Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts

C. Thi Nguyen

(This is a pre-print draft. Please cite the final version, published in *Synthese:* https://doi.org/10.1007/s11229-018-1692-0)

Abstract

I propose to study one problem for epistemic dependence on experts: how to locate experts on what I will call cognitive islands. Cognitive islands are those domains for knowledge in which expertise is required to evaluate other experts. They exist under two conditions: first, that there is no test for expertise available to the inexpert; and second, that the domain is not linked to another domain with such a test. Cognitive islands are the places where we have the fewest resources for evaluating experts, which makes our expert dependences particularly risky. Some have argued that cognitive islands lead to the complete unusability of expert testimony: that anybody who needs expert advice on a cognitive island will be entirely unable to find it. I argue against this radical form of pessimism, but propose a more moderate alternative. I demonstrate that we have some resources for finding experts on cognitive islands, but that cognitive islands leave us vulnerable to an epistemic trap which I will call runaway echo chambers. In a runaway echo chamber, our inexpertise may lead us to pick out bad experts, which will simply reinforce our mistaken beliefs and sensibilities.

Let's presume that cognitive independence is largely a fool's errand, and that modern cognitive life depends on trusting others — to fill in evidence, to corroborate and discorroborate our beliefs, to think through what we do not have the training or time to comprehend for ourselves (Hardwig 1985; Nguyen 2011). As Annette Baier puts it, our epistemic situation makes us trust other people, and trust makes us vulnerable (Baier 1986). Let's also presume that, as Elijah Millgram says, the diversity and specialization of modern scientific knowledge has sharply increased our degree of epistemic vulnerability.

Contemporary epistemic life has evolved so many particular and arcane epistemic specialties that we are often required to pick out experts in faraway fields in which we have no expertise ourselves. But we must depend on those experts, for modern practical arguments often involve assembling long strings of work by different expert specialists together to yield a single practical conclusion, says Millgram. Take, for instance, the design decisions of a nuclear engineer, whose practical conclusions emerge from arguments which meander across nuclear physics, theoretical physics, material science, statistics, and more (Millgram 2015). This all leads us to a rather thorny problem: how do we successfully identify and assess experts in fields in which we are not ourselves expert? And beneath that rather narrow question lurks some rather larger ones. How much control can we exert, and how vulnerable are we to epistemic powers outside of our control? How much can we *manage* our dependence on distant experts, and hence our epistemic vulnerability?

In this paper, I propose to study one particular vulnerability that arises from our dependence on experts – the possibility of our own errors compounding themselves through expert selection. This vulnerability arises most keenly on what I will call *cognitive islands*. Cognitive islands are those cognitive domains in which successfully identifying and assessing expertise in that domain requires that one already have expertise in that domain. You already have to be an expert to find the experts. Cognitive islands are an extreme case, and I propose to study them as a way to start thinking about the general phenomenon of *cognitive isolation*. On a cognitive island, general cognitive capacities and expertise in other areas will not help one to successfully identify experts. Only expertise in that specific domain will do. Plausible candidates for cognitive islands include morality, aesthetics, and perhaps even philosophy. I will suggest that our expert dependence is particularly troublesome on cognitive islands.

Cognitive islands are the places where we have the narrowest set of resources for evaluating experts, which makes our expert dependencies particularly risky.

Some have argued for something even more radical: that expert testimony is completely unusable on cognitive islands. Only the novice really needs expert advice, or so the argument goes, and on a cognitive island, they will be entirely unable to find it (Cholbi 2007). I will argue against this most radical form of pessimism and show that there are legitimate uses of experts even on cognitive islands – namely, all those cases where experts need the help of other experts from the same domain. However, cognitive islands do exact a heavy epistemic price. The cognitive isolation of such domains leaves us dramatically vulnerable to an epistemic trap, which I will call a *runaway personal echo chamber*. In such a runaway echo chamber, one's own flawed expertise will lead one to trust bad experts, which will reinforce one's mistaken beliefs and sensibilities.

Echo chambers have been analyzed before, usually with an eye to the epistemic vices which lead to their formation. For example, Jamieson and Cappalla have argued, in their analysis of the epistemic community surrounding Rush Limbaugh and Fox News, that an echo chamber has formed because right wing media leaders have consciously tried to alienate their listenership from all other sources. In their account, the listenership is sometimes complicit; they have often sought reassurance over truth (Jamieson and Cappella 2008). I do not deny their analysis of that phenomenon, but I am describing an additional, and, to my mind more frightening, possibility. Jamieson and Cappalla's echo chambers are a particularly social phenomenon – the result of malicious action, purposeful social isolation, and intellectual laziness. But if my analysis is right, then a different type of echo chamber can form on cognitive islands with no epistemic vice in sight. Under these cognitively isolated circumstances, we can simply be trapped despite our best efforts. Flaws in our beliefs and

cognitive abilities will reinforce themselves by influencing our selection of which experts to trust. This is not necessarily the outcome of epistemic laziness or maliciousness. These sorts of runaway echo chambers turn out to be a basic feature of human cognitive life — an outgrowth of our epistemic limitation and vulnerability, combined with the cognitively isolated nature of certain domains.

Perhaps there are no true cognitive islands, in which case we will be saved from this particular epistemic tragedy. But I suspect that cognitive islands do exist. There is much reason to think that the moral domain and the aesthetic domain are cognitively isolated in this way. And if Millgram is right, it might even turn out that a very large number of specialized domains are cognitive islands. But even if it turns out that cognitive islands are rare or non-existent in their fullest flowering, this investigation may still prove of use. For even if no domain is perfectly cognitively isolated, some domains are relatively more cognitively isolated than others. Studying the idealization will help us begin to think about more moderate cases.

What is a cognitive island?

In order to understand the nature of cognitive islands, let's start with a brief analysis of what isn't a cognitive island — what's on the cognitive mainland, so to speak. A domain is on the cognitive mainland if expertise is unnecessary to decide which experts to trust. On the cognitive mainland, a novice can successfully identify a genuine expert from among the posers. They can also assess an expert's degree of expertise, especially in those situations where different legitimate experts disagree. (I'll refer to these two processes – of identifying a genuine expert and of assessing their degree of expertise – collectively as the process of

evaluating an expert.)[1] How might this be possible? Let's start with Alvin Goldman, who considers various methods by which a novice might assess the credibility of an expert, in order to decide which expert to trust. Many common approaches won't work, says Goldman. Imagine an utter novice confronted with two climatologists arguing for radically different climatological interventions, each citing various atmospheric studies and complex statistical analyses. What methods does the novice have to decide which expert to trust more? The novice can't adequately assess the evidence for themselves, nor the climatologists' criticisms of one another. Assessing their apparent skill at debating — like their ability to say *something* in response to any point — is an indicator of expertise, but an imperfect one, because rhetorical skill is imperfectly correlated with genuine expertise. The most promising route for evaluating expertise, says Goldman, is in studying a purported expert's track record. The novice looks at how successful a purported expert has been in answering past questions (Goldman 2001, 93-108).[2] But how could a novice properly evaluate an expert's track record? Wouldn't that, itself, also require some degree of expertise?

