

Attention for Whose Sake? Why Social Epistemology Needs to Consider Attention

Edward Windle

July 2025

1 Introduction

It is our lot to pay attention to things for a purpose. You are paying attention to reading this paper for your own ends; I am paying attention to writing it for mine. Attention can be directed properly or improperly, and how we direct attention influences our evidence, beliefs, knowledge, and understanding, which can be said to serve utility to some desired goal. But what about attention directed at finding something to attend? Here, I argue this is harmful when not directed by an individual's desired goal. Specifically, I introduce the novel concept:

ATTENTIONAL FUTILITY

The directing of attention to find something to attend with a goal.

I will evaluate the epistemic implications of this concept with credal epistemology and use it to evaluate the way users of social media platforms direct their attention. I argue that social media platforms produce a state of attentional futility within users through *hostile attentional scaffolds* (Voinea et al. 2024), which have epistemic consequences evaluable with credal epistemology.


In this paper, I motivate five claims about the epistemology of attention: (i) attentional states are normative—i.e., attention can be directed properly and improperly which incurs epistemic implications; (ii) singular instances, habits and patterns of directing attention can be evaluated normatively; (iii) the way groups and collectives direct their attention is non-summative, and this generates its own kind of epistemic implications; (iv) attentional futility is made harmful through hostile attentional scaffolds; and, (v) that attentional futility is epistemically harmful. These claims bring together the literatures on credal epistemology, the epistemology of attention, and philosophy of technology for the first time to argue the novel point that attention is undermined from a credal perspective.

I begin in §2 by unpacking the Blackwell-Good value of information and utility theory approach to credal epistemology. I then address (i)-(iii) to es-

establish an evaluable conception of the epistemology of attention in §3. Next, §4 deals with (iv), making clear how attention is manipulated and exploited through affective systems to create states of attentional futility. Finally, §5 addresses (v) by exploring how attentional futility is harmful in individual cases, when habituated, and at the group level. Upshot: attention has serious epistemic implications. Claim: these implications need to and can be considered by epistemology.

2 Framework of Credal Epistemology

Here are some triangles:



How many triangles are there? Let ϕ be the claim that there are 27 triangles. You can have certain beliefs about ϕ , such as being confident of it. What determines your beliefs about ϕ (and many other propositions) is your *evidence* for it¹.

Why should we care about evidence? Good (1967) gave an intuitive answer: *value of information theorem* (VOI). This formalises the idea that a set of evidence is *valuable* iff the expected utility of making a decision is higher than the other option, providing the process of gathering the evidence isn't too costly. So, say you want to bet on the truth of ϕ : you may decide to either (1) gather information about how many triangles there are; or (2) just bet without gathering information. Assuming your goal is to be correct, you are better off with (1) than (2)—the evidence is valuable here because you are more likely to be correct if you gather information. Hence, we have the slogan: evidence is valuable.

We can be more precise about the strength of the beliefs an individual holds by characterising them as *credence functions*: a function that assigns a value representing confidence, where $0 \leq$ or ≥ 1 , where 0 is 0% accurate and 1 is 100% accurate, is to ascribe a *credence*—i.e., a *degree of belief* in some proposition² (Pettigrew 2024, p.12). So, we may say 'Sam is 100% sure of ϕ ; she counted the number of triangles and deduced a proof also', and 'Ally is 40% sure of ϕ ; she only counted the number of triangles'. Here, we ascribe a credence of 1 to Sam and 0.40 to Ally.

But how do we evaluate credences? Well, they can be evaluated with their *epistemic* and *pragmatic* utility. The former assesses how accurately an individual's credences represent the world; the latter assesses how accurately

¹For now, I am neutral about what evidence is beyond this claim.

²Regarding the definition of propositions, I will say nothing more than that they can be conceived of as a set of possible worlds to which an individual attributes credences.

an individual's credences guide their actions. The evaluation is of the *outcome* of an individual's credences (*ibid.*). Moreover, presuming a rational subject pursues credences with higher value, we can say that this approach is at least teleological: if subjects should follow VOI, they achieve more valuable epistemic states.

For epistemic utility, I take the value of an outcome to depend on the accuracy with which a set of credences represents the world. Standardly, an *epistemic utility function* EU will evaluate an individual's credence function C that is the set of credences an individual holds at a world w , and return $EU(C, w)$ —a numerical value representing the accuracy of C at w (*ibid.*, p.48).

