Hostile Epistemology

C. Thi Nguyen

I feel like the motivating cry for a lot of people in this room is: what the hell is going on right now? The world is suddenly full of conspiracy theorists, climate change denialists, anti-vaxxers, and white supremacists. In grappling with all this, some have taken an approach which has come to be known as *vice epistemology*. The misinformation crisis, according to this approach, can be substantially explained by thinking about people's *bad epistemic character*. As Quassim Cassam puts it, "epistemic vices are defined as character traits, attitudes or thinking styles that systematically obstruct the gaining, keeping or sharing of knowledge" (2018). Take, for instance, Cassam's analysis of Oliver the conspiracy theorist, who is a 9/11 truther. Oliver believes that 9/11 was an inside job — that it wasn't the result of a plane crash, but a controlled demolition by secret government agents. Oliver, says Cassam, isn't reasonable. His beliefs have been "conclusively refuted" by "reputable studies". He "should have known better but doesn't" (2016, 162). The best explanation for Oliver's beliefs is bad intellectual character. He is "gullible, cynical, and prejudiced" (2016, 163).

I'm worried about vice epistemology and what it tends to highlight. It's not that I think there's no such thing as epistemic vice. It surely exists, and is worth studying. But I worry that vice epistemology, when overemphasized, can encourage us to reach for individualist explanations first. It centers the epistemic failings and defects of the agent. And its fixes usually involve changing individuals: educating them better, bringing them up to be more intellectually virtuous. Some newer approaches have broadened the picture to include the possibility of group beliefs and group vices — but they still center some kind of epistemic agent. The vice approach focuses on the thing which believes, asks how it's supposed to behave, and then wonders why it's failed.

I would like to present an alternative framework, which I'll call *hostile episte-*

mology. Hostile epistemology is the study of the ways in which *environmental features exploit our cognitive vulnerabilities*.1 I am particularly interested in when those vulnerabilities are unavoidable — when they arise from the essential condition of our epistemic lives. We are essentially finite beings, with limited cognitive resources. We are perpetually reasoning in a rush, because there is far too more information than we could ever fully process. Our desperate attempts to cope with a cognitively overwhelming world will inevitably leave holes in our armor. And the world can take advantage of those vulnerabilities.

I often hear people saying things like, "We're in a post-truth world now; people just don't *care* about the truth anymore." I would like to explore another possibility: that often, people do care about the truth, but their good-faith efforts are being subverted by features of the world. And it can be easy to mistake the results of such exploitation for character flaws. It is easy to think that the victims of exploitation are just lazy or lack sufficient grit. Feminist work has been particularly clear on our tendency to misattribute the results of structural oppression to individual character flaws. Marilyn Frye (1983) argues that women's anger can be an entirely appropriate response to a systematically unjust world but it is often dismissed as a character flaw. We fault the angry person, we label them "oversensitive" or "hysterical".2 My worries about vice epistemology come in a similar key. I am worried when we reach a bit too readily for explanations in terms of vice and character flaws without considering the alternatives. My goal here is not to criticize vice epistemology so much as it is to fill out a counterbalancing approach.

There is a tendency for philosophers, I think, to seek a kind of epistemic invulnerability. We want a procedure that guarantees us the right belief. And we often act as if human beings could plausibly reach such invulnerability. At least, we condemn those who fail to achieve it, or regard them as defective. But such epistemic invulnerability is out of our reach. We live on far shakier ground. We are limited beings, forced to cope with a vast world far beyond our individual ability to understand. We are forced to take insecure methods. We must take cognitive shortcuts, use heuristics, and engage in quick and dirty reasoning. And we must trust one another other if we are to have any hope of making sense of the world. This is no way to do that from a place of perfect security. Our cognitive condition

---

1. I first introduced "hostile epistemology" in "The Seductions of Clarity" (Nguyen, 2021). The current presentation is in the same spirit, but (hopefully) improved in many small details. I trace the term "hostile epistemic environment" to Heather Battaly (2018), though our uses of the term differ slightly. She raises the term in a discussion of the epistemic vice of close-mindedness — where she argues that closed-mindedness is occasionally a virtue, in epistemically polluted environments.

2. In this spirit, Sue Campbell (1994) offers an analysis of bitterness. To call somebody "bitter", she says, is to dismiss their complaints. Instead of following their gaze to the world — to look at how that world might have wronged them — we direct our critical gaze at the person; we accuse them of a character flaw

forces us to overextend ourselves. This inevitably makes us vulnerable and exploitable. We can adapt to fight those exploitations, but there are no perfect adaptations. Sometimes, we are simply relocating our vulnerability for the battle at hand.

I use the term "defective" here, which I find quite uncomfortable, with some deliberateness. Cassam defines epistemic vice as those epistemic failings which arise from our nature, for which we might reasonably be held responsible for. There are, he grants, other failings of our nature we might not be held reasonably responsible for. These he labels "cognitive defects". They are, for him, something like blindness. A defect is a condition which obstruct our cognitive grasp, but which we aren't responsible for (Cassam 2019, 21). It is this basic choice space — vice or defect — that I want to resist. Much of our exploitability arises, not because we are failures or defective, but because we are doing the best we can with what we have.

Hostile epistemology

My goal here is to introduce a framework for hostile epistemology. Let me be honest: what follows is big and hand-wavey. I have come to realize that many of the different things I've been writing about are actually, at their root, about the exploitation of our essential cognitive vulnerability. Today I'm going to try and weave these pieces into a larger picture. This will involve a lot of impressionistic brushwork. I suspect that many of the particulars here will be familiar to many of my readers, but perhaps they might find something useful in the assemblage.

Let me offer a fuller delineation of the term. Hostile epistemology is the study of the ways in which environmental features can exploit our cognitive vulnerabilities and weaknesses. I mean "environmental features" quite broadly: it can include other people, communities, cultural practices, institutional structures and technologies. I will focus here on intentional and near-intentional forms of hostility, but hostile epistemology can also include non-intentional hostility. There can be natural hostility and accidental hostility. A hostile environment can be made intentionally — like a minefield. A hostile environment can also be the product of natural evolution. A field of wild thorny bushes is a hostile environment for humans; its hostility did not arise from any conscious intent, though there is an evolutionary mechanism which explains the thorns' defensive functionality.3 And

---

3. I am leaning here on Ruth Millikan's (1987) account of how biological traits can have be said to have a function through an unguided, non-intentional process of natural evolution — when the traits are reproduced because they are casually efficious at fulfilling that function. As A.W. Eaton (2020) draws our attention to how this line of argument can be used to show how cultural artifacts can have a

sometimes, hostility can be the result of a pure mismatch. Mars is a deeply hostile environment for humans, but not because anything intended or evolved to be that way. Hostility is also relative to our capacities. The deep ocean is incredibly hostile to humans, but friendly to anglerfish. And hostility is relative to our purposes. The French bakery near my house makes truly extraordinary almond croissants. This is conducive to my culinary interests, but it is hostile to my efforts to eat less sugar.