Not necessarily. A novice could accurately evaluate an expert's track record if some of the expert's actions impinged in some comprehensible way on the world outside their expert domain. In many domains of knowledge and skill, there are tests for expertise whose results are interpretable by the inexpert. Many expert realms come with stock accomplishments: weather forecasters must be able to accurately forecast the weather, mechanics should be

[1] Thanks to an anonymous reviewer for stressing the importance of the distinction between identifying and assessing expertise.

[2] I leave aside Goldman's discussion about whether we can look to the weight of numbers and trust the expert who has the most other experts on their side. First, Goldman's argument against using the weight of numbers has been challenged (Coady 2006)(Lackey 2013). Second, the approach presumes that we have already successfully identified all the genuine experts, and are only trying to ascertain which expert to trust more in a case of expert disagreement. The weight of numbers will not help us with dilemma I've proposed, which involves first identifying the genuine experts from the fake, before we can go about assessing their relative merits.

able to fix cars, and the like. In some cases, the success of such an accomplishment is clear even to a novice. If my car wouldn't start this morning, but my mechanic made it work again, then this is a mark in favor of their expertise. I need no particular mechanical expertise myself to appropriately make that particular judgment.[3]

In some domains, such tests are available on an immediate and individual basis. If you claim to be an expert axe-thrower, your claim can be tested simply by handing you a throwing axe and pointing to a distant target. For more esoteric domains, the testing may be more involved. Suppose that you are an aeronautical engineer. I have no way to assess your individual judgments and statements about airplanes — I can, in fact, barely understand them. But I trust the discipline of aeronautical engineering as a whole, since planes rarely fall down. Thus, I have reason to trust the internal standards of the discipline as a whole. The facts that you have been licensed by the professional association of aeronautical engineers and have gone through peer review to publish your research on aeronautical engineering carry weight with me because I have a method for evaluating your entire field. Let me call all these sorts of domains *obvious*: they produce, individually or collectively, some result whose successfulness is available to the inexpert observer.

[3] My discussion of inexpert-available tests here is similar in spirit to Goldman's discussion of track records, though it differs in some details. Goldman's solution to the track record problem is to focus on statements that begin as esoteric (firmly within the expert's domain) and then later become exoteric (accessible to the novice). His example is the diagnosis of a repairman ("Your whombulator is causing your figgle to spile") which is esoteric when initially offered but becomes exoteric after a successful repair. But I do not think these esoteric/exoteric transitions capture all the relevant cases. Simple prediction cases ("It's going to rain tomorrow") are not esoteric when made; they are simply not yet verifiable. Furthermore, Goldman's solution focuses on statements, which reflects his focus on cognitive experts whose abilities involve relationships to true statements; he excludes from his discussion experts of skill, such as the axe-thrower. But many useful tests involve no propositions, just successful practical demonstrations of skill: hitting a target, building a machine. I am trying to offer a more general account of expert and expert finding – partially because there are some convincing arguments that moral expertise is significantly more of a skill than a strictly cognitive domain. In any case, my discussion of inexpert-available tests is intended to include Goldman's esoteric/exoteric transitions along with some other important cases.

Of course, there are various degrees of accessibility for such inexpert-available tests. Axe-throwing has a relatively clear test, whose result is quite comprehensible to the inexpert. Other tests might require more substantive capacities, and more fragile chains of reasoning, to interpret. For example, checking up on my accountant might require some algebra on my part, which is a common, but not universal, skill. Since the accessibility of these tests comes in degrees, accordingly, the obviousness of domains also comes in degrees. For simplicity's sake, I'll say a domain is obvious when it has tests whose results are accessible to the average adult.

On the other hand, some domains are *subtle*: there are no such tests available to the inexpert observer. There are many plausibly subtle domains: abstract expressionist painting, number theory, and particle physics are decent candidates. In subtle domains, it takes significant expertise to tell the good results from the bad. To my inexpert eye, a good proof in number theory is indistinguishable from mechanically generated garbage. And it has often been observed that the artistically unsophisticated cannot tell great works of abstract expressionism from a child's painting. There are three problems which face the inexpert who seeks advice in subtle domains. First, in many cognitively legitimate domains, we are inundated with individual frauds. Even the most ardent moral realist must admit that the moral domain is full of spiritual charlatans and fake gurus. Second, whole domains have been charged with fraudulence. Take, for example, the oft-cited claim that the average wine-drinker cannot tell the fine stuff from the plonk. Some take this to be a sign that the whole practice of wine appreciation is rubbish. Similar charges have been leveled with varying degrees of effectiveness against new age healing, critical theory, Freudian psychology, and modern art. Third, even if the inexpert has somehow successfully identified the body of legitimate experts, in many cases the legitimate experts disagree with each other. How is the

inexpert to assess the various experts' degrees of expertise, and choose which ones to trust? How do we distinguish the genuine gurus from the fake, the real science from the pseudo-science, the top experts from the mediocre ones?

Even in subtle domains, the inexpert are not entirely without resources. Many domains are interlinked, so one can perform what Phillip Kitcher calls "indirect calibration" (Kitcher 1993, 320-3). That is, some fields are interwoven with other fields; thus, we may be able link up subtle fields with more obvious ones. For example, nuclear engineers rely on the results of particle physics, and nuclear engineering has some rather dramatic tests, whose failures are available to the inexpert. The evidence of functioning nuclear reactors and nuclear bombs leads me to trust nuclear engineers, and the nuclear engineers trust and depend on the work of the particle physicists. So, though I can't make any direct judgments about the reliability of particle physics for myself, I can connect it up to another field which I have some capacity to assess. I trust particle physics because the nuclear engineers trust it; here, the subtle field is *linked* to an obvious field. I can use that linkage in several ways. First, I can use a linked obvious field to establish trust in the subtle field as a whole. Second, once I establish that trust in the field as a whole, I have some reason to trust its licenses and rankings of its own experts.