Pragmatic utility is the utility of the outcome of a decision, where the choice was determined by your credences. The value of the outcome is determined by how useful it was in the end; in turn, it relates to what an individual wants (i.e., *preferences* or *goals*). So, a *pragmatic utility function* PU is applied to a probabilistic model of possible credence functions P at world w , which result from the set of choices from a decision problem D . This gives us $PU^D(P, w)$ (*ibid.*, p.14)

For example, let us say I am considering whether or not to buy a new reliable car to replace my old unreliable car. The pragmatic utility of my current credences will be determined by the utility of the choice I choose because of them. So, if I do not buy a new car and my old one breaks down, then the utility of my credences will be lower than if they had determined I buy a new car. Formally, if buying a new car is determined by credence function B at w , and not buying a new car is determined by credence function B' at w , then we can say:

$$PU^D(B, w) > PU^D(B', w)$$

Notably, this means external factors outside the control of the individual are relevant to the utility, and hence value, of a choice. In this case, the external factor is the car breaking down at world w .

In addition to their epistemic and pragmatic utility, credences can be evaluated also by considering how they are formed. Relevant for our purpose is an *evidence-gathering episode*, which not only determines the credences we have but the credences we will get (*ibid.*, p.15). For example, presume I am planning on eating either mushrooms or peppers later, to which I currently have no preference. Suppose also that between now and when I plan to eat, I can go gather evidence and find out whether the mushrooms or peppers are spoilt. This leaves two options for each decision that will alter my credences³: I gather evidence that the mushrooms or peppers appear spoilt, when they are not; I gather evidence that the mushrooms or peppers are not spoilt, when they actually are. Consider a world where the mushrooms appear spoilt, but actually are not. The evidence-gathering episode's value is determined by the utility of eating the peppers and never eating the mushrooms. Essentially, the evidence-gathering

³ Assume that if I discover neither are spoilt, then my credences will not change as a result of gathering evidence; and if I discover both are spoilt, then my credences for both drop to 0.

episode has a value determined by the utility of the option it determines you take. Another way of stating this: credences determine your preferences.

Evidence-gathering episodes play a role in determining *expected* epistemic and pragmatic utility. Consider an evidence-gathering episode ε that always links the agent to truth (i.e., it is *factive*). An individual's credences after ε will be higher and garner more utility. Therefore:

$$EU(\varepsilon, w) > EU(C, w)$$

The same applies also for pragmatic utility:

$$PU^D(\varepsilon, w) > PU^D(C, w)$$

For one thing, this establishes the VOI. Essentially, in these idealised circumstances, it is always better gather evidence (see *ibid.*, pp.13-18 and pp.47-52).

However, evidence-gathering episodes are not typically *factive*. Moreover, evidence can be too costly. Essentially, there are times when we should not gather evidence because:

OPPORTUNITY COSTS

‘[G]iven our cognitive limitations, the pursuit of evidence along one line of inquiry will necessarily involve siphoning resources away from other (possible) lines of inquiry’
(Begby 2023, p.401).

I will adapt this to the credal framework: gathering evidence provides opportunity for bettering our credences in expectation. And because what evidence we gather determines what utilities we get, we should allocate these resources diligently. So, while we can gather evidence on one set of propositions, we could have been better off gathering evidence about another. As for guidance, an individual should not gather evidence on one goal if they can allocate their evidence-gathering resources to a different goal with more utility in expectation. Note that this is not necessarily bad. Provided the measures for utility allow an individual to choose what they prefer, an individual may have a preference for the other goal with less utility overall. However, allocating finite resources to propositions or goals that are not in an individual's interests is certainly a way they *could* go wrong. Keep this in mind as we turn to attention.

3 What is Attention

Before addressing (i)-(iii), it should be noted that, though (i)-(iii) are greatly inspired by Gardiner (2022), this is all our accounts share in common. For Gardiner opts for virtue epistemology; I opt for credal epistemology. I do this for two reasons. First, her approach is restricted to the evaluation of habits and patterns at the individual or group level. However, credal epistemology allows for the evaluation of singular instances also (this will come in handy later). Furthermore, credal epistemology allows for the evaluation of more specific cases of

habitual attentional direction, instead of the broader and less precise teleology of virtue epistemology.

Let us start by clarifying what we mean by attention. Because attention as a cognitive function has multiple diffuse definitions, I will only focus on one here, namely the selection-for-action account⁴. This takes attention as functional and part of agency, described by Wu (2011) as ‘the subject’s *selection of information* (input) that *guides* or otherwise *informs* his or her *response*’ (my emphasis, p.101). So, an agent directs their attention. An agent is someone who *performs actions* and *makes decisions* based upon their goals and preferences, which are partially formed by how they direct their attention.