I think the not-quite-intentional forms of hostility are particularly important right now. Think, for example, about corporate makers of junk food. They make food that exploits our vulnerabilities; they create hostile nutritional environments. But they didn't necessarily intend to harm us or to make food that exploits our vulnerabilities. A corporation might have gotten there by simply pursuing profit. They might have tried different ingredients and approaches in food and went with whatever sold best, modifying their recipe through trial and error. Their interest in profit has steered them, unintentionally, into producing food which exploits our vulnerability to, say, massive infusions of sugar and salt.

When I talk about *epistemically* hostile environmental features, I'm thinking of things like: seductively clear belief systems like conspiracy theories; communities structured as echo chambers; gamified user interface design on social media; fake institutional credentialing systems. Notice: there is a way to read these examples in terms of epistemic vice. The person who loves conspiracy theories is cynical; the person who falls for fake institutional credentials is gullible; the person who gets taken in by an echo chamber is weak-minded. But there is also a hostile epistemology take. How might these false beliefs have arisen from the exploitation of our vulnerabilities? Here's one possibility: the world forces us to use shortcuts to estimate expertise, because we don't have the capacity to figure out which experts to trust for ourselves. So we must use rough signals to figure who to trust. Often, we look for signs of officialness: medical degrees from recognized institutions, scientists at universities we've heard of. But these signs can be readily faked to nonexperts. Using such signs is a fallible heuristic; it creates a vulnerability. But we are forced to use fallible heuristics, because we don't have the capacity to judge every kind of expert for ourselves.

Some of our vulnerabilities might be our fault, the result of carelessness, laziness or other vices. But I am particularly interested in thinking about our *unavoidable* cognitive vulnerabilities –the class of vulnerabilities that we inherit simply in virtue of being limited beings, forced to cope with a world too large for our brains. Limited beings must take shortcuts. Shortcuts are, by their nature, imperfect; taking them makes us vulnerable. And we might come to see the problems with one

---

function even if it was never intentionally made to fulfill that function — if the artifact is reproduced because of it is casually efficacious at fulfilling that function. Thus social traditions and institutions may perpetuate because they fulfill a function, even if none of the people involved think of those traditions or institutions as fulfilling a function.

shortcut, but we can, at best, replace it with a different shortcut. (The image that comes to mind is of a fighter using a shield: whenever they move it to protect one vulnerability, they expose another.)

Feminists, critical race theorists, Marxists, disability theorists, and many others have emphasized structural and systemic explanations for apparent character flaws, and I have found these approaches enormously useful.[4] But these approaches typically focus on the structural dynamics of oppression. And their analysis is usually aimed at exposing *contingent power dynamics*. Like: these people seized the capital first and used it to oppress these other people. People who look like this had sufficient power to control the media image of beauty, and biased it against people who look like that. These people created the urban environment to suit their bodies, and so rendered other people, with different bodies, apparently "disabled". The interest here is usually in how unjust social structures create vulnerabilities. As Carla Bagnoli puts it, in this sort of ethical approach to vulnerability, we focus on those agents whose vulnerability comes from social discrimination or oppression, and other contingent historical and social factors. The ethical approach focuses on this sort of contingent vulnerability, says Bagnoli, because the goal is moral and political redress (2016, 14-16).

But I want to explore here something a little different: the cognitive vulnerabilities that are an inescapable part of the human condition.[5] I want to find out which of our exploitable vulnerabilities are the result of a particular history of injustice and oppression, and which arise from our basic nature — especially our nature as finite beings. Of course in the actual world, these will be deeply entangled. All of us have limited cognitive resources, but unjust social dynamics will leave some of us with more cognitive resources, like time, and others with far less.[6] But still: I think it will be useful to work out what forms of vulnerability are inescapable for any cognitively finite being, if only to articulate the theoretical outer limit of our epistemic security.

---

4. The literature here is vast, but some key moments are Collins (1990), Young (2002), Fricker (2007), Pohlhaus (2012), Medina (2013), Dotson (2014), and Hendren (2020). Particularly useful recent papers which provide overviews of the relevant literature include Toole (2019), Pohlhaus (2020), and Sertler (2022).

5. Bagnoli (2016) is also interested in focusing on the inescapable vulnerabilities, and offers a careful taxonomy of them. I find also taxonomy rich and provocative, but it also brings on board certain (likely Kantian) conceptions of agency and agential constitution that I am not willing to accept. Briefly, she takes vulnerability to be a feature of human *embodiment* (16), while I think, there some forms of vulnerability arise from embodiment, others do not — they arise from cognitive finitude alone. So I do not adopt her taxonomy here, though our general approaches are very sympathetic. (Though her interest is in the moral implications of cognitive finitude, where mine is in the epistemic implications.)

6. I owe this point to Ezgi Sertler.

But this is all becoming impossibly abstruse; we should dive into the particulars. From here, my goal is to summarize some recent work I've done and then try to draw some larger lessons for hostile epistemology. There are, very roughly, two categories of vulnerabilities I've been thinking about: ones associated with *heuristics*, and ones associated with *trust.* A reminder: this is big picture work, and I will be roughly sketching arguments and conclusions that I and others have worked out in more detail elsewhere.

Heuristics and seductive clarity

Many of the seductive belief systems we're looking at — like conspiracy theories — seem to work by offering their believers an intense dose of the feeling of clarity. When you accept a conspiracy theory, suddenly all these random parts of the world fit into one great overarching theory. What was once chaotic becomes navigable, controllable, and sensical. It feels like you suddenly understand, that the world has become comprehensible and coherent.

I've tried to give an account of the seductive power of clarity.[7] By "clarity" here, I mean the *feeling* associated with understanding, which is not the same as genuine understanding. To genuinely understand is to properly grasp some structure in the world. The feeling of understanding is somewhat correlated with genuine understanding — just as the feeling of getting an intense workout is somewhat correlated with actually getting a useful workout. Such a feeling can be used as a heuristic for the real thing ("Feel the burn!"). But the feeling and the real thing can come apart. For example, the constant pursuit of the feeling of workout intensity can actually be counterproductive. The best path to physical fitness and health is from a mixture of low, moderate, and high intensity work — and plenty of rest. This can feel, to those hyperfocused to the feeling of intensity, like you're not doing enough.[8] But, still, "feel the burn" is still a decent heuristic, and it can actually take you pretty far. Similarly for cognition: as long as there is a moderate correlation between the feeling of clarity and actual understanding, then we can reasonably use the feeling of clarity as a heuristic guide. But there still is a gap between the feeling and the real thing, and that gap can be exploited.[9]

---

7. What follows is a rough summary of "The Seductions of Clarity" (Nguyen, 2021), with some added implications for hostile epistemology.

8. You can look at pretty much any scientifically informed fitness and strength training manual for this, but I learned it from Rippetoe (2005).

9. I have a private theory that Crossfit has gamed the "Feel the burn!" heuristic, and owes some of its popularity to optimizing for the feeling of intensity, rather than for actual fitness effect. But I don't want to admit this in public because, I must admit, the average Crossfitter can probably beat me up

Let's take a big step back. There is too much world. We can't investigate everything in the world; we can't even investigate everything directly relevant to our lives. Here are some things very relevant to my life that I failed to investigate this week: the reliability of the new brand of gummy vitamins I bought my kids, the professional background of the new doctor I saw at the walk-in clinic, whether the updated spreadsheet on my university's online grading portal is still calculating grades correctly.