But some domains are not so lucky. Some domains seem both *subtle* and *isolated* — that is, they possess neither an inexpert-available test, nor are they linked to any fields with such tests. These are what I call *cognitive islands*. Without an available test, or links to other fields with available tests, there is no way for a novice to judge the track record of an expert. One already has to be an expert to evaluate another's expertise. On cognitive islands, novices are cognitively helpless when looking for advice.

This problem is not new, though this particular language and framing is my own. There have been several recent attempts to argue that moral experts are hard to find for something like these reasons (Cholbi 2007; LaBarge 1997; McGrath 2008). A similar claim has recently been made for aesthetic experts (Meskin 2007). I suggest that these discussions can be usefully treated together, because the problems arise from a uniform set of conditions.

The moral domain and the aesthetic domain seem some of the likeliest candidates for being cognitive islands. First, both domains are likely subtle — neither seems to admit directly of any inexpert-available tests. The result of, say, great moral sensitivity is a morally good recommendation, but an immoral person wouldn't recognize a good recommendation as such. The same goes for, say, brilliance in free jazz. This is why it is so easy for an outsider to the delights of free jazz to accuse the entire field of fraudulence. One might suspect similar things of literary theory, philosophy, and even perhaps macro-economics. Second, the moral and the aesthetic are plausibly isolated domains. There seem to be no non-moral domains that depend on the results of the moral. And what non-aesthetic domain depends on the results and judgments of the jazz critic? In fact, Millgram has even argued that every academic discipline is, at the very least, subtle, because "in order to apply the standards of the __ologists, you have to be a __ologist yourself" (Millgram 2015, 45). The evidence, says Millgram, is clear: simply think about the helplessness of an academic dean in the faculty hiring process to decide, on their own, the legitimacy of a candidate in any field but their own. They cannot do so independently; they must depend on the recommendations of the relevant department. What's more, even department members from the same field often won't evaluate another prospective faculty member's work directly; instead, they often count the number of peer-reviewed publications, which outsources the judgment to experts in that person's particular sub-field (44-8). Millgram, in fact, thinks that we have plenty of historical

evidence for something approaching cognitive isolation of many empirical disciplines. What better explanation, he says, for the mortgage bubble crash, than that the domain of high-end finance is very hard to check by an outsider (38-9)? Note that, even though the domain of high-end finance might not be a perfect cognitive island, the relative difficulty and obscurity of its tests and linkages may make it function as something very like a cognitive island for a period of time. (We could call such a near-island a cognitive swamp.)

Most of the discussion that follows emerges from recent work on moral expertise and moral testimony. I will begin the discussion in those terms, but then show that the difficulty is generalizable. I will not here argue significantly for or against the claim that morality or any other particular domain is a cognitive island, nor will I attempt here to provide an exhaustive list of cognitive islands. My goal here is only to explore the epistemic significance of cognitive islands, wherever they might happen to be.

Some have argued for the radically pessimistic view that the nature of cognitive islands renders expert testimony useless. I argue against that radically pessimistic view. I argue, instead, for a more moderated pessimism: that experts are usable on cognitive islands, but they cannot act as safety nets. A person with a moderately reliable moral sensibility may be able to pick a good moral advisor, but a KKK member is probably going to pick a bad advisor, who will only take them further down the path of corruption. On a given cognitive island, our ability to evaluate experts depends on our own abilities in that domain, so the experts we choose to trust will be a reflection of and an extension of our own abilities.

In order to see why the radically pessimistic view fails, but the moderately pessimistic view succeeds, I will consider three different problems that might arise on cognitive islands: the credentials problem, the agreement problem, and the licensing problem.

Problem 1: the credentials problem

Much of the previous work in their terrain originates in a puzzle about when we're willing to defer to experts. There seems to be a peculiar asymmetry in our willingness to believe from testimony. We have no problem simply trusting empirical testimony, but we typically take there to be something wrong with simply deferring to moral or aesthetic testimony. There is, for example, nothing wrong with me believing that my car engine needs a new serpentine belt based purely on my mechanic's testimony, but there seems to be something wrong with my believing that Van Gogh's *Irises* is the most beautiful painting, or that medical euthanasia is wrong, based solely on the word of another (Nickel 2001; Meskin 2007; McGrath 2011; Hopkins 2011; Budd 2003; Jones 1999; Nguyen 2017). This intuitive asymmetry between the usability of moral and aesthetic testimony, on the one hand, and empirical testimony, on the other, demands an explanation. One approach has been to suggest that there is a requirement for autonomy in judgment in these terrains, though this approach has been problematized (Driver 2006; Nguyen 2010). Another approach has been to explain the asymmetry as a result of the epistemic difficulties with evaluating experts in these special terrains. It is this latter approach that is relevant to our discussion of cognitive islands.

Let's start with Michael Cholbi's argument, which attempts to establish that there can never be any useful moral advice on epistemic grounds. A moral expert, says Cholbi, is "someone who very reliably, though not necessarily infallibly, provides correct moral advice

in response to moral situations and quandaries" (Cholbi 2007, 324).[4] Cholbi grants that there can be moral experts, and even grants that they can be identified successfully by other experts. But precisely those who need the help of experts — the non-experts — thereby lack the resources to successfully identify experts. The situation is different for empirical domains, says Cholbi, because those domains have tests available to the inexpert. But there are no such tests for morality, whose results are readily interpretable by the entirely morally obtuse. The output of a moral expert is morally sound advice and morally good action, and recognizing these requires moral discernment. We could widen the scope and look at the purported expert's reasoning along with their output, but then the same problem recurs. Fake experts will give bad moral reasons, but sorting the good moral reasons from the bad also requires moral expertise (332). Similar arguments can be deployed for other moral capacities, like moral sensitivity.[5] The moral non-expert is no position to identify good moral output, and so in no position to identify genuine moral experts. Cholbi calls this the "credentials problem for moral expertise" (325).

Cholbi's conclusion here, translated into my terminology, is that the moral domain is subtle. But, as we have seen, if the moral domain were linked to another, we might still be able to identify moral experts. Cholbi does not raise the possibility of Kitcher-style indirect calibration. Let me propose the following slightly amended argument, which I will call the *strong credentials problem*: given that the moral domain is both subtle and isolated, a moral non-expert has no resources with which to identify a moral expert. Since only a non-expert

[4] Cholbi adds a further criterion to moral expertise — that the expert must also be motivated by their own prescriptions — but this criteria plays no part in the discussion of the credentialing problem.