A relevant structural feature of this definition is that it identifies attention as a process within the larger function of generating responses to information that attention helps select. So, attention directs cognitive processes to *specific evidence* that is used as a *means* to some further goal. Let μ be the total set of all available evidence. For all evidence δ ($\delta \in \mu$) selected by attention in an evidence-gathering episode, our definition of attention means that $\varepsilon \geq \delta$, where ε is the ideal factive evidence-gathering episode. So attention can give greater weight to some propositions over others. As such, attention can be both beneficial (by selecting what is needed for a specific aim), or detrimental (by selecting evidence that is misleading or false). Note that, according to VOI, this implies tentatively the normativity of attention—i.e., the more shared elements between δ and ε , the more accurate my beliefs; therefore, the better I directed my attention.

However, this definition is limited; it requires continuous action; hence, continuous response. This is *prima facie* false. Are we willing to concede that change in direction of attention *must* be intended to select *for an action*? No. Let us assume that it does: this means every direction of attention elicits a response. Consider the following⁵:

A driver considers taking an alternative route because her usual is very busy. She decides to gather evidence: she shifts her attention from the time it takes for the traffic to move; how many cars in front are taking the shortcut; an advertisement on the side of the road. Having attended to the advertisement for a moment, she shifts her attention back to the traffic.

If response is interpreted strongly, then the attentional shift to the advert and the information elicited by it must be considered as contributing to the decision she makes. This is not the case, for the attentional shift only influences her by detracting from potential evidence she *could* have gathered. Consider this with opportunity costs: directing attention to the advertisement costs time and attention which could be better directed to the decision problem. However, it does not necessarily cause a response, simply distracts—i.e., attention directed

⁴This is not to make any claims about the invalidity of any other accounts of attention.

⁵This example is greatly influenced by (Buehler 2018, p.4).

to the advertisement did not act to select for an action, but to distract from selection.

As such, two clarifications can be made. Firstly, response can be understood as a credence, and attention as a non-ideal factor that influences desirable credences positively or negatively through guiding evidence-gathering episodes. And because evidence gathering can shape epistemic and pragmatic utility, attention can affect these by guiding the evidence-gathering episode to one proposition over another. Clearly, this relates to opportunity costs. For attention is what directs the cognitive resources over time which are the cost. As such, we can claim that attention can be directed properly and improperly, which is to say it should guide evidence-gathering episodes to propositions that the individual either values more or find more useful. We can thus evaluate two distinct evidence-gathering episodes where attention directs you to different bodies of evidence. The difference in utility between outcomes of your credences in expectation can be compared, and we can claim: it was better for you to direct your attention to this over that. Upshot: this partly details (i) and (ii).

Secondly, we can think of directing attention as a behaviour—i.e., something we do. Beyond (not) contributing to our credences, characterising attention this way reveals two further dimensions: habits and patterns of directing attention are often what makes direction of attention normative (i.e., good or bad). Concerning habits, consider the previous case of the driver. If this is a singular case of distraction, then it does not necessarily have any real negative ramifications. However, say the driver is consistently distracted by advertisements; this increases the likelihood she will make mistakes while driving—a potentially fatal consequence. Alternatively, if she is rarely distracted whilst driving, then we can say the pattern of her attention is to attend to the road; therefore, she is more likely to drive well, reducing the risk of accidents and increasing the efficiency of her driving. So, the normativity of attention frequently manifests itself within the patterns and habits of how it is directed: we can have *good* and *bad* habits which may ‘*appear* as compulsion [but] require some degree choice’ (White 2024, p.43). Upshot: if someone habitually directs their attention improperly, then this can be harmful; conversely, if someone habitually directs their attention properly, this can be beneficial. Note, this is distinct from the previous case because it evaluates broader patterns of directing attention, while the previous clarification only made it clear how attention direction can be good or bad in singular instances. Upshot: this completes (i) and partly details (ii).

This account of attention also works well when considering how attention affects groups. Gardiner (2022) points out that groups can display attentional traits. To return to the driver, suppose she is no longer in traffic, but the road features many adverts and there are other vehicles. If the rest of the traffic also directs their attention to the advertisement, then less people are focusing on the road—therefore, there is a higher risk of an accident occurring. If they all wish to get to their destination most efficiently and decide to gather information about whether to take the shortcut, then they have a biased set of evidence. This may make redundant the point of even taking the shortcut because more

people may end up taking the shortcut than the normal route. So, whilst they direct their attention properly, the actual outcome for the group is negative. This is because attention in this case is not *well-distributed* for the desired end. Equally, a group of drivers may all unevenly distribute their attention to the advertisement; the outcome may not be a negative, for they distribute their attention at different times and each adjust to the mistakes of others. Their attention is well-distributed. This is termed by Gardiner as:

NON-SUMMATIVISM

‘A group can lack the [correct distribution of attention] even though each member’s attention is [properly directed]’. Conversely, ‘a group can [distribute attention properly] even when no member does’ (2022, p.9)

So, the attentional normativity of a group can be distinct from the attentional normativity of a group’s individual members. Upshot: this completes (ii) and (iii).

