

Suspiciously Convenient Beliefs and the Pathologies of (Epistemological) Ideal Theory

Alex Worsnip

Forthcoming in *Midwest Studies in Philosophy*, vol. 47
(issue on the Genealogy of Belief, ed. Yuval Avnur)

1. The phenomenon

Public life abounds with examples of people whose beliefs – especially their political or politically-relevant beliefs – seem *suspiciously convenient*.¹ Consider, for example:

Libertarian Billionaire. Connor is a billionaire who would stand to financially benefit greatly from the elimination of income taxes. Connor is interested in economic policy and arguments about its moral foundations, and he has become convinced of the view that government taxation of income is theft, and is thus illegitimate.

Judicial Originalist. Amy is a Supreme Court Justice who holds strongly conservative moral and political views. However, Amy also holds the view that when judges engage in statutory interpretation, they should set their personal moral and political views aside and should simply adhere to the plain original meaning of the statute. Yet when Amy arrives at her interpretations of the plain original meaning of the statutes, they consistently and reliably support rulings that happen to align with her own personal moral and political views.

Socialist Protectionist. Bernie is a self-identified socialist politician with deep ties to and emotional investment in the labor movement. Labor unions tend to support protectionist trade policies on the part of the USA, on the grounds that they protect the jobs of their members. But many economists argue that protectionist trade policies on the part of the USA also ossify extreme poverty in developing countries, and that the dismantling of these policies would help billions of the world's poorest people out of poverty. Bernie would not like to think of himself as contributing to the ossification of extreme poverty in developing countries. Fortunately for him, however, he thinks that these economists' arguments are bunk: liberalizing the USA's trade policies, he thinks, would hurt not only domestic American workers but *also* those living in poverty in developing countries.

For an example that strikes closer to home for philosophers, consider also:

I'm grateful to Yuval Avnur and Genae Matthews for very helpful comments on this paper, and to Corey Cusimano, Ava Geenen, Ram Neta, Jim Pryor, and Conner Schultz for helpful conversations about the topics discussed in it. I'm also very grateful to audiences at Johns Hopkins University, the Oxford Epistemology Group, the 2022 UNC prospective students event, a workshop on irrelevant influences on belief at Scripps College, and the University of Toronto.

¹ The term 'suspiciously convenient belief' is also used by Levy (2020), though he uses it in a narrower way to refer to beliefs *about the consequences of actions* that align suspiciously with one's *moral* beliefs. Of the below examples, only Socialist Protectionist looks to (arguably) be an example of this narrower phenomenon.

Religious Philosopher. Jennifer is a philosopher working in ethics and epistemology. She grew up in, and is deeply emotionally attached to, a strongly religious tradition, and maintains her faith now. In her work in ethics, she defends moral views that broadly align with the teachings of her religious tradition, though on purely secular grounds that don't appeal to religious texts or divine authority. In her work in epistemology, she defends the view that it's rational to be steadfast in one's views (including religious views) in the face of peer disagreement, and that the discovery of "irrelevant influences" on one's belief is, in itself, epistemically irrelevant.

Let's make some stipulations about these cases, so as to get a fix on the phenomenon of maximal philosophical interest. First, let's stipulate that these agents are sincere in their professed beliefs. So, for example, Connor really does believe that the taxation of income is unjustified theft; that's not just something he *says* to try to advocate for the policies that would financially favor him. Second, let's stipulate that, at least at a conscious level, the characters all take themselves to have good arguments or evidence for their beliefs; arguments and evidence of a sort that, they think, even those who don't share their particular interests and desires should be able to appreciate. So, for example, Connor takes his belief that income taxation is theft to be supported by the best arguments on the topic, arguments that even non-billionaires should (in his view) be able to appreciate. If asked to justify his belief, he wouldn't just say "oh, I believe that because it serves my interests"; rather, he'd produce some arguments that he genuinely takes to justify his belief. And third, let's stipulate that the arguments and evidence that they would produce for their beliefs, if asked, are fairly sophisticated and at least not *transparently* fallacious or stupid—though they may (or may not) fall short of fully justifying their conclusions.

Despite these stipulations, there still seems to be something suspicious about the beliefs in all the cases just described. Moreover, this suspiciousness is not just a matter of their seeming false, or its seeming that they are not in fact justified by strong arguments. (Note that the cases just mentioned cross the political spectrum, and all the views mentioned have serious adherents with sophisticated arguments for them.) Rather, the suspiciousness resides in the way that the beliefs in question seem so *convenient* for particular individuals that hold them, in the sense – roughly – that they align so neatly with their interests and desires. This suggests – as I'll substantiate in more detail in §1 – that it's at least pretty likely that the interests that the individuals in the above cases are not just following the arguments and evidence where they lead – even if that is what it seems to them that they are doing. Rather, their interests and desires are (perhaps subtly and subconsciously) biasing their inquiries. In doing so, these desires and interests play an important causal role in the production of the beliefs they end up holding. Of course, we can't be *sure* that this is what's going on in any of our cases, just from how I described them. But it's at least a hypothesis we should give some serious credence to in cases of this kind.

Suspiciously convenient beliefs are an instance of a broader category of cases that have attracted a great deal of attention from philosophers in recent years, namely beliefs that have been significantly shaped (or, at least, that we have good reason to suspect have been significantly shaped) by

“irrelevant influences,” considerations that don’t bear on their truth.² In the case of suspiciously convenient beliefs (and in some other cases, too), the suspected irrelevant influence is the believing agent’s interests and desires. (Notice that Religious Philosopher, somewhat ironically, is a case where the agent’s beliefs about – among other things – the epistemology of irrelevant influences are themselves plausibly shaped by irrelevant influences.) In turn, the category of irrelevantly-influenced beliefs belongs to a still broader, and still more widely-discussed, category: cases where we have “higher-order evidence” that suggests that our beliefs are rationally suboptimal, or were produced in a rationally suboptimal way.³

Many (though not all) of the issues I’ll be discussing about suspiciously convenient beliefs in this paper arise in very similar ways with other cases of irrelevant influences, and indeed with other cases of higher-order evidence more broadly. I’ve chosen to focus more specifically on suspiciously convenient beliefs as a case study here for two reasons. First, such cases are relatively underdiscussed in the contemporary analytic epistemology literature.⁴ Though my ultimate view is that the suspicious convenience of a belief is higher-order evidence (and thus, that it has epistemic significance just if higher-order evidence more generally has epistemic significance), it (as we’ll see) takes some work to show why this is so, and it is worth systematically spelling out the reasons why.