[5] Cholbi's argument here can be buttressed with Driver's observation that general features of rationality - like an aura of intelligence – are not well-correlated with moral expertise.

needs moral advice, there is no situation in which an advice-seeker could properly identify a good source of such advice. Therefore, there is no usable moral advice.

Notice that Cholbi's argument draws on the general features of cognitive islands. If the argument works for morality as he describes it, it should also work for other cognitive islands. If that is all correct, it would yield the most pessimistic of conclusions: that on all cognitive islands, experts are never useful. For example, if music is a cognitive island, then I can only identify a musical expert if I am one already, but if I am already a musical expert, I don't need any musical advice. But surely this is wrong. I have plenty of musical training and I seek the advice of other trained musicians all the time — about what albums to buy, or who might make a good jazz guitar, or about who the best voice instructor is in town (Nguyen 2017). What has gone wrong here?

The situation of moral advice seeking

Cholbi's key assumption for the conclusion of unusability is that experts don't need expert advice: "…One's own expertise obviates the need to seek out experts in the first place. Experts don't need the expertise of other experts" (324). But, as we've seen, the argument yields an unwelcome conclusion. What's gone wrong? Perhaps the problem here is an implicitly binary conception of expertise: either you are a total novice, in which case you're helpless, or you're an expert, in which case you need no advice. But this seems unreasonable. First, the typical advice-seeker would, at the very least, conceive of themselves as a *partial* moral expert. They have both some moral resources and some need for advice from other moral experts. Second, insofar as we don't think of experts as epistemically perfect, even experts need the advice of other experts. Thus, Cholbi may be right that the utter novice has

no resources to discover an expert. But the utter novice is not the only person who needs moral advice. It is important to realize that expert dependencies arise not only for novices. Partial experts and full experts also depend on other experts, even in their own fields of expertise.

The problem seems to be Cholbi's use of a rather extreme idealization as his paradigmatic case of moral advice-seeking: the complete moral incompetent seeking a reliable moral guide.[6] As an antidote, let's look at some more recognizable situations of moral advice-seeking. These cases will function as counter-examples to Cholbi's radical pessimism; but they are also intended to help us flesh out how we actually identify, assess, and use moral experts, which will, in turn, herd us towards the real problem.

CASE 1: FASTER

> My friend Samantha and I agree strongly on most cases. In fact, on cases to which we've both given our careful consideration, we almost always agree. But I have discovered that Samantha is simply faster than me in making these judgments. She is a faster reasoner. It takes me a month to get to the conclusion that she arrives at in moments. Even when she tells me her reasoning, it takes me a long while to understand. Now, an emergency situation comes up — a student has plagiarized and has claimed that they didn't fully understand how citation was supposed to work, and I have to decide right away whether to escalate it to the Dean or to give them another chance. I don't know what to do yet, but Samantha thinks I should give the student another chance, and I defer, because I believe her to be reliable.

This first case is the simplest. Samantha's moral beliefs and mine are identical in the long run. I don't think that Samantha knows more; I just think she knows it more quickly. The

[6] There seems to be a subtle slippage here. Cholbi claims as the target of his analysis the "non-expert", which is a fairly broad category that would include all the intermediate categories between novice and expert. Thus, he can be assured that his analysis is complete, because the categories of "expert" and "non-expert" are exhaustive. But his argument only works if we interpret "non-expert" to mean a complete novice, and this renders his analysis of experts and non-experts non-exhaustive.

difference between Samantha and me is not one of in-principle access to the moral domain, but merely a difference in speed. Thus, in this case, I can think that another person is significantly better than me at a particular aspect of moral reasoning, even though, in the end, we agree. This case already shows the problem with Cholbi's view. Here is a case where I can take myself to be a moral expert and yet still have a use for moral testimony from another. There are many similar cases in aesthetics — many of the music reviewers I trust are those who quickly see the aesthetic qualities in an album that it takes me hundreds of listens to discover.

The next case is adapted from Karen Jones' bias cases (Jones 1999).

CASE 2: BLIND SPOT

> I take myself to be morally astute in general, but I know that I tend to slip up when it comes to matters of racial sensitivity. Things I thought were harmless jokes turn out to be offensive, and I hurt people's feelings without intending too. This may be because I grew up in a racially monotonous environment and have little first-hand experience with racial discrimination. However, I have a friend who doesn't seem to share this blind spot. Now, I wouldn't trust just anybody about this — I think plenty of people are oversensitive about such issues or just plain humorless. But Jed is, in general, morally clear-sighted, has a good sense of humor, and doesn't overreact. I have a great idea for a Day of the Dead Halloween costume, but when I run it by Jed, he says that it's racially and culturally insensitive. I don't see the problem myself, but I trust Jed, so I go as Thor instead.

Blind Spot is a more complicated case than Faster. In Faster, I had access to all the same moral domains as Samantha, my access was simply a bit slower. In Blind Spot, I don't have access to the key moral sub-domain. How then can I evaluate somebody's expertise in that sub-domain? I can judge that Jed is morally apt in those moral sub-domains in which we overlap; this gives me good reason to think him similarly apt in my blind spot. This presumes that different moral sub-domains are not isolated from each other. But that isn't a problem — our initial presumption was only that the whole moral domain was cognitively isolated from

other domains, not that its internal sub-domains were isolated from one another. In fact, the very plausibility of the Blind Spot case suggests that morality is not so radically partitioned. Moral expertise is not the same thing as unqualified moral perfection.

We can find even more plausible uses for moral testimony when we look for uses beyond outright deference:

> CASE 3: CORROBORATE
>
> My spouse said something to me last night that I found hurtful. It was subtle, and I'm unsure about my response — maybe I'm just being oversensitive and overreacting. I call one of my best friends, Jenny, who I consider to be a sensitive and morally trustworthy person. I recount the conversation from last night. Jenny reacts strongly; she agrees with me that my spouse's comments were subtly malicious. My sense of being wronged solidifies, and I make up my mind to say something about it tonight.

It is, in fact, difficult to imagine a morally engaged life without this sort of corroboration. This case is particularly interesting because it is a case in which we use expert testimony, but it is not a case of deference. The discussion of expert testimony often seems to presume that deference is the paradigmatic use for experts. This may be true in legal cases of relying on scientific experts, but in moral and aesthetic life, the typical uses of experts are very different (Nguyen 2010, 2017).