Now let us consider utility theory’s evaluation of attention. First, epistemic utility. Once again, consider the original case of the driver. Say she can direct her attention to two things: the advert or the traffic. Imagine also that she only has enough time to attend to one of the two features of her environment. With epistemic utility, we give more value to those options that are more important to the individual. Seeing as her goal is to evaluate whether or not to take the shortcut, and this is decided by the state of the traffic, she should give more weight to the features that accurately represent the traffic. So, provided she is good at assessing traffic and directs her attention to the traffic, she will acquire accurate credences about that decision. In turn, her expected epistemic utility will be better if she attends to the traffic.

Second, pragmatic utility. The driver faces a decision: to take the shortcut or to continue on her existing route. She will be better equipped to make that decision if she has gathered information about the traffic ahead of her. The advertisement exists independently from that decision. From the perspective of pragmatic utility, she will do better in expectation if she pays attention to the traffic because it is more in line with her goal.

So, the way a subject directs attention can be seen as proper or improper with utility. It is also important to note that attention is manipulable. And precisely because it is formative to credences, manipulation of attention can be used to produce credences in individuals that a manipulator (e.g., individuals, corporations, institutions, etc.) wants to inculcate within them. For instance, the corporation that set up and designed the advertisement on the side of the road aims to capture the driver’s attention, not so they are distracted from their driving (though this may be an *unintended side-effect*), but to form a credence within them about the advertised product. Optimistically, this credence may later cause the individual to buy the product.

This is not to say that the individual is not compliant in this manipulation. If anything, their compliance is very much part of the reason companies invest in advertising (they would not buy a product unless they complied). Individuals

are compliant in directing their attention to manipulatory information, and how they direct it can be said, in part, to be of their own will. So, they have agency over how they direct their attention, and thus have responsibility to direct their attention properly.

So, what we have thus far is a clear evaluative framework for attention as a non-ideal feature of epistemic life in relation to credal epistemology. Attention: directs us to sources of information; helps form our credences; exhibits patterns and manifests in habits; can be evaluated by the utility of the credences they form; is (in part) volitional; and, finally, is manipulable. The credences that attention helps to form can be evaluated with epistemic and pragmatic utility. And, when attention is directed at different environmental information, we can compare the expected outcomes of two different attention-driven evidence-gathering episodes in relation to the same goal.

4 Exploiting Attention

In light of this definition and evaluative framework, we will now look at the more maligned exploitation of attention that capitalises upon the manipulable dimension of attention. To do so, we will focus on the case of humans directing their attention to social media.

4.1 Attentional Futility

The specific kind of attention direction we will be looking for is *attentional futility*: the use of attention to find something to pay attention to which in itself has little utility. There are two separable components of this definition. Firstly, there is the use of attention. Attention is used to achieve some end; it serves a purpose to some goal. The thing with attentional futility is that the individual's goal is to simply be in another attentional state. This is not necessarily bad. Say I am paying attention to a syllabus to find a specific topic to then pay attention to so I can revise an enjoyable topic: fine. There is a light at the end of the attentional tunnel; the attentional state that I am using attention to acquire has utility. My point is that the goal within a state of attentional futility is to be in another state of attentional futility. So, an individual in attentional futility *seeks further attentional states that have the same "purpose"*.

Secondly, it should be clear from the previous section that attention can be thought of as forming credences that have utility. The value of attention can be reduced to the evidence it acquires. With attentional futility, because the evidence is not for a specific proposition, the value of the attention will be wasted. As before, this is not necessarily harmful. In fact, many claim that much of our information acquisition occurs passively (Sosa 2017). However, what makes it bad is what attention is directed to: *hostile attentional scaffolds*. Henceforth, talk of attentional futility is attentional futility directed to these platforms.

So, let us see how this works.

4.2 Addiction to Futility

In singular instances, attentional futility is no real issue—the value of people’s credences is, after all, determined by the individual’s preferences. Nevertheless, we can still evaluate singular instances as negative. However, further harm is incurred when attentional futility is prolonged and habituated. Before delving into this, we will look at how the way social media platforms capture your attention does not promote utility and undermines their users’ goals.