I skipped these investigations because I am in a perpetual cognitive rush. But rushing is not necessarily a sign of carelessness; some degree of rush will always be forced on us by our cognitive finitude. I particularly like the philosopher of science William Wimsatt's (2007) articulation of this point. Wimsatt pushes back against the traditional methodology of epistemology, which has focused on how ideal, cognitively limitless beings should think. But, says Wimsatt, limited beings must reason in a way that is different in kind from ideal, unlimited beings. We can't just imagine how God would think and then try to approximate that. We need to come up with tools and strategies built for limited beings who must mercilessly manage their cognitive procedure. They will, of necessity, be imperfect. We need to acknowledge that inevitable imperfection, and then manage it.

The world constantly assaults us with new topics for inquiry. An unlimited being could investigate it all; a limited being must decide what to investigate and what to ignore. But how should we do that? How do we decide where we should spend our precious cognitive resources? We can't investigate everything to completion before we decide how to allocate our cognitive resources. We need a heuristic for generating *rough estimates* of what's worth investigating and what's not. This would function as a *thought-terminating heuristic* — a rough guide to when we should look away and save our energy for something else.

A heuristic is a rule-of-thumb used by rational agents to reason and decide under conditions of limited cognitive resources. A heuristic trades away cognitive perfection in return for speed and ease of use. It is a cognitive shortcut — imperfect, but fast. One way that Wimsatt puts it: heuristics gain efficiency by presuming a particular context; they are fast in that context, but will fail if the context changes.10 A heuristic I've been using lately when buying booze is: the better the graphic design, the worse the beverage. Right now, there are lots and lots of new alcohol companies run by people with backgrounds in marketing. They tend to focus their efforts on advertising and graphic design, and not making delicious booze. The best booze is usually made by small independent companies run by people obsessed with the fine culinary details of the drink, and those folks usually

---

pretty easily — which may speak to the fact that the heuristic is, even if imperfect, still pretty effective.

10. Wimsatt (2007. *76-77*). Wimsatt here is drawing on a long tradition of work on heuristics from decision theory, political theory, and cognitive science; I find his articulations particularly helpful.

don't care much about graphic design. My graphic design heuristic turns out to be pretty good in the current environment. It's far quicker than doing meticulous research on every new booze brand. It is also deeply imperfect. For one thing, it generates false negatives. For years, I missed out on a gorgeous local rye, because the distiller's sibling happens to be an artist and made them lovely labels. And using the heuristic creates a vulnerability. In the case of my graphic design heuristic, if enough people start using my same heuristic, then marketing-savvy booze companies could respond by carefully designing defiantly plain labels. (In fact, that's probably already happening.)

Using a thought-terminating heuristic would create an enormous vulnerability. There is inevitably a gap between the heuristic and the full fact of the matter. If somebody could game that heuristic — if they could package a system of beliefs with the trigger for our thought-terminating heuristic — then they would have an incredible power over our beliefs. They could package a system of false belief with a signal that it wasn't worth investigating any more. If they could give us the *feeling* that we understood it, then we would end our investigations and turn our attentions elsewhere. This is the exploitation of a vulnerability which arises from the heuristic nature of our reasoning. And this sort of vulnerability is unavoidable precisely because we are forced to take shortcuts to solve a cognitive resource management problem.

I take a lot of inspiration here from contemporary work in practical reasoning.11 A common theme in the practical reasoning literature that limited beings need techniques to *settle the mind* — to cut off the inquiry process before it is perfectly complete. To settle your mind is to make a choice to avoid considering new evidence. Michael Bratman offers very useful accounts of how we use mind-settling techniques to manage our cognitive load.12 We need to make some decision about some future course of action. Let's say it's one of those rare nights when my spouse and I have been gifted some free childcare for the night. Should we go to a movie? A nice restaurant? Get stupidly drunk with our friends like we used to?

11. Traditionally, epistemologists have spent relatively little time worrying about the implications of cognitivel limitation, while practical reasoning theorists have worried about it obsessively. There is, however, an exciting new body of work in epistemology running against this trend, which often calls itself "non-ideal epistemology". Of particular use for the project at hand is the literature which has concerned itself with the interaction of strategies for limited cognition with biases, prejudices, and stereotypes. See Beeghly (2015), Johnson (2020), and Begby (2021b). Of particular relevance is Jessie Munton's (2019) discussion of visual perception. For cognitively limited beings, the visual system must fill in missing information using prior experiences. Such a limited being might be following an unimpeachable procedure, but still be biased due to biases present in their prior experience. Munton points out that such biased can be produced to through intentionally gerrymandered experiences – which I take to be an example of the hostile exploitation of our vulnerabilities.

12. The original statement of the account is (Bratman, 1987). My understanding of Bratman is heavily influenced by Richard Holton (2009), which offers very useful extension of a Bratman-style project.

We could endlessly deliberate about it, as new bits of information comes to mind, and spend the entire day texting about the different options. But we will quickly pass the point of diminishing returns on that deliberation, and start wasting our cognitive resources. So what we do, says Bratman, is to form an *intention*. An intention, for Bratman, is a plan or a policy which settles the mind. If we decide this morning that a movie sounds nice, we form an intention and so defeasibly close our mind to future deliberation. We decide to ignore most new relevant information. We might re-open deliberation in really extraordinary circumstances — like if we come down with a stomach flu or suddenly discover that we're bankrupt. But for run-of-the-mill considerations, we avoid restarting deliberation.

Settling the mind creates a kind of *cognitive inertia*. And limited rational being benefit greatly from the deliberate deployment of that cognitive inertia to manage their cognitive resources. You might have thought that refusing to consider new evidence was blatantly irrational. But for limited beings, the decision to close your mind can be an entirely reasonable response to cognitive load problems.

I'm suggesting that the thought-terminating heuristic is another strategy of cognitive resource management. But unlike Bratman-style intentions, a thought-terminating heuristic must often be deployed before we have thought much about a topic at all. Again: it might seem irrational to ignore to dismiss things without having carefully considered them, but limited beings actually need the capacity to ignore topics they haven't thought very much about. This capacity is rational, but also highly exploitable.

This has been, so far, an abstract philosophical argument to an abstract philosophical conclusion: that cognitively limited beings need a thought-terminating heuristic. There is nothing so far that tells us what that heuristic would actually be like. And we shouldn't expect an abstract philosophical argument to tell us which heuristics we actually use, since heuristics are temporary, tailored to specific contexts, and malleable. But we do have some empirical evidence about which heuristics are actually in use right now.

There is significant empirical evidence to suggest that many people use as their thought-terminating heuristic, the feeling of clarity — that is, the feeling associated with genuine understanding. What might that feel like? Here we can draw from recent philosophical work on the nature of understanding. Traditionally, epistemologists have seemed to think that the real goal of inquiry is *knowledge*, for which the paradigmatic example is some kind of secured true belief in an individual proposition. But to some, this traditional picture seems too piecemeal. It doesn't capture how inquiry is often directed at seeing the relationships between individual propositions. Philosophers of science and philosophers of education have recently argued that the real target of inquiry is something more holistic — that we want is not just knowledge, but *understanding*. Loosely, to understand is to possess a model of some part of the world which integrates various piecemeal bits of

knowledge into a connected whole.