The Corroborate case suggests an interesting variation on the Blind Spot case. Along with outright blind spots, there are also what we might call "blurry spots" — areas where we are not utterly blind, but where we are less than reliable. For most of us, it's likely that we have one very particular, but very important blurry spot: ourselves. Most of us have some trouble seeing moral flaws and errors in ourselves, or noticing applicable moral criteria.[7] When I am

[7] I owe this case to an anonymous reviewer.

wrapped up in working on a philosophy paper, I can become so absorbed in my work that I fail to notice my rudeness to other people. Others can (and do) help me by pointing out my moral error and my obliviousness. My response here isn't outright deference to the testimony of others — I apply the criteria for rudeness myself, once another has brought it to my attention. But notice that there is a substantial degree of trust in the negative case. When somebody tells me that I haven't, in fact, morally screwed up, I am trusting their greater degree of perceptivity in my blurry spot. In fact, if we are surrounded by morally trustworthy friends, we can legitimately extract some small moral corroboration from the fact that they haven't said anything lately.

Two relevant analyses can strengthen my case here. First, Karen Jones argues that trustworthy people will be pro-active in aiding those that depend on them (Jones, 2012). Second, Sanford Goldberg argues that when surrounded by an appropriate social network of other epistemic agents with adequate coverage, one may rely on them to convey relevant information in a timely manner. Thus, one may extract from the silence of an adequate network the knowledge that nothing relevant and important has occurred. For example, I know that no meteor hit a major city yesterday, because if one had, somebody would have said something by now (Goldberg 2011). What I've argued for here, then, is the existence of moral coverage networks. Most of us have blurry spots with regards to the morality of our own actions, so we depend on our moral coverage networks to check our worst tendencies. And, provided that we have good moral coverage networks, we can use their silence as a form of implicit corroboration inside our blurry moral self-regard. Such blurry spots also surely extend outside the moral domain. One might think that in generating philosophical arguments, most of us have blurry spots about the quality of the arguments we've generated, as well as for the arguments supporting those beliefs that are personally important to us.

Thus, not only do we rely on others to point out what we've missed, but we take their active agreement, and even their mere critical silence, as indicative: it is positive evidence that we have not, in fact, missed something in our blurry spot.

Strangely, the moral domain is precisely the place where Cholbi's argument seems to have the least teeth. Who thinks of themselves as a complete moral novice? Among the candidates for cognitive islands, morality is the one domain in which almost everybody claims some degree of expertise, even if it is only partial. Cholbi's version of the problem actually seems more relevant in other fields. It seems much more likely that somebody might conceive of themselves as knowing absolutely nothing about music and yet need musical advice. And if any of the more abstruse intellectual fields, like macroeconomics, turn out to be cognitive islands, then the strong credentials problem will have very significant force there. There are plenty of self-acknowledged novices in macroeconomics who might like some expert advice about various policy proposals during an election year.

For the cases involving the credentials problem, we've learned something very useful. First, the radically pessimistic conclusion — that there is no use at all for expertise on cognitive islands — is wrong. The strong credentials problem shows only that the absolute novice cannot identify experts on cognitive islands. But that still leaves on the table all those plausible uses of experts by other experts and by partial experts.

Problem 2: the agreement problem

The strong credentials problem fails to support the deeply pessimistic conclusion that there is no use for expertise on cognitive islands. Instead, let me suggest a distinct, but related problem. In reflecting on the above cases, one might begin to notice a worrying similarity. In

all these cases, I assess the quality of another person's moral judgments by checking their pronouncements against what I take to be the truth — which, in turn depends on an exercise of my own moral judgment. Similarly, I evaluate another's moral reasoning by comparing it with what I take to be good moral reasoning. In short, my appraisal of another's expertise seems to depend on the degree to which they agree with my own judgments and reasoning. This threatens to place an upper bound on the degree of expertise I can attribute to others. If I assess another's expertise based on agreement, what basis could I possibly have for judging another to be more expert than myself? Must I not judge any departures from my own reasoning as failures on their part? This would lead to a severe limitation: on cognitive islands, I might be able to judge somebody to be an epistemic peer, but never an epistemic superior.

So here is a second possible problem for inhabitants of cognitive islands: call it the *agreement problem.* This is a problem, not for the initial identification of other experts, but for assessing another's degree of expertise relative to one's own. Suppose that I am an expert on a cognitive island; I must evaluate others' expertise by their agreement with my own views. Now suppose a purported expert offers me a conclusion on something about which I already have a belief. Either they agree with me, in which case their belief is non-independent of mine and has no weight; or they disagree with me, in which case I ought to downgrade my judgment of their expertise relative to mine, at least for this particular claim. The result is a weaker version of the pessimistic conclusion of unusability: I only have a use for fellow experts in those cases where I have no belief of my own; there are no uses for experts where I have my own belief. Thomas Kelly suggested something similar in an early paper – that moral

disagreement shouldn't matter, because disagreement counts as a disqualifying condition for epistemic peer-hood (Kelly 2005).[8]

But we can dissolve the agreement problem too, because we can reject the premise that agreement is the only means to assess another's expertise. Consider the following cases:

CASE 4: CHALLENGER

I have an academic advisor, Scott, who constantly challenges my considered judgments with new arguments. After considering the new arguments, I usually recognize them as being better than my standing basis for belief and revise my beliefs. In this case, I am still using my own reasoning to assess the reasoning of another. Still, the fact that the better arguments consistently *originate* with my advisor is a reason to think that they are more expert than me.

CASE 5: EXPLAINER

I have all sorts of moral intuitions about which I feel strongly, yet I am often at a loss to explain the complicated ways that my intuitions change in subtly different cases. I often worry that they are just incoherent. I explain these judgments to my friend, Judith, and explain my difficulty in accounting for them. Judith not only agrees with my intuitions but also consistently articulates and explains my intuitions far better than I can, and those articulations resonate with me.

CASE 6: SENSITIVE

My friend Xiu is extremely aesthetically sensitive. Often, I will confront some artwork and think I've seen all there is to see. But when I talk with Xiu, she points out all sorts of relevant details I didn't notice. Once Xiu points them out, I see that they are important and relevant, and they transform my view and understanding.

[8] I've discussed Kelly's account at greather length in (Nguyen 2011).

Certainly, in all these cases, agreement plays an important role, but it is not my only resource. In Challenger, I eventually come to agree with Scott's judgments. But the fact that Scott consistently originates these judgments, and the fact that my original judgments come to seem, in retrospect, naive and ill-considered, gives me reason to believe that Scott is more expert than me. In Explainer, Judith and I agree in our intuitions. But her ability to produce explanations for these intuitions and articulate the reasoning behind these intuitions gives me reason to think that her understanding is deeper than mine. And in Sensitive, agreement does initially establish that Xiu is some sort of aesthetic expert. But the fact that Xiu notices things more quickly than me and can successfully guide and improve my aesthetic perceptivity is a reason to take her to be even more of an expert.