To help us better grasp the the platform’s aim in this scenario, let us start with understanding the people directing their attention: they are agents; have credences; and are *users*—they are using (and hence attending to) some digital platform. This involves consuming content, and content is tailored to capture you attention (Nguyen 2021). Content is pushed this way because of the *attention economy*, whereby the majority of digital platforms exist because they attract consumers’ attention (Wu 2017)⁶. So, it is the “logic” of the attention economy to capture users’ attention and for users to direct their attention towards the platform they are using.

The specific reason social media platforms aim to capture attention is because they rely on personalised advertisements to generate profit, and personal advertisements require information about the user (Zuboff 2019). As such, gathering more data equals more personalised ads, and more data is acquired by maximising the amount of time users spend on their application. To achieve this, algorithms are designed to attract users’ attention and keep them attending. In other words, keep them engaged, where engagement is interpreted as a proxy for high-level goals, scored by the ‘moment-to-moment actions of the user’ (Narayanan 2023, p.18). High-level goals are simply goals that an individual wishes to achieve. In the case of the driver: getting to her destination as efficiently as possible.

We can outline some preliminary ways social media platforms are designed specifically to capture attention. One is through platform’s feedback to user interaction to stimulate sensorimotor engagement, such as pull-to-refresh and infinite scroll. Another: algorithms, such as recommender systems that suggest content to be consumed, can provide increasingly radical content to keep the user watching (Alfano et al. 2021). In one sense, these features can act to manipulate users’ goals and beliefs away from those states that benefit them (Burr et al. 2018). So, keeping users engaged is necessary for social media companies to profit. And we have reason to think that attention is fundamental to engagement. For social media platforms, this is brilliant: all they need to do is get people to attend to their platform and, in turn, hijack a users’ goals by grabbing their attention (which is even better for business! For a user whose attention is hijacked has less attention to direct elsewhere).

⁶While it is my aim to evaluate only the way content is presented and not the content itself, evaluating the content itself may be a direction of further discussion regarding attentional futility.

So far, I have highlighted two parts of social media design in relation to attention. First, while users direct their attention, the design of the digital platform is acting to keep you directing your attention. Second, companies direct attention to their platform by employing techniques that prioritise user engagement over user benefits; i.e., they potentially undermine user’s attentional states in favour of profit. In this scenario, the company does not consider users’ goals, and only aims to make them direct their attention to their platform. Essentially, they do not consider what is ‘best’ for a user, only what achieves engagement.

But we need to say more about the interactions users have with these platforms. To start, we can understand social media platforms as *attentional scaffolds*, which draws from the extended cognition thesis that cognitive processes and mental states may extend to artifacts external to the subject (Clark & Chalmers 1998). Artifacts or scaffolds trigger cognitive states. Specifically, attentional scaffolds are ‘designed external structures, artifacts or processes that systematically influence users’ focus, management, and direction of attention, either benefitting or harming their ability to achieve their goals’ (Voinea et al. 2024, p.686).

Attentional scaffolds can be either beneficial or detrimental. Conveniently, we can characterise benefit and detriment in terms of utility. So, a beneficial attentional scaffold would be one where attention is directed towards information that influences credences to have greater epistemic or pragmatic utility in line with the user’s goals and preferences. For instance, meditation apps may help users’ direct their attention more properly and improve mental health (Flett et al. 2018, cited in Voinea et al. 2024, p.687). Similarly, emergency and government notifications that signal emergencies draw users’ attention to information about potential danger, thus acting in the users’ interest of survival. Why are these valuable? Credal epistemology and utility say: because they provide more utility in expectation regarding an outcome desired by a user.

Detrimental attentional scaffolds can simply be anything that draws your attention to something that serves no utility, such as an alarm clock that rings at the wrong time (*ibid.*). However, detrimental attentional scaffolds become hostile when they ‘instrumentalize users’ emotions and biases to capture their attention and to steer it towards goals that were not deliberately chosen by users and which do not serve their interests’ (*ibid.*). So, a hostile attentional scaffold has *beneficiaries* and *victims*, where the beneficiary undermines a victim’s cognitive resources for their own benefit. Similarly, we can evaluate hostile attentional scaffolds by how they affect users’ utilities. Consider the design of social media platforms’ recommender systems: the recommended content is independently optimised by predicting whether a specific user is likely to engage with it (Narayanan 2023). The algorithm presents the immediate possibility of a *reward* to the user. Reward is emotional response to the content, and social validation through ‘likes’ and ‘shares’. The reward-based structure essentially *hijacks* attention, directing it towards the next possible reward, which is the next piece of content. This content can be understood as evidence, but evidence not related to any proposition desired by the user. Because the immediate reward

structure of social media platforms only anticipate the next instantiation of attention with the exclusive goal of retaining engagement, the attention spent on the content of social media platforms is outweighed by the costs in cognitive capacities and time that are incurred when attention is directed at these platforms. This is the result of users' attention being spent without a personal goal. The *reason* attention is ruled out from utility is because it has failed to contribute to credences useful for the user. The system undermines users' direction of attention. In other words, their attention direction is futile.