Moreover, and second, suspiciously convenient beliefs are among the cases of irrelevant influences (and higher-order evidence) that are most interesting from a socio-political perspective. Though much of the interest in the topics of irrelevant influences and higher-order evidence derives from these politically charged cases, there’s a tendency – albeit not an exceptionless one – in the literature to flee from such cases when getting into the details of epistemological theorizing. Allegedly to simplify or avoid unnecessary distraction, they are replaced by maximally apolitical examples involving the splitting of restaurant checks,⁵ the effects of grad school attendance choices on one’s beliefs about the analytic-synthetic distinction,⁶ or imaginary drugs that interfere with one’s ability to perform deductive inferences.⁷ While there’s nothing wrong with considering these cases alongside

² Key papers include, among others, White (2010), Schoenfield (2014), Avnur & Scott-Kakures (2015) and Vavova (2018).

³ See fn. 9 for some of the key literature. It’s worth noting that evidence that suggests that one’s beliefs are irrational or were produced in a rationally suboptimal way is only one kind of higher-order evidence. Most obviously, evidence that one’s beliefs *are* rational is also higher-order evidence. But also, at least by most characterizations of higher-order evidence, evidence about what one’s evidence supports – which may not concern any existing belief one has, nor the functioning of one’s rational capacities, also counts as higher-order evidence. (For an attempt to give a more precise and overarching characterization of higher-order evidence, see Chen & Worsnip (forthcoming).) The issues about the rational impact of these other kinds of higher-order evidence are in some respects quite different from the issues about the rational impact of evidence that suggests that one’s beliefs are irrational or irrationally produced, and some confusion has been generated in the literature by the fact that some (tacitly) use ‘higher-order evidence’ *only* to refer to the latter. Since suspicious convenience is the kind of higher-order evidence that suggests one’s beliefs are irrational or irrationally produced, though, it’ll be the rational import of *this* kind of higher-order evidence that’s of relevance in this paper.

⁴ As noted in n. 1 above, Levy (2020) is a partial exception, though the phenomenon he discusses is importantly narrower. There is also a tradition of discussion of related phenomena in the continental philosophy literature, often taking inspiration from historical figures such as Marx, Nietzsche, and Freud, and often under the name “the hermeneutics of suspicion.” See, e.g., Ricoeur (1970); Leiter (2004).

⁵ This case was introduced by Christensen (2007), and has been discussed by countless others.

⁶ This case was introduced by Cohen (2000). Notably, Cohen *does* discuss it alongside political examples, but it’s his grad school case that has become a central touchpoint in the literature (see, e.g., White 2010, Schoenfield 2014, Vavova 2018, and countless others).

⁷ This case was introduced by Christensen (2010), and has been discussed by countless others.

others, I think the defenders of views about higher-order evidence should also be willing to test their theories on cases with more political bite, and to face up to the potentially troubling results their theories generate in such cases.⁸ This paper aims, among other things, to shift the debate a little closer toward this approach.

2. The question

A (perhaps *the*) central question in the general debate about higher-order evidence is this: when you get evidence that your belief or credence in some proposition p is (or is likely to be) irrational, should this, in itself, have an impact not just on your belief or credence in the higher-order proposition that your belief or credence in p is irrational, but also on your belief or credence in the first-order proposition p itself?⁹ Assuming (plausibly; I'll explain further in §3.3) that evidence that your beliefs are biased by your desires and interests is evidence that they're irrational, this question can be fairly straightforwardly applied to suspiciously convenient beliefs in particular. The more specific question, then, is this: suppose you realize that *your* present beliefs line up conveniently with your desires and interests, and are thus likely due to bias. Should that realization, in and of itself, usually make you less confident in the content of those beliefs?¹⁰

For short, let's call the affirmative answer to these questions Significance, and the negative answer Anti-Significance. For maximal clarity, I'll sometimes call the view that higher-order evidence in general is epistemically significant (with respect to one's first-order beliefs) General Significance, and the view that the suspicious convenience of one's beliefs in particular is epistemically significant Specific Significance (and likewise, *mutatis mutandis*, for Anti-Significance, General and Specific).

⁸ What I'm saying here chimes with Srinivasan (2020) and Johnson King (2022), who make similar points about the methodology of the debate between internalists and externalists (though reach opposing substantive conclusions about this debate). They dub their preferred approach to epistemology "radical," though I don't think one needs to be a political radical, or even to be broadly of the left, to appreciate the methodological point.

⁹ Those who argue that it should include Christensen (2010), Horowitz (2014), Avnur & Scott-Kakures (2015), Schoenfeld (2018), González de Prado (2020), and Lord & Sylvan (2021). Those who argue that it shouldn't, or at least shouldn't always, include Titelbaum (2015), Tal (2020), and *perhaps* Lasonen-Aarnio (2014) and Weatherson (2019: chs. 7-8), though the latter cases are less clear (see ns. 27 and 53 on the complications of interpreting Weatherson's view in particular).

¹⁰ I'll assume throughout that we have both coarse-grained, outright doxastic attitudes (i.e., outright belief, suspension of judgment, and outright disbelief) and fine-grained, graded doxastic attitudes (i.e. credences, or levels of confidence). When I talk about becoming less confident in your belief, I mean that your credence in the proposition that you believe (or previously believed) goes down. In some but not all cases, the loss of confidence may cause, constitute, or otherwise go along with your ceasing to outright believe the proposition entirely. (I don't assume any particular account of the relationship between outright belief and credence.) I'll also assume that we can employ both an on-off, outright notion of *justification*, where an (outright) belief is either justified or not, and a graded notion of justification, where an (outright) belief can be *more* or *less* justified. One's credences, I assume, generally ought to track one's graded justification: as one loses justification for believing p incrementally, one's credence for p should diminish.

It will often be easier to talk in term of the graded notions, since fine-grained differences in evidential support will make a difference to one's fine-grained justification for outright believing, and to the credences one should have, but not always to what one is on-off justified in believing. So, if your evidential position with respect to your belief worsens somewhat, so that you lose *some* graded justification for your belief but not all of it, this may or may not mean that you're no longer on-off justified in believing it (depending on exactly how much graded justification you had before, exactly how much you lost, and exactly how much graded justification is needed to be on-off justified in believing in your present context). But it will always mean that your credence or level of confidence should go down.