In all these cases, I initially identify expertise through agreement. But after the initial recognition, I have more fine-grained resources for further assessment. I can reflect on how our agreement came to be. And these *conditions* of agreement can, in turn, give me reasons to rate another as more expert than me, even on a cognitive island. Thus, even the weaker argument for the unusability of experts on cognitive islands fails. Experts and partial experts have both the resources to find and good reasons to use other experts. So do cognitive islands actually present any special difficulties for our dependence on fellow experts?

Problem 3: the public license problem

I have argued against the strong credentials problem and the agreement problem. But the foregoing discussion will help us to see the intractability of a third problem for cognitive islands: that of establishing an uncontroversial credential.

Let us call an uncontroversial, generally accepted method for evaluating experts a "public license." A public license can be extraordinarily useful, for it can serve as the basis for adjudicating disputes between parties. If disputing parties can agree about who the experts are and about their relative degrees of expertise, then we can use experts to adjudicate the dispute. But establishing a public license will require a method for evaluating expertise on which we can all agree.

Sarah McGrath has argued that there is a problem with establishing a public license for moral experts. On a superficial pass, her argument might seem quite similar to Cholbi's. McGrath begins by noting that there is no good "independent check" on moral expertise. An independent check would be a non-controversial method for evaluating an expert. Unlike the medical sciences, "we possess no analogue to an eye exam, by which we might determine whose moral vision is askew and whose is in good working order" (McGrath 2011, 96-9). But notice that establishing McGrath's "independent check" is a much more difficult goal than that of merely discovering an expert for ourselves. Cholbi's question was about whether an individual who needed advice could ever identify a moral expert. McGrath's is after something else: a non-controversial method for evaluating experts which could be used to settle disputes. Cholbi is after a privately justifiable criterion, whereas McGrath is after a publicly acceptable license.

On a cognitive island, expert evaluation is highly dependent on the evaluator's particular abilities. Expert evaluation on cognitive islands will thus be highly controversial. Two people with differing moral sensibilities will identify different bodies of experts, or rank their relative degrees of expertise differently. Arch-conservatives and ultra-liberals won't be able to settle the health care debate by turning to a common body of experts, for their choice of experts will differ as violently as their positions. Even when evaluators might manage to

settle on a procedure for identifying experts and manage to agree on the same body of experts, their differing sensibilities can lead them to assess those experts' relative ranking differently, which would also prevent any use of those experts for adjudication. For example, though various public policy types may all agree on who counts as an economics expert, libertarians and progressives will likely pick out different sub-sets of those experts as the better and more trustworthy sub-set. Thus, though it is possible for an individual to successfully evaluate a moral expert, it is impossible to create a public license for moral expertise. It may be possible for sub-communities to form with sufficiently similar moral sensibilities as to create an internal license, but so long as morality is a cognitive island, and so long as there is substantial moral disagreement, there can be no universal public license for moral expertise.

McGrath's argument works where Cholbi's doesn't. And the problem is clearly generalizable from the moral domain to any other subtle and isolated domain. Only unsubtle or non-isolated domains can offer the hope of public adjudication by experts. But on cognitive islands, there are no such uncontroversial means for identifying and ranking experts. People with different moral beliefs, commitments, and sensibilities will pick different experts in whom to put their greatest trust. The public license problem for cognitive islands is, in fact, insurmountable.

We've now arrived at a very interesting place. Considering the various arguments for the unusability of experts has helped us zero in on the particular difficulties of expert dependence on cognitive islands. On cognitive islands, the evaluation of others' expertise depends on the substantive exercise of one's own abilities. I recognize the expert explainer by coming to understand their explanations for myself. I judge somebody to be sensitive when I am led by them to notice something for myself. But this opens the door for a different sort of

epistemic tragedy. On cognitive islands, we cannot use experts as safety nets; they cannot help us to guard against our having gone completely off the rails. Let's return to the moral realm, where the tragedy is clearest. If my moral abilities were totally off track, then I will select the wrong people as my trusted experts, and they will only further confirm my poor beliefs. Enraged racists will pick more fanatically enraged racists as their advisors. The utterly selfish may find that Ayn Rand echoes their own views and articulates them with convincing clarity, and they may come to treat her works as holy writ. If one also thought that economics was a subtle domain, one might also suspect such bootstrapping would occur there. It might look something like this: the free market libertarian economists would treat corroborations from their fellow libertarians as mattering, but dismiss Marxist economists as having gone wildly off the rails, and vice versa.

The consequence of cognitive islands is not that experts are entirely unusable, nor is it that we cannot ever have the grounds to judge somebody more expert than ourselves. Instead, the real problem is in the possibility of runaway bootstrapping. My evaluation of experts must involve a substantial exercise of my own expertise. Thus, the experts I put my trust in will be a reflection of, and an extension of, my own abilities, commitments, and sensibilities in that domain. A certain degree of epistemic bootstrapping is inescapable when using experts on cognitive islands. There are many uses of moral experts, and all other experts, on cognitive islands, but these uses also leave one vulnerable to pernicious self-reinforcement.

Runaway echo chambers

Cognitive islands are not entirely unnavigable, but they are treacherous going. If my relevant faculties have some grasp of the truth, then I can use them to identify and assess other experts and use those experts to help refine my grasp of the truth. But if my faculties are completely off-track, then using them to evaluate other experts will simply lead me further astray. Think about how the various cases worked, especially the agreement cases. Somebody comes to the same conclusion as I do, but faster. Somebody points out a consideration that I missed, and I come to see its relevance. In these judgments, my abilities are not just repeated back to me; the nature of my epistemic relationship with another expert introduces other factors which give their abilities some extra weight. But those factors are also evaluated with my abilities; thus, the expertise of others, while not a mere *echo* of my own, is an *extension* on my own. My trust in others is grounded in my trust in myself.

In the empirical domains, the obviousness of output and the interconnectedness of domains work as checks on self-reflective bootstrapping. But on cognitive islands, other experts cannot function as an independent fail-safe, precisely because the experts one selects are an extension of one's own purported expertise. On cognitive islands, one needs to already have a foot in the door. With a foot in the right door, one can bootstrap one's way in further by picking out good experts. But for somebody with a foot in the wrong door, reliance on experts will not help. Bad morality will simply compound itself through badly aimed trust.