5 No Utility for Attentional Futility

Above shows how attention can be spent without purpose. However, to refine attentional futility, we need to know why this is harmful.

First, let us frame attentional futility in opportunity costs terms. Attentional futility is restricted to the specific goal of being in another state of attentional futility. Fundamentally, this has little utility and uses (among others) the resources of attention and time. As such, the more time an individual spends in a state of attentional futility, the less resources they have to spend on propositions more important to the user. In other words, if there is something else to be doing other than being in a state of attentional futility, then you ought to do it.

What is the epistemic problem of attentional futility? The way hostile attentional scaffolds make you pay attention limits resources available for other tasks. Let us evaluate how.

5.1 Individuals

Take a paradigm case. 'Doomscrolling' is a media habit (Ghersetti & Westlund 2016) where initially goal-driven behaviour is directed into 'repetitive, automatic browsing cued by environmental factors' (Sharma et al. 2022, p.2). Whilst the user may have a goal in mind when opening the platform, she need not continue to have this goal and scrolling is likely to continue when the initial goal has been forgotten or satisfied because of the aforementioned hostile design underlying platforms. The time users spend on the platform beyond the initial goal can be seen as passive information acquisition.

First, passive information acquisition from social media is not useful. What are users attending to? Content: fed to them so they engage. What do users make of this content? By both credal and empirical metrics: not much. The overarching credal point is that attention to doomscrolling is wasted because there is more utility in other acts. The negativity of this comes from the opportunity for an agent to direct their attention to other activities that have higher utility.

Say an agent has two options: doomscrolling or doing nothing. Given limited time and attention, she can only do one. The epistemic utility of each option will be the epistemic utility of the credences she would have after directing her attention to doomscrolling or nothing. In this specific case, she would

always do better by doing nothing. This is because the utility to be gained from spending your time doomscrolling is always outweighed by the cost-as-attention-directed—so the attention and time is always *wasted*. Moreover, because the immediate reward process as defined earlier can act to recommend more radical content, engaging with the platform increases the likelihood the user incurs undesirable epistemic harm (e.g., by being recommended videos pushing extreme and polarising political opinions). On the other hand, doing nothing only uses time and there is no increased possibility for epistemic harm. So, because measures of epistemic utility often give weight to those propositions that are more important to the individual, and it is in the users’ interests to avoid epistemic harms, doing nothing is better in expectation than doomscrolling.

This case highlights an important point: the engagement of resources reduces agency. The idea is this: less attention means you have less resources to direct to other propositions in your environment. Therefore, the user—whose attention is rapt up engaging with the hostile attentional scaffold—has less resources to go and consider other propositions in their environment. So, whilst doing nothing leaves all these resources free, doomscrolling does not. Aggregate this with algorithms, recommender systems, etc., and the user has less agency to decide how they direct their attention and what to. This helps the platform and harms the user.

To further push the point against doomscrolling, consider a user who does not see social media as a hostile attentional scaffold because she, say, only follows and engages with users and content creators she trusts; hence, she trusts the information on her feed and sees doomscrolling as an informative and useful task. Regarding how this goes wrong, consider the Gettier case of the *fake barn* raised by Goldman (1976): imagine someone driving through the countryside identifying barns in the field. Beyond her knowledge, some of these barns are facades. So, even if she is justified to believe she identifies a barn and her belief is true (from her perspective), she does not have knowledge because she only gets lucky when she forms the belief and drives past a real barn. Analogously, the user may scroll past many posts which may indeed improve her credences in expectation. However, beyond her perspective, some of the information is false; nevertheless, she takes it on board just as if it were true. So, as the content’s truth value is difficult to distinguish, doomscrolling negatively impacts the expected utility of her credences. At the level of the user, her attention is being directed to the content and the next reward. Just because the attention directed gives the illusion of utility, it does not vindicate her from attentional futility. This typifies a case where a user should not be gathering information.