My own sympathies are squarely with Significance, both general and specific. In the remainder of this paper, I will do several things. First (§3), I'll spell out in detail the *prima facie* case for Specific Significance – that is, for the epistemic significance of suspicious convenience. Second (§4), I'll consider how this *prima facie* case might be resisted. I'll suggest that the only viable way to do so is to reject General Significance, and so I'll examine some of the strategies and lines of thought that have been used to argue for General Anti-Significance. I'll argue that these strategies tend to presuppose a form of methodological “ideal theory,” in a sense to be made precise.¹¹ And I'll argue that this kind of ideal theory is pernicious and distorting, and that the case of suspiciously convenient beliefs brings this out particularly acutely. Finally (§5), I'll conclude with some speculations on possible upshots for norms of public discourse.

3. The *prima facie* case for Specific Significance

Though I've said that the debate about suspiciously convenient beliefs is an application of more general debates about higher-order evidence, it takes some work to fully explain why this is so. To do so, we first need to say in exactly what sense suspiciously convenient beliefs are convenient, and in what sense this makes them suspicious. That will put us in a better position to see how the suspicious convenience of a belief is higher-order evidence – and thus, how it is epistemically significant if higher-order evidence more generally is epistemically significant. This will yield a *prima facie* case for Specific Significance, one that employs General Significance as a premise (alongside others).

3.1 *Convenient how?*

As I've already said, suspiciously convenient beliefs are *convenient* in the sense, roughly, that they align so neatly with the believer's interests and/or desires. There are two (distinct but compatible) senses in which this can be so, and which of the two provides the more appropriate model may vary from case to case. The first possibility is that given one's interests and/or desires, the *truth* of the belief in question would be convenient for the believer. For example, there's plausibly a good sense in which if it's *true* that all taxation is unjust, that's convenient for the billionaire's interests (since he then doesn't have to face down a tradeoff between justice and his own interests). And plausibly, there's a good sense in which if steadfastness is justified, that's convenient for the religious philosopher's desires: in particular, her desire to hang on to her religious beliefs.

One might push back on whether this description fits the cases, however. For example, one might suggest that the billionaire doesn't need it to actually be *true* that taxation is unjust – he just needs taxes to be kept low. This end will be served by others *believing* that taxation is unjust, regardless of whether this is true or not. This, however, leads us to the second possibility: that given one's interests and/or desires, it's psychologically comfortable to *hold* the belief in question: that is, to *think*

¹¹ The terms ‘ideal theory’ and ‘non-ideal theory’ are most familiar from political philosophy (see Valentini 2012 for an overview) and ethics (e.g., Murphy 2003). But in recent years they have been increasingly discussed as methodological approaches to epistemology (e.g., Saint-Croix 2018; Carr 2021; McKenna 2023; Greco 2023). In all these contexts, the terms are used with a bewildering variety of non-equivalent meanings.

that it is true. While the billionaire could dishonestly advocate for lower taxation while believing that justice requires higher taxation, this is not psychologically comfortable for him: it's much nicer for him to believe that he has justice on his side. Similarly, the Supreme Court Justice could just consciously see herself as lying about what the law is in order to advance her own normative agenda. But plausibly, judges don't like to think of themselves as doing this, and so it's much more psychologically convenient for her to convince herself that the law really does say what she says it does.

Also under the broad umbrella of this second possibility, it can be that it's psychologically comfortable to hold a belief given one's *other* beliefs and a desire to avoid facing down difficult tradeoffs.¹² For example, plausibly, it's psychologically uncomfortable for the politician in the Socialist Protectionist case to face down a tradeoff between what benefits domestic workers and what benefits the global poor, and given that the politician believes that protectionism benefits domestic workers, believing that it *also* benefits (rather than harms) the global poor allows the politician to avoid perceiving or having to adjudicate such a tradeoff. The same may be true of those who oppose significant large-scale government intervention in the market and also deny the reality of catastrophic phenomena that would, if real, apparently be solvable only with such intervention, such as large-scale anthropogenic climate change. Again, such people face no tradeoff between the evil of regulation and the dangers of climate change if there are no such dangers.

3.2 *Suspicious how?*

So much for what it is for a belief to be convenient in the relevant sense. How are beliefs that are convenient in this sense *suspicious*? The short answer is that their convenience lends plausibility to the hypothesis that the interests and desires for which they are convenient played a role in generating the belief. To see why this is so, let's first examine how interests and desires *can* play a role in generating beliefs. There are (at least) two ways that this can happen, roughly corresponding to the two different kinds of convenience just explored. The first is "wishful thinking," where a desire *that p be true* plays a role in generating the belief that *p*. The second is "motivated reasoning," where a desire *to believe p* (often called a "directional goal" in the psychological literature) plays a role in generating the belief.¹³ In what follows, I'll work on the assumption that motivated reasoning is generally the more salient hypothesis, but everything I say would go through for wishful thinking too.

Two points about motivated reasoning bear stressing for our purposes. First, the influence of one's desires on one's beliefs in motivated reasoning is typically indirect. We usually can't just desire to believe something and then immediately believe it as a result.¹⁴ Rather, the causal influence goes via

¹² For a classic discussion of the psychology of tradeoff avoidance (which was first explored in the context of international politics), see Jervis (1976: 128-143). Jervis calls the tendency to try to see *all* considerations as favoring the same side of an issue, such that no tradeoffs arise, "belief overkill" (see also Baron 2009 for a more recent discussion).

¹³ Kelly (2002: 176-7) stresses the importance of a distinction very close to this one. The classic discussion of motivated reasoning is Kunda (1990). For a more recent overview, see Kahan (2016). For philosophical discussions see, e.g., Avnur & Scott-Kakures (2015), Carter & McKenna (2020), Williams (forthcoming).