Let's call this a *runaway personal echo chamber*. One is threatened with a runaway personal echo chamber when:

1.) one relies on other experts to check and reinforce one's own abilities and beliefs,

2.) one evaluates others' expertise through a deployment of one's own abilities and beliefs,

3.) there is no check on failure for either individual or collective expertise whose application is independent of one's own expert abilities.

Runaway personal echo chambers are particularly threatening because we depend significantly on social methods of self-check and self-correction. But if our ability to evaluate experts in a domain is an extension of our own abilities in that domain, then what we're doing is a kind of cognitive bootstrapping. Even if we find cognitive bootstrapping acceptable, it makes us extremely vulnerable. The corrective process is dependent on our initial abilities being, in general, more right than wrong. But if it is the other way around — if we are more deeply flawed in our abilities, then we will put our trust in the wrong experts and be reinforced in the wrong ways. Some examples: if your moral compass is such that you find various pick-up artist instructors admirable and select as your moral guides people who encourage you to lie and manipulate others to achieve greater quantities of sexual success, then social corroboration and discorroboration effects will only make you more of a sociopath.

I call this a "personal" echo chamber because each such echo chamber is generated around a single individual by the exercises of their own abilities. Contrast this with the social echo chambers described by Jamieson and Cappella. They describe a complex process of balkanization and polarization, by which Limbaugh and Fox News isolate their listeners within an intellectual enclave. This involves, among other things, actively and forcefully alienating their listeners from other sources of information through constant claims about the corruption of the mainstream media and through the creation of a distinctive shared language. These predatory social echo chambers typically originate from a malicious epistemic manipulation by external agents (Jamieson and Cappella 2008, 126-236). The

personal echo chambers I've described arise, not maliciously, but inevitably from the structure of cognitive islands. Certainly, personal echo chambers and social echo chambers can reinforce each other, but they are conceptually separable processes.

And I call them "runaway" because of the lack of a safety net. Again, social echo chambers, such as those described by Jamieson and Cappella, do not only occur on cognitive islands and so are not inherently runaway. In Jamieson and Cappella's study, the right-wing echo chamber made many claims about many unsubtle and non-isolated domains, including climate change, economics, and medicine. Those claims are open to some forms of inexpert testing. Here, we need some other explanation of why they persist in the face of obvious contrary evidence. According to Jamieson and Cappella, the best explanation of the persistence of their social echo chambers is various forms of epistemic vice — both an exaggeration of in-group/out-group differences and informational distortion on the part of the echo chamber's creators, and a preference for comfort and familiarity over critical thinking (75-8,177-242). Similarly, many have described the technologically mediated social epistemic phenomenon called 'filter bubbles', which occur when an algorithmic news filtering system, like Facebook or Google, presents news to a user based on what the viewer has previously clicked on or 'liked'. This creates a bias in a user's newsfeed, often without that user's awareness or consent (Pariser 2011; Watson 2015; Miller and Record 2013). Again, these accounts usually require epistemically irresponsible behavior on the part of institutions and individuals.[9]

But the process I've described for runaway personal echo chambers requires no such epistemic vice. Morality, let us suppose, is a cognitive island. One places one's trust in experts

[9] I will offer an extended discussion of the relationship between Jamieson and Cappella's views, the filter bubble view in a future paper.

through an exercise of one's own moral sensibility. If that sensibility is deeply flawed, then one will pick fellows, interlocutors, and other purportedly reliable advice-givers in a deeply flawed way. By following seemingly legitimate processes of corroboration and self-checking, one will only amplify one's flaws. We might even posit a technological acceleration of this process. The more power and autonomy one has to sort through a wide variety of moral experts via social media and choice of political websites, the more quickly one will be carried away by one's personal echo chamber. Furthermore, if similarly-minded people find each other, then their runaway personal echo chambers may reinforce each other, as corroboration effects feedback on confidence levels, and vice versa. In that case, those people will naturally generate a runaway interpersonal echo chamber with no epistemic vice in sight. Compare this to Jamieson and Cappella's predatory social echo chambers, which are particular and contingent social phenomenon of which one may or may not happen to run afoul. But the runaway echo chambers I've described here are a natural and inevitable occurrence, which arise from our limitations as epistemic agents combined with the dangerous conditions of cognitive islands. On any true cognitive island, runaway echo chambers should arise around any individual with flawed abilities or starting beliefs, or a set of individuals with similarly flawed abilities or beliefs.

Such runaway echo chambers are, in some sense, just an especially sharp-edged case of a more general problem. If all my senses are systematically misled, then I have no hope of using them for reflective self-improvement. As Fred Dretske argued, any good epistemic theory ought to leave room for the possibility of beings like us getting it completely wrong. The brain in the vat will arrive at all the wrong beliefs, says Dretske, even if they follow all the best epistemic procedures (Dretske 2000, 593-4). Similarly, if somebody has all the wrong epistemic procedures and beliefs, then their use of experts will lead them further astray in

any domain, not just on cognitive islands. So, even on the cognitive mainland, the same sort of trap is available, if one starts completely off the rails. But the problem of skewed self-reinforcement is ameliorated to a significant degree on the cognitive mainland. The more obvious the output and the more interconnections to obvious domains, the more safety nets one has. On the cognitive mainland, if one particular cognitive ability is entirely errant, that flaw can be caught by any of one's other cognitive abilities also on that mainland or their social extensions. But on a cognitive island, the relevant cognitive abilities are the sole ultimate check on their own well-functioning. Thus, cognitive islands drastically increase our vulnerability to epistemic tragedy.

Prospects for amelioration

The effect I'm describing is distinct from, but compatible with, the belief polarization effect described by Thomas Kelly. According to Kelly, there is good empirical evidence to suggest that we subject beliefs we disagree with to greater scrutiny than beliefs we agree with. If this is so, we are likely to generate good counterarguments to beliefs we disagree with. Rational agents will then naturally be led to increase their certainty in their own beliefs, and disagreeing agents will be increasingly likely to polarize their beliefs (Kelly 2008). But notice, as Kelly suggests, that this form of belief polarization may be ameliorated through simple means: being aware of that effect and, by force of will, subjecting one's own views to equivalent levels of scrutiny. Is there such a simple solution to runaway echo chambers? I think not. Runaway echo chambers do not emanate from such a simple and curable form of

neglect. So long as we are dependent on experts, we must evaluate them, and on cognitive islands, that evaluative process is essentially vulnerable to self-entrapment.