5.2 Habits

The next step is to look at how attentional futility manifests itself negatively in habits. From the previous paragraphs we can say this much: attentional futility is epistemically costly and misleading. If you partake consistently in an epistemically costly and misleading act, then you have less time to spend doing epistemically fruitful and accurate acts. Simply, this is the idea of opportu-

nity costs outlined earlier. Because habituation leads you to spend more time (hence attention) on something, habituating a state of attentional futility minimises your utility in expectation. As such, you should not habituate attentional futility—a time-consuming and harmful improper direction of attention.

Equally, we may evaluate habits of directing attention by asking the utilities of the credences you end up with. Say there is a university student with a goal of doing as well as she can in her degree. She can develop the habit of either spending her spare time doomscrolling or reading beyond the set syllabus. The case for both epistemic and pragmatic utility is clear. As she likely gives greater weight to doing well at university than the goals of attentional futility, she ought to give greater weight to the information concerning her university degree than the content of the social media platform. So, providing she understands and considers the extra reading properly, she will get accurate credences the more time she spends reading beyond the syllabus. If she creates a habit of this, then she will get increasingly accurate credences regarding her syllabus. Therefore, she will do better in expectation if she creates this habit and ought to do this. Conversely, creating a habit of doomscrolling (and hence attentional futility) will serve no goal that makes her do better in expectation. Hence, she should not doomscroll.

With pragmatic utility, say she faces a decision (e.g., an exam) where she will be much better equipped by knowing more beyond the syllabus than the content on social media. The content on social media will likely be unrelated to her goal, but grant that some of it is. Either way, she will still be better off in expectation if she develops the habit of reading beyond the syllabus. For both the way the platform presents content and her state means doomscrolling fails to serve in her interests and her attention would be better spent attending to the syllabus.

5.3 Groups

Finally, let us consider attentional futility and groups. Following our conception of non-summativism, there are four notable options: (1) some members engage in attentional futility—the group’s outcome is bad; (2) some members engage in attentional futility—the group’s outcome is good; (3) all members engage in attentional futility—the group’s outcome is bad; (4) all members do not engage in attentional futility—the group’s outcome is good.

Both (1) and (2) can be addressed at once. First, take the group’s credences to be the total credences of its constituent members, which is updated as they gather evidence. Assume that each member can direct their attention to one thing within a situation. In light of attentional futility, members who focus on, say, doomscrolling will not grow the utility of their credences. For instance, assume every member starts with the same priors. Those in a state of attentional futility are left with their priors; those who direct their attention elsewhere have gathered information for some goal and thus have more useful credences in expectation. In turn, though the group can perform well or poorly (as non-summativism made clear earlier), this is independent of the individuals who

spend their time with attentional futility who contribute only their priors (which in this case are just the same as everyone else's). Essentially, while they could have spent time gathering evidence that improves the groups' credences overall, they did not. In opportunity costs terms: they did not maximise the potential of their resources and thus limited the opportunities of the group as a whole.

I think it is already clear why the outcomes of (3) and (4) have their respective normative attribution from the prior discussion on both individual and habitual cases of attentional futility. Nevertheless, there is an important outcome that synthesises habits and group-level opportunity costs evaluation, specifically regarding (3). Take a step back and consider that an individual's credences after an evidence-gathering episode will be the priors for some future evidence-gathering episode. The habituation of attentional futility will mean that significant opportunity costs are incurred. Okay—we've already said this.

But what about on a group level? Consider a group of students who have, through a habituation of attentional futility, incurred a great deal of opportunity costs so their priors at this point are significantly worse off than if they had never habituated it. Simply, they will be less prepared than if they had instead directed their attention to other propositions. This is one part of the problem. The other part is that credences determine your preferences. Recall the mushroom-pepper example: if I learned one was spoilt, I would choose the other. In the case of hostile attentional scaffolds also: if I engage with one form of content, the algorithm suggests more of that content. If I consistently learn that mushrooms are spoilt, I will just start picking only peppers. Analogously, if I am consistently shown one form of content, I will likely continue to engage with that form of content. The problem I am alluding to here is:

FEEDBACK LOOPS

A self-reinforcing dynamic where an input results in a response that intensifies the original cause.

(See Burr et al. 2018)

This may act to destabilise the user individually by reinforcing inaccurate credences a user has by recommending more content that ratifies those credences in the interest of engagement⁷.

On the group level, feedback loops may shift standards and values of behaviour. For instance, if a group of users all engage with one specific type of content then their evidence will be biased. Aggregate this and the group has a less diverse body of evidence and biased credences. Tentatively, this aggregation of biased credences may shift a groups' goals and preferences over time, undermining their original interests to those that make them engage the most. Once again, it seems their attention can be hijacked by the attention economy⁸.