¹⁴ This is stressed both in the psychological literature on motivated reasoning (e.g. Kunda (1990: 482) and in the philosophical literature on "believing at will" (e.g. Williams 1973; Winters 1979; Hieronymi 2006).

a kind of biased or selective gathering and processing of evidence.¹⁵ There are a few different ways in which this can go. First, we can be selective in which evidence we gather or attend to in the first place, seeking out evidence and arguments likely to confirm our desired views, and avoiding those likely to disconfirm them. Second, we can overweight some pieces of evidence (those seeming to confirm our views) and underweight or dismiss others (those seeming to disconfirm them). And third, we can conduct biased assessments of what any one piece of evidence supports – taking it to support our desired views while ignoring or dismissing other explanations of the evidence that favor other, competing views – and thus miscalculate the *valence* of our evidence – whether it tells for or against some conclusion. Once this biased gathering and processing of evidence is complete, we then form the desired belief on the basis of the evidence (and judgments about the probative value of that evidence) that results. In short, the belief-forming process in motivated reasoning still works by considering and responding to evidence, but in a biased way.

Second, and relatedly, the influence of our desires on our beliefs in motivated reasoning is typically *subconscious* and unavailable to introspection.¹⁶ This means that the experience of someone who has engaged in motivated reasoning is typically as of having followed the arguments and evidence where they lead in an unbiased manner; there is an “illusion of objectivity.” Moreover, such a person will be able to produce a (putative) justification for their view, appealing to evidence and arguments, and can be expected to sincerely expect that this justification could persuade a dispassionate observer.

These two points are important for our topic because, as I stipulated, the subjects in our opening vignettes do sincerely *think* that the arguments and evidence support their views, and that neutral parties should be able to accept them. The indirect and subconscious nature of motivated reasoning, and its unavailability to introspection, makes clear that this is not at all inconsistent with the possibility that *in fact*, the beliefs in question were produced by motivated reasoning (and are not in fact supported by the total available evidence).

With this summary of motivated reasoning in place, we’re in a better position to see why the convenience of beliefs lends plausibility to the hypothesis that they were generated by motivated reasoning (or similar processes such as wishful thinking). First, there is a large body of empirical evidence that motivated reasoning is extremely prevalent – and moreover, that it is prevalent across a wide range of topics,¹⁷ and agents with a range of political ideologies, levels of education, and levels of cognitive sophistication.¹⁸ This constitutes background base rate data that raises the probability of the hypothesis that any one of our individual political beliefs is due to motivated reasoning. Moreover, the hypothesis is particularly plausible for convenient beliefs, because they are by their very nature exactly those beliefs that we’d expect the believed agent to hold if motivated reasoning were operative.

¹⁵ Kunda (1990: 483, 493-5).

¹⁶ Kunda (1990: 483). For a more general discussion of the unavailability of motivation to introspection, see e.g. Wilson (2002, esp. ch. 5).

¹⁷ For example, beliefs about the efficacy of the death penalty (Lord, Ross & Lepper 1979); beliefs about sports events (Gilovich 1983); beliefs about the health effects of smoking (Kassarijan & Cohen 1965) and of caffeine (Kunda 1987); beliefs about the seriousness of medical conditions (Ditto, Jemmott & Darley 1988); beliefs about climate change (Kahan, Jenkins-Smith & Braman 2011); and a whole gamut of beliefs about one’s own abilities, traits and behavior, including moral evaluations of behavior one has engaged in in the past (Ross, MacFarland & Fletcher 1981; Kunda 1987; Kunda & Sanitioso 1989; Sanitioso et al 1990).

¹⁸ Lord, Ross & Lepper 1979; Ryan & Aziz 2021.

By contrast, the hypothesis that one's desires and interests did not play a role in generating one's convenient belief posits a (relatively) implausible coincidence, one whereby unbiased reasoning just happened to lead the agent to exactly the belief that we would expect them to have were motivated reasoning operative. Of course, it is not *impossible* that this is exactly what has happened – coincidences do happen, and for any suspiciously convenient belief, there's some chance it was in fact produced by unbiased reasoning. But we don't need the coincidence to be *impossible* in order for convenience to be *pretty* good evidence that motivated reasoning was at work – it just needs to be relatively unlikely. Moreover, the (would-be) coincidence gets greater and greater, and less and less plausible, the more that one's suspiciously convenient beliefs form a pattern. For example, the justice in Judicial Originalist doesn't just have *one* belief about what the law is that aligns neatly with her substantive normative agenda – she has a pattern of continually arriving at such beliefs. For all these beliefs to be the product of unbiased reasoning would be a massive unexplained coincidence. And the would-be coincidence gets still greater and still less plausible when there are *cross-agent* patterns of suspiciously convenient beliefs—say, when *multiple* Supreme Court Justices keep arriving at beliefs about what the law is that align neatly with their (perhaps differing) substantive normative agendas. This makes it even more likely that motivated reasoning (or something similar) is at work—likely, on *all* sides of the relevant dispute—and in the absence of any special reason to think that some particular individual Justice is immune to such bias, that further strengthens the case for thinking that this individual is engaging in motivated reasoning. Thus, I conclude, the convenience of a belief is evidence that it was produced by motivated reasoning (or some related process) wherein it has been effectively reverse-engineered to align with the agent's agenda (or similar). It's in this that its suspiciousness resides.

The appeal to motivated reasoning helps us to meet a challenge that Roger White (2010) poses for the view that irrelevant influences are epistemically significant more broadly. White asks why considerations of irrelevant influences have any epistemic significance beyond that of peer disagreement.¹⁹ He asks us to imagine two scenarios (577): one in which there is significant disagreement about whether *p* but it is randomly distributed, and one where there is significant disagreement about whether *p* but the distribution of opinion is strongly correlated with some factor about individuals' backgrounds or interests. He submits that the disagreement provides no stronger reason to reduce confidence in the latter case than the former. This, I think we can now see, is a mistake—at least for many such correlations. The key point is that such correlations (at least often) make the hypothesis that motivated reasoning was at work—indeed, that *both* sides' beliefs were produced by motivated reasoning—more plausible.²⁰ And, as I'll argue shortly, that makes it plausible

¹⁹ Cf. also Avnur & Scott-Kakures's (2015) response to White.

²⁰ Incidentally, the psychological hypotheses that White himself canvasses (and dismisses) as possible explanations of the correlation in question do not come close to the most sophisticated and plausible such hypotheses. In the case of students who end up concurring with the party line about a philosophical issue at their home institution, he writes: "It is not at all clear why philosophers should tend to believe as their professors do. [...] Perhaps students were bullied into accepting the local orthodoxy. Or perhaps they uncritically looked up to their teachers as gurus. But these extreme stories are not realistic, and surely the worry [about irrelevant influences] doesn't depend on anything of the sort" (578). But even if it's rare for things to go quite that straightforwardly, there are myriad more sophisticated and plausible hypotheses available. For example, many students are likely to have a (subconscious?) desire to fit in with the local way of thinking, not to be branded as obtuse by other students and professors, and so on, and these desires can provide them with directional goals to believe the local orthodoxy, which (subconsciously) sets off a process of motivated reasoning aimed at finding arguments for that view.

that both sides' beliefs (or levels of confidence in those beliefs) are irrational. By contrast, mere (randomly distributed) disagreement indicates at most that *one* party to the disagreement is being irrational,²¹ and not necessarily in a way involving motivated reasoning.