The discussion here is complex and ongoing, but we can note a few generally accepted features of understanding (Elgin 2002; Kvanvig 2003; Grimm 2006, 2012; Khalifa 2012; Strevens 2013). When we understand something, we grasp a whole structure and not just individual nodes. We see how the parts hang together; we gain an internal model of some bit of reality. As John Kvanvig puts it, when we understand, we grasp a coherence relationship which grants us the ability to move between nodes, to see connections and to accommodate new information. There is also, as Catherine Elgin notes, a characteristic moment of *coming to understand*. When we gain understanding, our system of categories shifts to accommodate new information. And, as Michael Strevens says, when we really understand, we can transmit that understanding. We possess an internal model which we can pass on.

Let us suppose for the moment, that we are hostile epistemic manipulators. As I've said, hostile epistemology doesn't necessarily involve intentional manipulation. But it will be a useful starting point to think about what intentionally manipulative agents would do: the kinds of strategies they would take, and the kinds of structures they might employ. This will help us to recognize such other hostile structures with similar functional structures, even when they arise non-intentionally.

So suppose we are hostile manipulators, and we want to *simulate* understanding — to attach the feeling of clarity to some false belief system. How would we do it? We would want to maximize the feeling of all the pieces fitting together. We want to maximize the feeling of a system of categories suddenly shifting to accommodate all the information, the feeling of the ability to see connections and move between nodes. And this might seem like a mighty tall order. But we hostile epistemic manipulators have a significant advantage. Unlike our more epistemically sincere compatriots, we are not bound by the truth. We aren't trying to present a belief system that actually captures the world itself. Freed from the constraint of getting it right, we are free to optimize our system for the *feeling* of clarity.

I actually adapted the above argument from my mother. As a young child, I bugged her every day for McDonald's apple pies and Kentucky Fried Chicken because they were more delicious. My mother would respond: of course fast food is more delicious. Those people don't care about you, so don't need to make their food nutritious. They're free to just make it as yummy as possible. But *her* cooking had to balance nutritiousness and deliciousness, so it could never compete with the unqualified deliciousness of fast food. I am making the same suggestion about epistemic manipulators: freed from the constraints of the truth, they can optimize their epistemic systems for maximally deliciousness. They can make cognitive junk food.

And this is precisely what a conspiracy theory can do. A conspiracy theory can

seduce us with its clarity.13 A conspiracy theory gives you an internal model that fits a thousand disorganized pieces into a neat causal and conceptual scheme. We suddenly see a blooming whole. And we can pass on this intensely coherent vision. But conspiracy theories are just the most obvious example. Consider bureaucratic schemes of justification. Some universities have focused on a single narrow metric, like improving "student success", where student success is defined through a specific formula, like a weighted average of graduation rate, graduation speed, and post-graduate employment rate. Such an institutionalized justificatory scheme grants a sense of clarity. It grants a feeling of understanding of goals and justifications, and a vision of how all those goals intersect. And notice that here such metrics, plausibly, are not made from hostile intent. Their clarity and ease-of-use are often in service the need for large-scale information aggregation; we can manufacture such systems out of a simple interest in institutional efficiency (Porter, 1995; Bowker and Star, 1999). But such metrics also *function* to exploit our vulnerabilities, since such informational simplification and standardization makes it easier to assemble clear and coherent sweeping explanations.

If people are, in fact, using clarity as a thought-terminating heuristic, then we should expect to see people accepting clear-feeling belief systems without adequate investigation. And this is precisely what we do see. We see it with conspiracy theories, obviously. And the empirical work on bureaucracy, metrics, and quantification culture is full of examples of people and institutions latching onto some oversimplified metric, and letting it guide all kinds of policies and actions, without adequate investigation into the quality and accuracy of that metric.14

I've offered the account of clarity here just as an example, to paint a picture of what such heuristic exploitation might look like in the real world,. But for our current purposes, it doesn't really matter whether I've got the thought-terminating heuristic exactly right. What really matters is that we are the kind of beings who need a thought-terminating heuristic. And such a heuristic will inevitably create a vulnerability.

But how are we supposed to adapt, and how responsible are we for the exploitation of our heuristics? Wimsatt is helpful here. Says Wimsatt: cognitively limited beings cannot hope to get a perfect theory before we act. We need a system by which we can act on imperfect theories, and improve on them. We can't avoid imperfection or mistakes. What we need is a way to learn from our mistakes — to be able to use our mistakes as data to improve our rough theories. What we need

---

13. As many have pointed out, a conspiracy theory is not necessarily false (Coady, 2012; Dentith, 2018). After all, there have been many genuine conspiracies in the world. What I am suggesting here is that conspiracy theories can have extra cognitive seductiveness in virtue of their clarifying features, beyond what the evidence warrants.

14. For starters, see Porter (1995), Merry (2016), and Espeland and Sauder (2016). I offer an extended discussion of the literature in Nguyen (2022).

is the capacity for *error metabolism* (22-24). (Science has such error metabolism built into its bones, says Wimsatt. Traditional a priori philosophy, on the other hand, does not take our cognitive limitation in mind, and has no useful system of error metabolism.)

Let me accept Wimsatt's view. The rational limited being can't wait for perfect theories before they act. They must make rough stabs at heuristics, make errors, and then metabolize those errors. Here, then, is what I would do, if I were an epistemic manipulator. First, I would engineer a seductively clear belief system, to trigger the thought-terminating heuristic. That might get people to accept the belief system without further inquiry. But that is not enough. Reasonable people will still re-start inquiry if they discover evidence that their belief system is leading them into errors. So I need to conceal the evidence of such errors. How? Here's one strategy: I might embed, in that belief system, a further set of beliefs that preemptively dismissed other sources that might reveal those mistakes.15 I might construct an *echo chamber*, which dismisses as untrustworthy any outsiders who don't share our belief system.16

I could also make use of the pleasure of clarity. There is an intense pleasure in the "a-ha" moment, the moment of epiphany. As Alison Gopnick (1998) puts it, epiphanies are something like mental orgasms. So I could stack the pleasure of clarity with other cognitive pleasures. I've offered a few accounts of other seductive cognitive pleasures; let me briefly summarize them here. First, there is a seductive pleasure to being in a community of like-minded people. And that pleasure is amplified in an *echo chamber* — a social structure that automatically dismisses as untrustworthy anybody that all who disagree (Nguyen 2020b). Second: there is a pleasure to righteous moral outrage. And that pleasure can be exaggerated and simplified, to produce what Bekka Williams and I call *moral outrage porn*. Moral outrage porn is when we *use*, instrumentally, representations of moral outrage for pleasure, while avoiding the costs and consequences of actual moral action (Nguyen and Williams 2020). Third: we can increase the pleasures of our investigative activities through gamification — the adoption of a system that offers us points and game-like signs of victory, into a real-life activity. In gamification, we adopt a simple, pre-fabricated point system as our values; in exchange, we gain the joys of hyper-clear value. Consider Twitter. Twitter offers us a set of very clear points: Likes, Retweets, Follows. But those points don't capture all the potential values and goals in communication, like connection, understanding, empathy or information. Those measures just capture popularity. Twitter offers us a trade. If

15. See Endre Begby (2021) for an account of how a belief system may be structured to pre-emptively dismiss counter-evidence. Interestingly, Wimsatt (2007, 80-89) describes cases in science where a heuristic model makes certain simplifying assumptions, and those simplifying assumptions make relevant errors invisible.