So, the potential epistemic harms of attentional futility to hostile attentional scaffolds are clear. Individual's and group's goals can be undermined by it

⁷This phenomena has been studied as 'filter bubbles' Pariser (2011).

⁸There are clear socio-political ramifications that may arise here.

and, in opportunity costs terms, it is always better to direct your attention at something else. In slogan form: there is no utility for attentional futility.

6 Conclusion

The primary point I have been pushing here is that there are serious epistemic ramifications for individuals and groups who have the purpose of attention hijacked for the interests of a hostile attentional scaffold. With credal epistemology and utility theory, I have shown how the direction of attention can be evaluated in relation to opportunity costs, with specific ways of directing attention either serving or undermining an agent's goals. By developing and evaluating the novel concept of attentional futility in relation to hostile attentional scaffolds, I have made clear how attention can be undermined and why this is negative in light of epistemic and pragmatic utility, and opportunity costs. Specifically, attention can be hijacked to serve the goals of the attention economy instead of the individual who is directing their attention.

In total, the key takeaway here is that attentional futility is a harmful epistemic state to be in. However, it should be noted that the harms are not limited to the epistemic. There are clear political and social harms, such as polarisation and value manipulation, that can be brought about by the grabbing and undermining of users' attention with hostile attentional scaffolds. Much more can be said about attentional futility, and I have contributed to this literature by developing a framework in credal epistemology to evaluate both attention and attentional futility.

References

- Alfano, M., Fard, A., Carter, J., Clutton, P. & Klein, C. (2021), 'Technologically scaffolded atypical cognition: the case of youtube's recommender system', *Synthese* **199**(1-2), 835–858.
- Begby, E. (2023), 'Opportunity costs and resource allocation problems: Epistemology for finite minds', *Canadian Journal of Philosophy* **53**(5), 400–417.
- Buehler, D. (2018), 'A dilemma for ?selection-for-action?', *Thought: A Journal of Philosophy* **7**(2), 139–149.
- Burr, C., Cristianini, N. & Ladyman, J. (2018), 'An analysis of the interaction between intelligent software agents and human users', *Minds and Machines* **28**(4), 735–774.
URL: <https://doi.org/10.1007/s11023-018-9479-0>
- Clark, A. & Chalmers, D. (1998), 'The extended mind', *Analysis* **58**(1), 7–19.
URL: <https://doi.org/10.1093/analys/58.1.7>

- Gardiner, G. (2022), Attunement: On the cognitive virtues of attention, *in* M. Alfano, J. D. Ridder & C. Klein, eds, ‘Social Virtue Epistemology’, Routledge.
- Ghersetti, M. & Westlund, O. (2016), ‘Habits and generational media use’, *Journalism Studies* **19**(7), 1039–1058.
- Goldman, A. I. (1976), ‘Discrimination and perceptual knowledge’, *The Journal of Philosophy* **73**(20), 771–791.
- Good, I. J. (1967), ‘On the principle of total evidence’, *The British Journal for the Philosophy of Science* **17**(4), 319–321.
- Narayanan, A. (2023), ‘Understanding social media recommendation algorithms’, 23-01 Knight First Amendment Institute.
URL: <https://knightcolumbia.org/content/understanding-social-media-recommendation-algorithms>
- Nguyen, C. (2021), How twitter gamifies communication, *in* J. Lackey, ed., ‘Applied Epistemology’, Oxford University Press.
- Pariser, E. (2011), *The Filter Bubble: What The Internet Is Hiding From You*, Penguin UK.
- Pettigrew, R. (2024), The value of information and the epistemology of inquiry.
- Sharma, B., Lee, S. S. & Johnson, B. K. (2022), ‘The dark at the end of the tunnel: Doomscrolling on social media newsfeeds’, *Technology, Mind, and Behavior* **3**(1: Spring 2022).
- Sosa, E. (2017), *Epistemology*, Princeton University Press.
- Voinea, C., Marin, L. & Vică, C. (2024), ‘Digital slot machines: Social media platforms as attentional scaffolds’, *Topoi* **43**, 685–695.
URL: <https://doi.org/10.1007/s11245-024-10031-0>
- White, D. J. (2024), ‘Paying attention to attention: Psychological realism and the attention economy’, *Synthese* **203**(2), 1–22.
- Wu, T. (2017), *The Attention Merchants: The Epic Scramble to Get Inside Our Heads*, Vintage.
- Wu, W. (2011), Attention as selection for action, *in* C. Mole, D. Smithies & W. Wu, eds, ‘Attention: Philosophical and Psychological Essays’, Oxford University Press, pp. 97–116.
- Zuboff, S. (2019), *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, PublicAffairs, New York.