3.3 *Why suspicious convenience is higher-order evidence*

We have just seen that the convenience of a belief is evidence that it was produced by motivated reasoning. The next step is to show that this, in turn, is evidence that the belief in question is irrational (which, on standard characterizations of higher-order evidence, makes it a variety of higher-order evidence).

The core idea here is that evidence that your belief was produced by motivated reasoning is at least some evidence that it is not as well-supported by the arguments and evidence as you take, or previously took, it to be. Why should that be so? The answer is that, as noted earlier, motivated reasoning typically involves overweighting the evidence that *prima facie* favors one's belief, underweighting the evidence that *prima facie* tells against it, and/or misevaluating the valence of particular elements of one's evidence (i.e., taking them to be evidence for one's view when they are actually evidence against it).²² All of these possibilities ultimately lead to a misevaluation of what one's evidence supports, and thus are liable to lead to beliefs that are not in fact supported by one's evidence. And holding beliefs that are not in fact supported by one's evidence is a paradigmatic kind of irrationality. (The point is even clearer on a graded notion of belief, or credence: overestimating the force of one's evidence for p will lead one to be at least somewhat more confident in p than one's evidence justifies one in being, while underestimating the force of one's evidence for p will lead one to be at least somewhat less confident in p than the evidence justifies one in being.)

Now, it needs to be admitted that there are also *some* cases of motivated reasoning, or of something very close to it, that need not involve failing to believe what one's evidence supports. First, as noted above, motivated reasoning may sometimes work simply through one's engaging in selective, biased evidence-gathering. But engaging in selective, biased evidence-gathering *prima facie* seems compatible with going on to believe what one's (admittedly: selective, biased) evidence supports. It's a vexed issue, about which there's some debate, whether one is thereby *epistemically* irrational if one does this.²³ Second, it also seems like there might be cases where motivated reasoning causes one to

²¹ And, as many have observed (e.g. Kelly 2010), even that is true only if the "uniqueness thesis", according to which (very roughly) at most one of a number of incompatible attitudes toward p can be rational on a single, fixed body of evidence, is true.

²² Some subjectivists about evidential support would deny that there is any fixed, objective extent to which a given body of evidence supports a belief, independently of agents' priors. I disagree with this subjectivist view. But (as Yuval Avnir pointed out to me) even these subjectivists should be able to see that motivated reasoning involves the intrusion of truth-irrelevant motivations into one's reasoning.

²³ For example, Kelly (2002: 169-70) claims that there's no epistemic irrationality here. On his view (see also Kelly 2008: 622-3), questions about what evidence to gather are necessarily practical, not epistemic. However, Kelly's view is not uncontroversial. First, Salow (2018) argues that it should not be possible to deliberately manipulate your evidence to *genuinely* favor your desired beliefs, because so long as you are aware that you've manipulated your evidence in this way, this undercuts the would-be force of the evidence in question. Second, several philosophers have argued that choices about what evidence to gather *can* be epistemically evaluated: see, e.g., Worsnip (2019: 241-3), McWilliams (2021, forthcoming) and Flores & Woodard (2023).

be convinced by the arguments in favor of p and to be unconvinced by the arguments against p , but as a matter of luck, the arguments in favor of p actually *are* stronger than the arguments against p , and suffice to decisively support believing p . Again, it's not totally clear whether one should count as irrational if one believes p in such a case.²⁴

Despite these complexities, though, it remains the case that *many* cases of motivated reasoning are cases of irrationality. And given the opacity of the processes of motivated reasoning to the motivated reasoner, the cases that *are* instances of irrationality seem to be subjectively indistinguishable from the cases that aren't (if there are indeed cases that aren't). Thus, if you have evidence that you've been engaged in motivated reasoning, it seems that you have at least *some* evidence that you've been irrational too. So, if convenience is evidence of motivated reasoning – as I argued in the previous subsection that it is – it's also at least some evidence of irrationality.

Again, the foregoing reasoning helps to meet another challenge from White (2010), this one concerning why allegedly problematic irrelevant influences generate epistemic pressure to reduce confidence while others clearly don't (574, 582, 597). White points out that in many cases, a factor irrelevant to my belief's truth helps to explain why I hold the belief in a way that seems utterly unproblematic. For example, he points out that if he had never studied philosophy he would not believe that Hume was born in 1711; but even though the fact that he (Roger White) has studied philosophy is utterly irrelevant as to the question of when Hume was born, this doesn't seem to threaten the justification of his belief that Hume was born in 1711. But the explanation of why this "irrelevant influence" doesn't threaten belief, while others do, is simple. It's that the fact that White studied philosophy (and formed his belief that Hume was born in 1711) doesn't in any way suggest that his belief was produced by irrational process (such as motivated reasoning), and so isn't higher-order evidence that his belief is irrational. By contrast, the convenience of a belief *is* evidence that it was produced by motivated reasoning, and so *is* higher-order evidence that the belief is irrational. Similar things, I think, are true of the other irrelevant influences that do seem epistemically threatening.²⁵

²⁴ It might be thought that the right thing to say about this kind of case is that even though believing p is "propositionally" (or "ex ante") rational, one's token belief in p isn't "doxastically" (or "ex post") rational, because of its motivated nature (cf. Carter & Kenna 2020: 707). But it's not even totally clear that one's belief in this case is *doxastically* irrational (cf. White 2010: 598-600). On the standard view of doxastic rationality, one's belief is doxastically rational just if it is *based* on the reasons that make it propositionally rational. But it seems that in this case, one's belief *is* in a good sense based on the reasons that make it propositionally rational – after all, it's based *on the very arguments* that decisively support believing p . It's just that one's interests and desires are involved in the explanation of *why one finds those arguments compelling*, and thus bases one's belief on them. To hold that the belief is doxastically irrational, then, we'll need to either explain how one's belief doesn't *really* count as being based on the arguments that one takes to support it (perhaps instead "really" being based on one's interests and desires), or modify the standard account of doxastic (ir)rationality.