16. Nguyen (2020b) contains my account of echo chambers as communities of structured distrust.

we take on its simple scoring system as our goal, then we get all the thrill of a game: clear and visible progress, inarguable success. But to do that, we need to give up on our more subtle and personal values, and take on the simple, pre-fabricated values built into the system (Nguyen 2021b; see also Nguyen 2020c).

This suggests an explanation for why conspiracy theories, echo chambers, moral outrage porn, and gamified social media are so often found together. They are all powerful tools for an epistemic manipulator willing to leverage cognative pleasure, so we should expect to find them deployed together for that purpose. It is something like why we find guns, knives, and camo together in the world. They have a functional symbiosis.

The vice epistemologist might be tempted to blame the believer here. Taking on such pleasurable beliefs systems, they might say, is a sign of a weak mind — one who prefers pleasure over truth. But the hostile epistemologist can offer a more complex view. Pleasure can, in many contexts, be a good heuristic. We seem to experience pleasure when we are thinking well — when we assemble the pieces, when we understand. We experience pleasure when we have clarified for ourselves our real goals and values. Pleasure is often correlated with cognitive success, so it isn't implausible to use pleasure as heuristic for cognitive success. But, like any other heuristic, it's exploitable.

Here's a rough analogy. Let me help myself to a simplified evolutionary psychology style explanation. In the pre-industrial nutritional landscape — as the evolutionary psychologists put it, in the environment in which we evolved — it makes perfect sense that we take pleasure in fat, salt and sugar, and less immediate pleasure in, say, leafy greens. The pre-industrial world was usually scarce in sources of ready fat and sugar and abundant in leafy greens. In that context, following our culinary pleasure is a good heuristic for nutritive eating. But our context has changed. The modern landscape is full of profit-seeking large-scale food corporations which have filled our environment with cheap, easily manufactured nutrition-free crap, optimized for quick pleasure. The pleasure-heuristic no longer works; hostile forces have gamed the gap between pleasure and nutritional content. In this new context, we now have good reason to revise our heuristics.17

Similarly: following one's cognitive pleasure can be a useful heuristic. Pleasure in cognition can be correlated with cognitive success. We can pursue this pleasure

---

17. Please don't mistake the argument that culinary pleasure is a poor heuristic for nutrition for a general claim that culinary pleasure is bad. It seems plausible that culinary pleasure is intrinsically valuable *and* that it is the subject of nutritional gaming. There is, however, another argument on might make:

that such companies can also game us by offering us cheaply producible quick pleasure, which make replace deeper pleasures — but that is a topic for another day.

for its own sake, but we can also use its presence as heuristic for cognitive success.[18] But things change once our environment fills with the products of manipulative actors, who are incentivized to carefully engineering cognitive products optimized for cognitive pleasure instead of for truth. We need to update our heuristics in response.

How might we update them? Let's go back to the food heuristics. I recently realized that I have subconsciously evolved a counter-heuristic in response to the contemporary nutritive landscape. I have become suspicious of easy pleasure in food. When a food is just too yummy, I have learned to stop myself and ask: what exactly is in this stuff? Who made it, and what is it doing to me? This is not because yumminess is intrinsically bad, but because our environment is now full of entities which profit from the manipulation of yumminess signals. So I have changed my heuristic, tried to tune into the particular kind of hollow, mechanical, addictive culinary pleasure that is the typical product of corporate profiteering.

I suspect that we should develop an analogous counter-heuristic towards ideas. If an idea is just too delicious — if it goes down incredibly easy, confirms everything I want to know, and makes me feel great — then I should be suspicious. Not because the avoiding pleasure is a good universal principle of rationality, nor because pleasure is intrinsically bad, but because this is a good epistemic heuristic for our current context.[19]

Trust

Another key vulnerability for cognitively limited beings arises from trust. Much of the current misinformation crisis seems to derive from misplaced trust — trust in anti-science celebrities, trust in conspiracy theory forums and propagandistic media networks sources.[20] There is a vice epistemology reading of this,

---

[18] Perhaps a more plausible reconstruction is not as a direct heuristic, but a heuristic-like useful motivation. A being motivated to pursue cognitive pleasure would be likely to accrue cognitive successes, in the right environment. See Nguyen

[19] But it's not all suspicion and self-denial. As I've argued elsewhere, the quality of intellectual playfulness can be protective against hostile epistemic environments. Many of the worrisome belief systems and belief communities we've looked at function as a kind of epistemic trap, like echo chambers and conspiracy theories. They are problematically self-reinforcing. Intellectual playfulness is the disposition to take pleasure in cognitive exploration for its own sake – in trying out new perspectives and ideas for fun. Such playfulness can function as something of an epistemic insurance policy against being caught in an epistemic trap (Nguyen, 2022d). Such playfulness would be particularly useful countertendency to cultivate in a world full of epistemic traps.

20. For contemporary empirical work on the prevalence of large-scale propaganda, see Benkler et. al. (2018). I offer a discussion of this piece, in which I compare the polarization narrative and the

which Cassam offers: such people are gullible and lazy. But there is also a hostile epistemology reading. Because we are so cognitively small, in order to cope with the world, we must trust each other, and that trust makes us profoundly vulnerable. That trust can be exploited, even when we have done our due diligence.

The world is too large for one human mind to grasp. The scope of the sciences is vast because the world is vast. We cannot think for ourselves; as John Hardwig (1985) puts it, we are epistemically dependent on one another. The basic epistemic condition of the current era, says Elijah Millgram (2015), is hyperspecialization. Science has become so complex that we need specialists who devote themselves to chemistry, statistics, ethics. In fact, we need is hyper-specialists — specialists in arcane sub-disciplines of sub-disciplines.

My spouse is a chemist; she specializes in the analysis of organic compounds using one specific methodology, mass spectrometry. As she will freely admit, she possesses no expert knowledge of physics or statistics. Actually, she tells me, she has precious little expertise about most other branches of chemistry, or about the details of the statistical software built into the instruments she uses every day. She is master of a teeny-tiny patch — as are we all. So, chains of scientific reasoning have to pass through many specialist domains. The analysis of climate change includes contributions from applied physics, chemistry, meteorology, geology, statistics, and computer modeling. But since the demands of hyper-specialization are so intense, no single person can actually follow the whole argument over its course. Instead, they have to trust one another. The industrial engineer has to trust the results of the chemists and materials scientists; the chemists have to trust their instruments and information-processing software packages, and the formulas passed to them by the statisticians. And on and on, for thousands of sub-disciplines.21

Science depends on non-experts trusting experts. When my spouse trusts the statistical software package that analyzes the data from her mass spectrometer, she counts as a non-expert trusting an expert in statistics. But here is the essential dilemma: how does a non-expert identify which expert to trust? I've argued that, left to their own resources, an individual non-expert cannot perfectly identify the right experts. Their attempts are always vulnerable to exploitation. And this is vulnerability is created by the basic structure of our epistemic condition. We need to trust

---

propaganda narrative of the misinformation crisis in Nguyen (2021c).