If what generates the trap is trusting experts via an exercise of our own abilities, perhaps we should abandon self-trust in the process of selecting experts. Perhaps we should open ourselves up to trusting everybody. Baumgaertner has suggested that perhaps one way to escape from echo chambers is simply to dispense with filtering our sources altogether. Baumgaertner considers the epistemic principle of Impartiality, under which epistemic agents would permit themselves to epistemically encounter, and be influenced, by all agents, and not just ones that have been filtered for reliability and expertise. Unfortunately, Baumgaertner's results are quite disheartening. In his computer models, following Impartiality yields a positive result only if the model contains the assumption that agent confidence in their beliefs is correlated with agent correctness — that is, if we assume that purported experts' expressed confidence in their beliefs is always a good indicator of their reliability. If we remove that wildly implausible assumption, we see far worse epistemic results — rather than promoting true beliefs, the process only promotes beliefs that happen to be deeply entrenched (Baumgaertner 2014). Still, Baumgaertner retains hope that some version of Impartiality can be made to work.

Given the perniciousness of runaway echo chambers, if there are any genuine cognitive islands, then we may need to consider such solutions, no matter how uncomfortably radical. Since the problem of these runaway effects arises from pernicious self-reinforcement, our only hope of checking the problem may be to counterbalance our autonomous and self-trusting information filtering with a process relatively untainted with our own cognitive abilities. But it isn't clear that we would ever be capable of sufficiently distancing ourselves from our own cognitive abilities to actually do so. Since runaway echo chambers originate in

one's own cognitive flaws, so long as one trusts one's own abilities more than those of others and uses those abilities to sort out who to trust, then the effect should take hold. The alternative is a kind of epistemic egolessness, which seems a terribly heavy price to pay, if not an outright impossibility. Perhaps some solution between these two alternatives may be found. But until then, a sadder conclusion seems to prevail: that our epistemic finitude, combined with the peculiar conditions of cognitive islands, makes runaway echo chambers inevitable and inescapable.[10]

Bibliography

Baier, Annette. 1986. Trust and anti-trust. *Ethics* 96 (2): 231-260.

Baumgaertner, Bert. 2014. Yes, no, maybe so: A veritistic approach to echo chambers using a trichotomous belief model. *Synthese* 191 (11): 2549-2569.

Budd, Malcolm. 2003. The acquaintance principle. *The British Journal of Aesthetics* 43 (4): 386-392.

Cholbi, Michael. 2007. Moral expertise and the credentials problem. *Ethical Theory and Moral Practice* 10 (4): 323-334.

Coady, David. 2006. When experts disagree. *Episteme: A Journal of Social Epistemology* 3 (1-2): 68-79.

Dretske, Fred. 2000. Entitlement: Epistemic rights without epistemic duties? *Philosophy and Phenomenological Research* 60 (3): 591-606.

Driver, Julia. 2006. Autonomy and the asymmetry problem for moral expertise. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 128 (3): 619-644.

Goldberg, Sanford. 2011. If that were true I would have heard about it by now. In *The Oxford Handbook of Social Epistemology.* Oxford: Oxford University Press.

Goldman, Alvin I. 2001. Experts: Which ones should you trust? *Philosophy and Phenomenological Research* 63 (1): 85-110.

Hardwig, John. 1985. Epistemic dependence. *The Journal of Philosophy* 82 (7): 335-349.

Hopkins, Robert. 2011. How to be a pessimist about aesthetic testimony. *Journal of Philosophy* 108 (3): 138-157.

Jamieson, Kathleen Hall and Joseph N Cappella. 2008. *Echo Chamber: Rush Limbaugh and the Conservative Media Establishment.* Oxford: Oxford University Press.

Jones, Karen. 1999. Second-Hand moral knowledge. *Journal of Philosophy* 96 (2): 55-78.

———. 2012. Trustworthiness. *Ethics* 123 (1): 61-85.

[10] I'd like to thank Shannon Mussett, Eric Stencil, Mary Beth Willard, Laura Guerrero, Anthony Cross, Thomas Hurka, Elijah Millgram, and the anonymous reviewers for their insightful commentary on this paper.

Kelly, Thomas. 2005. The epistemic significance of disagreement. *Oxford Studies in Epistemology* 1 (167-196).

———. 2008. Disagreement, dogmatism, and belief polarization. *The Journal of Philosophy* 105 (10): 611-633.

Kitcher, Philip. 1993. *The Advancement of Science.* New York: Oxford University Press.

LaBarge, Scott. 1997. Socrates and the recognition of experts. *Apeiron: A Journal for Ancient Philosophy and Science* 30 (4): 51-62.

Lackey, Jennifer. 2013. Disagreement and belief dependence: Why numbers matter. In *The Epistemology of Disagreement: New Essays.* Eds. Jennifer Lackey and David Christensen. Oxford: Oxford University Press.

McGrath, Sarah. 2008. Moral disagreement and moral expertise. *Oxford Studies in Metaethics* 3:87-108.

———. 2011. Skepticism about moral expertise as a puzzle for moral realism. *Journal of Philosophy* 108 (3): 111-137.

Meskin, Aaron. 2007. Solving the puzzle of aesthetic testimony. In *Knowing Art.* Ed. Matthew Kieran and Dominic McIver Lopes. Dordrecht: Springer.

Miller, Boaz and Isaac Record. 2013. Justified belief in a digital age: On the epistemic implications of secret internet technologies. *Episteme* 10 (02): 117-134.

Millgram, Elijah. 2015. *The Great Endarkenment: Philosophy for An Age of Hyperspecialization.* Oxford: Oxford University Press.

Nguyen, C. Thi. 2010. Autonomy, understanding, and moral disagreement. *Philosophical Topics* 38 (2): 111-129.

———. 2011. An ethics of uncertainty.

Nguyen, C Thi. 2017. The uses of aesthetic testimony. *The British Journal of Aesthetics* 57 (1): 19-36.

Nickel, Philip. 2001. Moral testimony and its authority. *Ethical Theory and Moral Practice* 4 (3): 253-266.

Pariser, Eli. 2011. *The Filter Bubble: What the Internet Is Hiding From You.* Penguin UK.

Watson, Jamie Carlin. 2015. Filter bubbles and the public use of reason: Applying epistemology to the newsfeed. In *Social Epistemology and Technology: Toward Public Self-Awareness Regarding Technological Mediation.* Ed. Frank Scalambrino. London: Rowman & Littlefield.