²⁵ Interestingly, White concedes that in a case he calls "Coin in the Head" (598), the presence of an irrelevant influence *does* seem to put epistemic pressure to reduce confidence in one's belief. In this case, one learns that one's moral beliefs about some matter are determined by a coin flip in one's head: if the coin flip comes up heads, one believes p ; if it comes up tails, one believes not- p . White's explanation of why the irrelevant influence in this special case is epistemically relevant is that in this case, it reveals that the subject "can't really be engaging in moral *reasoning* at all in this case" (599). But suppose we modify the case such that the coin flip doesn't *directly* determine one's moral belief, but sets off a process of reasoning that will eventuate in the belief in question: if it comes up heads, it will set off a process of reasoning that leads to one's believing p (perhaps: the *best available process of reasoning there is* that leads to that conclusion; if it comes up tails, it will do the same but for a process of reasoning that leads to one's believing not- p). Intuitively, the discovery of the coin in the head is just as disturbing in this modified case. So I think White's diagnosis that what makes the difference is just whether

3.4 Putting it all together: an argument for Specific Significance

What we've established so far is that the fact that one's belief is convenient is some evidence that one's belief is irrational. Presumably, then, the realization that one's belief in p is convenient should increase one's confidence in the (higher-order) proposition *my belief in p is irrational*. But Specific Significance says more than that: it says that this realization should decrease one's confidence in the (first-order) proposition p . The final premise that we need to add to get that conclusion is none other than General Significance – that is, the claim that higher-order evidence is epistemically significant with respect to one's *first-order* beliefs, in the sense that when one gets higher-order evidence that one's (first-order) belief is irrational, one should decrease one's confidence in that belief. With this premise in place, we can present the overall case for Specific Significance that has been cumulatively building up in this section as follows:

- (1) The fact that your belief is suspiciously convenient is at least some evidence that it was produced by motivated reasoning.²⁶
- (2) Evidence that your belief was produced by motivated reasoning is at least some evidence that it is irrational.
- (3) When you get evidence that your belief is irrational, you should reduce your confidence in your belief. (*General Significance*)²⁷

So,

- (4) The fact that your belief is suspiciously convenient should lead you to reduce your confidence in it. (*Specific Significance*)

Once the cases for them have been laid out, I think it is very hard to resist either premise (1) or premise (2) in the argument just summarized.²⁸ Thus, the best path to denying the argument for Specific

reasoning is involved at all is mistaken. Moreover, the modified case is one that parallels motivated reasoning pretty closely, so once we see that the irrelevant influence is relevant in this case, it's only a short step to a similar conclusion about irrelevant influences that indicate motivated reasoning, which is a *much* bigger category than the sort of case White concedes involves an epistemically relevant irrelevant influence.

²⁶ Unlike some of the literature on irrelevant influences, this argument makes no stipulation that in the cases of interest, you *learn with certainty* that your belief is due to an irrelevant influence—just that you get *some evidence* that your belief is due to an irrelevant influence. I think that this is the best way to set up the debate about irrelevant influences and higher-order evidence more generally. *Learning for certain* that our beliefs are due to irrelevant influences is relatively rare in real life; but, as the present argument shows, it is not required to get the case for the significance of irrelevant influences going.

²⁷ It's important that as I've stated it, General Significance only says that you should *reduce your confidence* in your belief when you get evidence that it is irrational (cf. fn. 10 on how to interpret this). I don't think it's true that such higher-order evidence always ought to make you give up your (outright) belief entirely; indeed, I'm not even committing to a view whereby this is so when the higher-order evidence is strong enough to decisively support the (higher-order) proposition that your belief is irrational (see Worsnip 2018 for some reasons for doubt about this view). Nevertheless, even as I've stated it, General Significance seems to be what Weatherson (2019: 120) calls a "level-crossing principle," according to which there is a necessary connection between a first-order and a second-order epistemic claim. Weatherson holds that there are no true level-crossing principles, but it's not clear to me that his arguments against them speak against principles like General Significance that take on commitments only about the need to reduce confidence somewhat, rather than to give up belief entirely.

²⁸ As noted above, White (2010) might be read as, in effect, denying one or both premises. I've already explained why I think his arguments misfire in the final paragraphs of section 3.2 and 3.3 respectively.

Significance is to deny premise (3) – i.e., General Significance. (This is why, as I indicated earlier, I think that the dispute about Specific Significance will ultimately come down to the dispute over General Significance, though it took some work to see why this is so.)

4. Resistance to the argument for Significance (General and Specific)

4.1 *Anti-Significance*

The most common position taken by those who deny General Significance is that the weight that you should give to the arguments and (first-order) evidence depends on how strongly they *actually* support your belief – not on whether you have (higher-order) *evidence* that they do or don't support your belief.²⁹ Applied to the case of suspiciously convenient belief, the view is this. If the arguments and first-order evidence don't actually support your belief much, you shouldn't give them much weight – but this is so regardless of whether your belief is suspiciously convenient. If they do actually support your belief a lot, you should continue to give them a lot of weight, the fact that your belief is suspiciously convenient notwithstanding. Either way, then, the convenience of your belief makes no difference to how confident you should be. The norm that these opponents of Significance (both General and Specific) endorse, then, is simply this:

Anti-Significance. You should hold on to your suspiciously convenient belief iff it's actually supported by the first-order arguments and evidence.

Analogous positions have been defended in a range of other debates. For example, proponents of the “right reasons” view of disagreement hold that whether you should lose confidence in a belief disputed by some peer of yours depends only on whether the first-order arguments and evidence (i.e., the right reasons) supported your beliefs in the first place.³⁰ And recently, some philosophers have claimed that whether you should lose confidence in beliefs that were produced in “echo chambers” just depends on whether the echo chamber is a “good” one – one that promotes well-supported views and denigrates ill-supported ones – or a “bad” one – one that does the opposite.³¹

4.2 *Ideal (epistemological) theory*

In the two subsections following this one, I'll aim to explain how many of the common ways of supporting Anti-Significance, more or less tacitly, rest on a form of epistemological ideal theory, and

²⁹ Broadly speaking, this is the view taken by all the opponents of General Significance cited in fn. 9: Titelbaum (2015), Tal (2020), and perhaps Lasonen-Aarnio (2014) and Weatherston (2019). White (2010) also takes a similar view with respect to irrelevant influences in particular, though because he is a conciliatorist about disagreement (White 2009), it seems that he cannot accept it with respect to *all* higher-order evidence.