21. Millgram is particularly worried about the accumulation of errors, as techniques and technologies get exported from one speciality to another, and simplified for export at every crossing. Every discipline will have to simplify its results for export: they create a simple principle to follow, or tool to use. That principle applies in most circumstances, but it has defeating conditions. Millgram's worry is that knowledge of the defeating conditions doesn't get exported, so people in other sub-disciplines won't be able to recognizing when they are applying the principle or tool inappropriately (Millgram 2015, 1-76).

experts precisely because they understand things which we do not. But how can we check their knowledge where it extends beyond our understanding? (Many solutions have been proposed to solve the expert identification problem, but, as I've argued elsewhere, they mostly don't work. 22)

What do we do? We take shortcuts. We use signals of competence and trust-worthiness — what Miranda Fricker calls "indicator properties".23 I trust this person with my life because they are wearing a lab coat and appear in the right medical spaces; I trust that person about climate history because they are speaking in an authoritative tone of voice, have many detailed anecdotes, and are on stage at a fancy venue. One might respond that we are using people's roles and uniforms as proxies for some institutional approval. The fact that this person has appears in the medical office, or on stage at some Harvard public lecture, indicates that they have been vetted by the right professional institutions. But what makes us trust one institution over another?

Relying on institutional vetting simply pushes the expert identification problem one step back. Now we have to identify the right institutions to trust, rather than particular experts. But the same problem recurs. We are relying on institutions because they are bigger and can understand more; they can grasp what we individuals cannot. But since what we want from those institutions is knowledge beyond our grasp, we don't have the means to perfectly select the right institution. So we have to rely on approximate indicators. This is another shortcut forced on us by the size of the world, which can, again, be exploited by faked signals of institutional competence. This vulnerability arises from the basic structure of the enterprise: our attempt, as individuals, to join together and understand a world beyond our individual grasp.

One might think that the right response is expert monitoring and regulation — say, by implementing schemes of transparency, in which experts justify their actions to the public. But, as I've argued elsewhere, we cannot hope to secure our expert trust through transparency. Schemes of transparency ask experts to explain themselves to non-experts. But expert reasons are, by their very nature, inaccessible to non-experts. So schemes of transparency will either force the experts to lie, yielding a poor guide to expertise. Or transparency will force experts to confine themselves to actions which are explicable to non-experts, which will prevent

---

22. My original argument is in (Nguyen 2020a). Here I was mostly pessimistic about the expert-identification problem, with some optimism about some cases: namely those cases where there was a *litmus test*, in which experts could create some easy, publicly comprehensible result. In more recent work (Nguyen 2022a), I have become even more pessimistic on this topic, because I have begun to take more seriously the problem of non-experts *misidentifying the litmus test*. What pushed me here was the literature on bureaucracy and metrics, in which it became clear that public metrics for judging often failed, because the public metric oversimplified the standard of success.

23. (Fricker 1998). In *Epistemic Injustice,* Fricker does briefly note that our need for quick signals arises inevitably from our cognitive finitude (2007, 17).

them from fully acting on their expertise (Nguyen 2022a).

The point here is not to get rid of transparency, but to see that transparency isn't a perfect solution to the expert identification problem. We need some degree of transparency and monitoring, because otherwise we will be helpless in the face of corruption and bias from fake experts. But if we try to go towards a scheme of total transparency — if we only trust those experts who can make all their reasoning explicable to us non-experts — then we will undermine their expertise. The whole reason we needed to trust experts in the first place was because they could go where we could not. Trust lets us go beyond ourselves. Trust lets us partake of the fruits of an expertise which we do not, ourselves, contain. Too much transparency robs us of that potential for agential expansion. Transparency chains experts to acting within our sight, to doing what we can understand. If we want to access the power of collective inquiry — if we want to leverage the ability of networked experts to go beyond the grasp of individual understanding — then we must trust beyond our ability to verify.

Let's go back to Oliver the conspiracy theorist. Cassam says that Oliver is gullible because he believes theories that have been fully discredited. But what institutions did the discrediting? Here is our alternate explanation, by way of hostile epistemology. Oliver could have fallen into a credentialing trap. Oliver needs to use imperfect signals of expertise to figure out who to trust. A poser institution might have gamed those signals and successfully masqueraded as an expert. Now Oliver believes a false network of institutions which are actively discrediting all other sources. ("Don't trust the mainstream media, they're all liars!") Oliver follows along. But that isn't, by itself, bad policy. We might be tempted to criticize Oliver, to say that Oliver should decide who to trust for himself, instead of outsourcing it. But we all outsource our trust. We are all trusting some institution to tell us which experts to trust and which to distrust. I believe that certain people do good science and others do bad science, but I can't verify that for myself. I trust science journalists, credentialing institutions like universities, and professional accreditation systems like the American Medical Association, to vet my experts for me. We not only must trust others' expertise, we must also trust others to tell us which experts to trust.

I'm not saying here that Oliver is definitely the victim of exploitation. I am simply arguing that exploitation is another available explanation. We needn't jump immediately to the conclusion of epistemic vice. I suspect the tendency to so readily blame Oliver arises, again, from that dream of invulnerability — from the fantasy that we could perfectly secure out trust. But trust is actually antithetical to invulnerability.

In her foundational paper, "Trust and Anti-trust", Annette Baier says that vulnerability is at the core of trust. What it is to trust, she says, is to entrust something of ours to somebody else's care — to put ourselves in another's power. And that

vulnerability is on full display in our relationship to experts. I am profoundly vulnerable when I take the medicine my doctor recommends, when I fly in a plane that's been checked over by mechanics. It is tempting to try to eliminate that vulnerability — to look for some procedure, like transparency, that will guarantee that our trust is always well-placed. But perfectly secure trust is only a fantasy. Baier's insight is that we trust because we cannot do it all on our own; so what it is to trust is to let the trusted person operate outside of one's sight. In some cases, they are literally outside of our sight – like when I leave my child with the babysitter and drive away. But experts operate even more deeply out of our sight. I may be able to physically inspect every step in the manufacture of antibiotics, but I cannot myself verify that the manufacturing was done properly. To trust an expert is to let them operate outside of our understanding. We need to trust because our intellectual vision — our capacity to understand — does not go everywhere.

The exploitable vulnerability of trust is worsened by time pressure. Sometimes we need not only to put ourselves in another's power; we need to cast all doubt from our minds. As I've argued elsewhere, there is a specific type of trust which involves, not just entrusting ourselves to another, but also suspending all suspicion about that relationship – ceasing any questioning. Trust often involves taking on an *unquestioning attitude*. I don't question when the *New York Times* tells me a major winter storm hit the Midwest, I simply believe it. I don't question the items on the shared Google Docs shopping list I share with my spouse; I simply buy whatever is on the list, without a second's pause. I don't question my Google Calendar when it tells me I have a department meeting; I just go. Such an attitude is required, again, because of our limited cognitive resources. I don't have time to question every external source and resource. We need the capacity to settle our minds, not just about particular beliefs, but about whole external resources (Nguyen 2022b).

