-from-repeat-misinformers-but-didnt-report/>.
- Ryan, S. (2018). Epistemic environmentalism. *Journal of Philosophical Research*, 43, 97–112. <https://doi.org/10.5840/jpr201872121>
- Ryan, S. (2021). Fake news, epistemic coverage and trust. *The Political Quarterly*. <https://doi.org/10.1111/1467-923X.13003>
- Whittle, B. (2012). Belief, information and reasoning. *Philosophical Perspectives*, 26, 431–446.

**Publisher’s Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.