³⁰ Proponents of this view include Kelly (2005) [but note that Kelly 2010 endorses a different, less extreme view] and Titelbaum (2015).

³¹ Proponents of this view include Lackey (2021) and Fantl (2021). See Sheeks (forthcoming) for pushback.

to try to bring out why – at least in this context – this is a problem. To do that, let me give us a sharper fix on what ideal theory in epistemology is.

The term ‘ideal theory’ can refer to a bewildering variety of different methodological views. Here I will understand it as being comprised of two (related and mutually supporting) commitments. The first is:

IT₁. The appropriate way to theorize is to begin by describing the ideal epistemic agent. Though we are non-ideal agents, this will still be relevant to us because we should try to approximate the ideal epistemic agent as best as possible.³²

IT₁ has a direct analogue in the political philosophy literature in which talk of ‘ideal theory’ is more familiar – namely, the view that the appropriate way to theorize is to begin by describing the ideal political system, and that this will be relevant to use because we should try to approximate it.³³

The second commitment is:

IT₂. What happens when imperfect, real people try to follow some epistemic norm *N* is irrelevant to the assessment of *N* qua norm. What’s relevant is only what would happen if people *succeeded* in following *N* perfectly (even where this is not practically feasible).

Again, this is close to a standard meaning for ‘ideal theory’ in political philosophy, according to which political theorizing need not be constrained by considerations of “feasibility” – or at least, not by those that are generated by real-world agents’ contingent imperfections (presumably it’s OK to take into account whether a political system is possible in a weak sense, such as whether it’s metaphysically possible).³⁴

4.3 *IT₁ and Anti-Significance*

In seeing how the case for Anti-Significance is related to these two ideal-theoretic commitments, let’s begin by noting how the assumption of IT₁ prima facie supports Anti-Significance. At least arguably, a truly ideal epistemic agent would always be able to look at the first-order arguments and evidence bearing on their belief, and infallibly assess whether they support the belief in question. Thus, it’s tempting to conclude, an ideal agent could *only* be led astray by higher-order evidence (or, at least, the sort of higher-order evidence that indicates irrationality): since their determinations of what the first-order evidence supports are always correct, any evidence suggesting that these determinations are incorrect is *ipso facto* misleading, and thereby liable to lead them astray. Thus, for ideal agents, Anti-Significance is a good norm to follow; and hence, if we’re trying to approximate the belief-forming procedures of ideal agents, perhaps we should follow it too. To put what is essentially the same point

³² In principle, a proponent of ideal theory might reject this second claim, and clarify that their ideal theory has no consequences for what real-world agents should do or believe. I will set this view aside, since it simply has nothing to say about how *we* should respond to the fact that our beliefs are suspiciously convenient.

³³ Valentini (2012: 654) identifies this as the third of three different meanings for ‘ideal theory’ in political philosophy.

³⁴ This is the second of Valentini’s (2012: 654) meanings for ‘ideal theory’.

in a different way, we will be led to more accurate beliefs if we perfectly respond to our first-order evidence than if we respond to this *plus* evidence that casts doubt on whether these perfect responses to the first-order evidence are in fact perfect.³⁵ Thus, again, a perfectly rational agent – who can perfectly respond to their first-order evidence – should opt for the former strategy.

To my mind, the fact that IT_1 has this consequence (assuming that it does) brings out quite sharply why IT_1 should be rejected. For the obvious reply is that when it comes to higher-order evidence, it is highly relevant that we ourselves are *not* ideal epistemic agents – more particularly, that we know that our attempts to just directly assess the quality of the arguments and evidence are themselves often affected by motivated reasoning and the like. For this reason, higher-order evidence will often not lead *us* astray, but rather be salutary, allowing us to mitigate our prior mistakes. Rationality requires imperfectly rational agents to take account of their own imperfect rationality, and try to mitigate it in various ways – whereas a perfectly rational agent, ipso facto, has no imperfect rationality to take account of. This is a crucial difference between us and ideal epistemic agents, and explains why we should not try to approximate their belief-forming procedures – i.e., directly assessing the arguments and evidence and ignoring higher-order evidence. Here, there is a non-equivalence between two things that may initially seem equivalent, namely *what it's rational to believe* and *what an ideally rational agent would believe*.

Some proponents of Anti-Significance are aware of this potential response to arguments for Anti-Significance that explicitly rest on IT_1 . In an intriguing recent paper, Eyal Tal (2021) attempts to construct a more sophisticated argument for Anti-Significance that gets around this response, and that applies even to non-ideal agents. Tal's argument, slightly simplified, is as follows. Either one's existing belief is rational given one's (first-order) evidence, or it isn't. If it *is* rational, then higher-order evidence suggesting that this belief is irrational is misleading. If it *isn't* rational, then this higher-order evidence is superfluous: since one's belief is *already* irrational, it's already the case that one should give it up. Tal then contends that if one knows that some piece of (putative) evidence is *either* superfluous *or* irrational, then epistemic rationality requires one to ignore it. And even a non-ideal agent – who cannot infallibly tell what their first-order evidence supports, and so cannot (always) determine whether the higher-order evidence is superfluous or irrational – can tell that it must be one of the two. Hence, we non-ideal agents are epistemically rationally required to ignore (putative) higher-order evidence that their belief is irrational.³⁶

While Tal's argument does not (explicitly) *presuppose* IT_1 as such, I'll argue that the premises of his argument can't all be true for non-ideal agents, and so that contrary to what he claims, the argument does not show that non-ideal agents can rightfully ignore higher-order evidence. Here are two of the crucial premises of Tal's argument:

³⁵ Cf. Schoenfield (2015).

³⁶ As Tal explains, his argument (if successful) also works for higher-order evidence that your belief *is* rational. But it is much less obvious that it gets going with (all instances of) the sort of higher-order evidence that concerns what your evidence supports without concerning some existing belief; Tal doesn't consider this sort of higher-order evidence. This is one instance of the problem mentioned in fn. 3 wherein different conceptions of higher-order evidence are at work in the literature, and this can lead to ships passing in the night.