Unquestioning trust has a very specific function. Most of the time, we run with an unquestioning attitude towards our internal sub-parts. I don't question the deliverances of my memory, my sight, my hearing. Mostly, I just believe my memory about when I'm supposed to pick up the kids from school, or whether there's coffee in my thermos. And it would be wildly impractical to constantly question such internal sub-parts. We have settled our mind to accept the deliverances of our internal cognitive capacities until we have a good reason to become suspicion. The unquestioning attitude permits us to expand the boundaries of that suspended doubt, to offer the same attitude of open and ready acceptance towards external bits as we have with internal bits. To trust something is to let down the barriers, to let it inside. There is, once again, a trade-off. Taking on an unquestioning attitude towards an informational source something is like running an open pipeline into your cognition. Things from the outside can simply insert beliefs and directives

for action, because we have taken down the barriers. Unquestioning trust is enormously efficient, but makes us highly vulnerable.

With trust, we need to overextend ourselves in two ways to cope with the world: we need to trust sources we cannot adequately vet, and we need to settle our mind and end questioning about certain sources. Each of these moves takes us closer to doing together into a kind of unified, quick-moving collective agent. But each of these moves also makes us more vulnerable. And we cannot secure that vulnerability, because we don't have the resources to ensure that we are unifying ourselves with the right external sources. The very conditions that force us to trust also make us unable to secure that trust.

Conclusion

Let's sum up. We are cognitively limited beings trying to cope with an overwhelming world. To do so, we must take risky strategies. We must take cognitive shortcuts, use heuristics. We must trust one another, even though we cannot perfectly secure that trust. These strategies create exploitable vulnerabilities.

Let me offer some big picture conclusions, which circle the same central thought. First: in the current mood in social epistemology, we've given up the possibility that an individual can know the world by themselves. We've accepted that we need collective efforts and group inquiry. Only institutions and communities can hope to know the world. Suppose that's all true. Still, accepting the essentially communal nature of knowledge doesn't solve the problem entirely because *the individual still needs to pick which groups to trust.* If we accept that only groups are the true experts, we still haven't solved problem of proper expert identification. That problem remains whether the candidate expert is an individual or a group. We still need to be able to tell good knowledge communities from bad. But the world is full of hostile epistemic actors trying to fake the signals of epistemic quality. The individual doesn't know enough to perfectly check up on a whole knowledge community. And they cannot, because of the basic conditions of the enterprise. The very reason we needed to engage in such a massive collective first place was that the world was too large for our individual grasp. In order to form these groups, individual actors must choose which groups to link up with. But how can the individual, with their inadequate understanding, select which group to trust? I am tempted to say that the true skeptical problem of our era is not, "How do I know that I'm not dreaming?", but "How do I know I've picked the right groups to trust?"

Second: You might have hoped to get out of epistemology some ideal picture of how to reason — something substantive and relatively static. Maybe we'd get

some universal principles of good reasoning, or a list of the character traits of the ideally virtuous epistemic agent. Or, at least, maybe we could settle on some good heuristics to use. But what I've said so far suggests a different picture. Cognitively limited beings need heuristics and other risky strategies to cope with the world. These are exploitable; hostile forces can game the gap between the heuristic and the ideal. We need to update our heuristics to cope. But since we are still cognitively limited beings, we can never get to some final principles. We'll just end up with newer heuristics, better suited to our newer hostile environment. But these new heuristics are also game-able and exploitable — albeit in a different way — and so the cycle begins anew.

What does that like? Let's go back to the culinary analogy. Once, we might have operated with a heuristic of following simple culinary pleasure; but then our landscape begins to fill with the products of hostile forces, which exploit our heuristic by optimizing food products to hit that simple pleasure. Suppose many of us change course and adopt a new heuristic: that healthy food will be less pleasurable pleasure. We might start using, say, the dryness of low-fat food, or the presence of bitterness and some unpleasant flavors, as a heuristic for healthy food. Of course, if enough of us do that, then that new heuristic can be targeted, gamed, and relentlessly exploited. Corporations could make food products that maximize that horrible feeling of penitent fatless dryness, or for certain kinds of bitter and unpleasant flavors, to give it an aura of healthiness. (I suspect this explains the health food landscape I grew up with.) We might then revise our heuristic again, looking for more complex signals, which, again, can be gamed. We might learn to look for particular ingredients, like whole grains. Of course, that signal can also be gamed — as corporations have learned to add the minimal pinch of whole grains that can qualify them for the "made with whole grains" label.24

What might this look like in the cognitive landscape? I have come to be suspicious of the very clean, over-produced, slickness of corporate information delivery systems — full of memorable mottoes, sharp quips, and slick framings. (You know: the TED talk feel.) As a response, I have developed a heuristic distrust of that too-smooth style, and others have too. But hostile forces have issued a counter-response: a flood of grungy, indie-looking YouTube videos — just some person in a basement, speaking their truth into their webcam. And I am now developing a counter-counter-heuristic suspicion towards that indie-look. And I expect to develop a new heuristic — and that hostile forces in the world will want to exploit that, too.

---

24. Some examples at https://www.gaplesinstitute.org/5-misleading-food-labels/ -- my favorite example of heuristic gaming here is that many people have learned to look for whether sugar is listed early in a nutrition label, which is a sign of high concentration. Food corporations have gamed this by using many different types of sugars, each of which individually is in low concentration, and so appears late in the label.

The right picture of our epistemic lives, then, doesn't involve some final ideal resting place. Even when we evolve and, say, become suspicious of manufactured pleasures, then hostile forces can simply game those suspicions. We come up with new heuristics, which will be subject to new forms of exploitation, which will force us to invent to new heuristics, which will lead to new exploitation strategies. There is no end — only the constant struggle to adapt. Epistemic life for cognitively limited beings isn't a quest for some static ideal. It is a dynamic and neverending cycle of response and counter-response — of heuristics, then exploitation of those heuristics, and then new improved, heuristics and better, more advanced exploitation. Limited beings in a hostile epistemic environment are locked in an unending *epistemic arms race*.

Finally: who's responsible for the exploitation of cognitive vulnerabilities? There's no easy answer here. If you buy the big picture of hostile epistemology, then you'll think that people aren't blameworthy or defective simply in virtue of having cognitive vulnerabilities. We all have to take cognitive shortcuts, to save on cognitive resources in one spot to spend in another spot. What we are responsible for responding appropriately to signs of error, and revising our heuristics. But we will only perform this error metabolism if we have access to the evidence of our errors. And that evidence can be hidden. If we are brought to trust and distrust wrongly, we if have been convinced to settle our minds in certain directions — then we can miss, or dismiss, the evidence of our error. And we cannot even be maximally vigilant for errors all the time. We cannot look in every direction, check under every rock, for signs of our mistakes. Finite beings need to manage their alertness, ration their efforts. Our attention is of necessity focused, and much can hide outside the narrow the spotlight of our vision.

BIBLIOGRAPHY

Bagnoli, Carla. 2016. "Vulnerability and the Incompleteness of Practical Reason." In *Vulnerability, Autonomy, and Applied Ethics*. Routledge.

Battaly, Heather. 2018. "Can Closed-mindedness be an Intellectual Virtue?" *Royal Institute of Philosophy Supplement* 84: 23-45.

Beeghly, Erin. 2015. "What is a Sterotype? What is Stereotyping?" *Hypatia* 30 (4): 765-691.

Begby, Endre. 2021. "Evidential Preemption." *Philosophy and Phenomenological*

Research 102 (3): 515-530.

----. 2021b. *Prejudice: A Study in Non-Ideal Epistemology*. Oxford: Oxford University Press.

Benkler, Yochai, Robert Farris, and Hal Roberts. 20918. *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics.* Oxford: Oxford University Press.

Bowker, Geoffrey and Susan Leigh Star. 1999. *Sorting Things Out: Classification and Its Consequences*. Cambridge, MA: The MIT Press.

Bratman, Michael. 1987. *Intentions, Plans, and Practical Reason.* Cambridge, MA: Harvard University Press.

Campbell, Sue. 1994. "Being Dismissed: The Politics of Emotional Expression." *Hypatia*, 9 (3): 46-65.

Cassam, Quassim. 2016. "Vice Epistemology." *The Monist* 99 (2): 159–180.

Cassam, Quassim. 2018. "Epistemic Insouciance." *Journal of Philosophical Research* 43:1-20.

Cassam, Quassim. 2019. *Vices of the Mind.* Oxford: Oxford University Press.

Coady, David. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Hoboken: Wiley-Blackwell.

Collins, Patricia Hill. 1990. *Black Feminist Thought: Knowledge, Consciousness, and the Politics of Empowerment.* Hyman.

Dentith, Matthew R. X. 2018. "The Problem of Conspiracism." *Argumenta* 3 (2): 327–343.

Dotson, Kristie. 2014. "Conceptualizing Epistemic Oppression." *Social Epistemology* 28 (2): 115–38.

Eaton, A.W. 2020. "Artifact and their Functions." In *The Oxford Handbook of History and Material Culture,* eds. Ivan Gaskell and Sarah Anne Carter. Oxford: Oxford University Press.

Elgin, Catherine Z. 2002. "Creation as Reconfiguration: Art in the Advancement of Science."*International Studies in the Philosophy of Science* 16 (1): 13–25.

Espeland, Wendy Nelson, and Michael Sauder. 2016. *Engines of Anxiety: Academic Rankings, Reputation, and Accountability*. New York, New York: Russell Sage Foundation.

Fricker, Miranda. 1998. "Rational authority and social power: Towards a truly social epistemology." *Proceedings of the Aristotelian Society* 98 (2):159–177

Fricker, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Clarendon Press.

Frye, Marilyn. 1983. "A Note on Anger." In *The Politics of Reality: Essay on Feminist Thought*. The Crossing Press.

Gopnik, Alison. 1998. "Explanation as Orgasm." *Minds and Machines* 8 (1): 101–118.

Grimm, Stephen R. 2006. "Is Understanding a Species of Knowledge?" *British Journal for the Philosophy of Science* 57 (3): 515–535.
———. 2012. "The Value of Understanding." *Philosophy Compass* 7 (2): 103–117.

Hardwig, John. 1985. "Epistemic Dependence." *Journal of Philosophy* 82 (7): 335–349.

Hendren, Sara. 2020. *What a Body Can Do.* New York: Riverhead Books.

Holton, Richard. 2009. *Willing, Wanting, Waiting*. Oxford: Oxford University Press.

Johnson, Gabbrielle. "The Structure of Bias." *Mind* 129 (516): 1193-1236.

Khalifa, Kareem. 2012. "Inaugurating Understanding or Repackaging Explanation?" *Philosophy of Science* 79 (1): 15–37.

Kvanvig, Jonathan L. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Medina, Jose. 2013. *The Epistemology of Resistance: Gender and Racial Oppression,*

Epistemic Injustice, and Resistant Imaginations. Oxford, New York: Oxford University Press.

Merry, Sally Engle. 2016. *The Seductions of Quantification: Measuring Human Rights, Gender Violence, and Sex Trafficking*. University of Chicago Press.

Millgram, Elijah. 2015. *The Great Endarkenment: Philosophy for an Age of Hyper-specialization*. Oxford: Oxford University Press.

Millikan, Ruth. 1987. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.

Munton, Jessie. 2019. "Perceptual Skill and Social Structure." *Philosophy and Phenomenological Research* 99 (1); 131-161.

Nguyen, C. Thi. 2020a. "Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts." *Synthese* 197: 2803-2821.

———. 2020b. "Echo Chambers and Epistemic Bubbles." *Episteme* 17 (2): 141-161.

———. 2020c. *Games: Agency as Art*. Oxford: Oxford University Press.

———. 2021a. "The Seductions of Clarity." *Royal Institute of Philosophy Supplements* 89: 227-255.

———. 2021b. "How Twitter Gamifies Communication." in *Applied Epistemology*, ed. J. Lackey, Oxford: Oxford University Press: 410-436).

———. 2021c. "Was it Polarization or Propaganda?" *Journal of Philosophical Research* 46: 173-191.

———. 2022a. "Transparency is Surveillance." *Philosophy and Phenomenological Research* 105 (2): 331-361)

———. 2022b. "Trust as an Unquestioning Attitude." In *Oxford Studies in Epistemology Volume 7*. Gendler et. al., eds. Oxford: Oxford University Press: 214-244.

———. 2022c. "Playfulness vs. Epistemic Traps." In *Social Virtue Epistemology*, eds. Alfano, Klein & de Ridder. Routledge.

Nguyen, C. Thi and Bekka Williams. "Moral Outrage Porn." *Journal of Ethics and Social Philosophy* 18 (2): 147-172

Pohlhaus, Jr., Gaile. 2012. "Relational Knowing and Epistemic Injustice: Toward a Theory of 'Willful Hermeneutical Ignorance." *Hypatia* 27 (4): 715-735.

Porter, Theodore. 1995. *Trust in Numbers*. Princeton: Princeton University Press.

——. 2020. "Epistemic Agency Under Oppression", *Philosophical Papers* 49 (2): 233-251.

Rippetoe, Mark. 2005. *Starting Strength: A Simple and Practical Guide for Coaching Beginners.* Aasgaard Company.

Sertler, Ezgi. 2022. "Epistemic Dependence and Oppression: A Telling Relationship." *Episteme* 19 (3): 394-408.

Strevens, Michael. 2013. "No Understanding Without Explanation." *Studies in History and Philosophy of Science Part A* 44 (3): 510–515.

Toole, Briana. 2019. "From Standpoint Epistemology to Epistemic Oppression." *Hypatia* 34 (4): 598–618.

Young, Iris Marion. 2002. *Inclusion and Democracy*. Oxford: Oxford University Press.

Wimsatt, William. 2007. *Re-engineering Philosophy for Limited Beings: Piecewise Approximations of Reality*. Cambridge, MA.: Harvard University Press.