- (5) If your belief in p is irrational, then higher-order evidence suggesting that it is irrational is superfluous.
- (6) If you know that some piece of evidence is either misleading or superfluous, you are epistemically rationally required to ignore it.

I'll argue that depending on how we interpret 'superfluous', one or the other of these premises is false for non-ideal agents. On one meaning of superfluous: a piece of evidence is superfluous if it *serves no useful function for you* given the evidence already available to you. If we interpret 'superfluous' this way, then I think (6) is plausible. But on this interpretation, (5) is true only for ideal agents. For, as we've already seen, higher-order evidence *does* play a useful function for non-ideal agents, even when their first-order evidence already requires them to give up a belief: specifically, it can make them take that possibility (viz., that the first-order evidence doesn't actually support their belief) seriously when they have failed to ascertain it. (An ideal agent, by contrast, wouldn't have any use for higher-order evidence of this kind, since they would already be able to infallibly tell that the first-order evidence doesn't support their belief; indeed, they would presumably have already given up their belief as soon as that became true.)

On a different interpretation, a piece of evidence is superfluous (with respect to some proposition) if it doesn't change the rational attitude to that proposition, because it supports the same attitude that one's first-order evidence already did. If we interpret 'superfluous' this way, then (5) is nigh-on indisputable. But on this interpretation, I think that it's (6) that is true only for ideally rational agents. For the mere fact that the higher-order evidence is superfluous in this sense only suggests that it can safely be disregarded if I'm certain that I *already have* the rational attitude, and (in contrast to ideal agents, or at those aware of their own ideality) this is not an assumption that non-ideal agents are entitled to make. Thus, knowledge that evidence is *either* misleading or (in this sense) superfluous does not entitle, still less require, non-ideal agents to disregard it. On the contrary, precisely because higher-order evidence that is superfluous in this sense can (as we just saw) still be salutary for non-ideal agents by drawing their attention to the as-yet-unnoticed (likely) irrationality of their belief on their first-order evidence, it seems if anything that they are required *not* to disregard higher-order evidence that is, or very well might be,³⁷ superfluous in this sense.³⁸

Thus, on either interpretation of 'superfluous', the argument goes through only for ideal agents.

4.4 IT_2 and Anti-Significance

³⁷ And remember that it's not like higher-order evidence is superfluous (in this sense) and misleading in an equal number of cases. Good evidence that you are irrational, almost by definition, will more often than not be "right" in that you *are* irrational.

³⁸ Tal would dispute this last claim. He thinks that the fact that higher-order evidence is salutary in this way (for imperfect agents) constitutes a prudential, rather than an epistemic, reason for them to adjust their beliefs in light of it. I disagree, and the dispute here turns in part on some issues that I'll touch on in discussing IT_2 below. Regardless of this, the point that his argument has failed to establish its positive conclusion that one (epistemically) *should* disregard higher-order evidence stands.

To begin exploring the relationship between Anti-Significance and IT_2 , consider what happens when people try to follow Anti-Significance. People can only follow a norm like Anti-Significance via applying their judgments about when the right-hand side is satisfied—that is, about whether their suspiciously convenient beliefs are supported by the first-order arguments and evidence. But if I sincerely hold a belief, I typically do think I have good arguments and evidence for it (even when, in fact, the arguments are not as good as I think). Moreover, this is especially likely to be the case when motivated reasoning is work, since it isn't consciously detectable and works via one's convincing oneself that the arguments and evidence are good. Thus, the *de facto* result of people trying to follow Anti-Significance will be that people hang on to their beliefs in the vast majority of cases, including many where the belief is not actually supported by the arguments and evidence. Even if people re-examine their beliefs, and whether they are supported by the first-order evidence and arguments, in light of their suspicious convenience, there's no reason why motivated reasoning wouldn't just kick in again.³⁹

Now, the proponent of Anti-Significance will likely react to this roughly as follows: “This isn't an objection to my view. For Anti-Significance doesn't sanction hanging onto one's view merely when one *thinks* it is supported by the first-order arguments and evidence. It only sanctions hanging onto one's view when it is *actually* supported by the first-order arguments and evidence. Indeed, when one thinks one's view is supported by the first-order arguments and evidence but it actually isn't, my view recommends giving it up! So it can't be an objection to my view that it's bad for people to hang on to their beliefs when they *think* it's supported by the first-order evidence and it isn't. My view simply doesn't sanction this; on the contrary.”

But notice: this response tacitly relies on IT_2 . It assumes that what happens when people *actually* try to follow Anti-Significance (and apply their own judgments about whether its right-hand side is satisfied) is irrelevant to the assessment of Anti-Significance. What's relevant is only what it actually sanctions, or what would happen if people followed it perfectly. And this just is IT_2 .

Again, my view is that the role that IT_2 plays in this context brings out the case for its rejection. Pronouncing what happens when people actually try to follow epistemic norms as irrelevant to whether to accept these norms—as IT_2 counsels—allows such norms to perform a pernicious, covert ideological function. (Here we see that epistemological ideal theory can perform a pernicious ideological function in a roughly analogous way to that which Mills (2005) argues that ideal theory in political philosophy can.) For example, in the case at hand, it allows Anti-Significance to enable and legitimize the rhetoric—which ones sees both in philosophy and (albeit perhaps in slightly simpler terms) outside of it—that casting doubt on a believer's motivations doesn't cast any doubt on the rationality of their views themselves, and that to do that what's needed is a *first-order* argument directly aimed at showing the conclusion to be false (or, at least, ill-supported by the evidence). Anti-Significance thus acts as a *resource* to those who want to defend, and avoid having to give up, their own blatantly suspicious beliefs, even in cases where these beliefs are in fact not supported by the evidence

³⁹ Contrast White (2010: 576), who suggests that the presence of irrelevant influences should merely occasion a reassessment of the first-order arguments and evidence for one's views, and not lead to a direct downgrade in confidence in itself. This is especially unpromising for cases of motivated reasoning, since the reassessment in question is also subject to motivated reasoning; but more broadly, it elides the way in which what seems compelling, or seems to be supported by the first-order arguments and evidence, can itself be a function of etiology.

(and so are not in fact sanctioned by Anti-Significance).⁴⁰ The problem is not that Anti-Significance legitimates such beliefs *de jure*, but that it is exploitable by those who wish to legitimate them *de facto*.⁴¹ In my view, we should not want our epistemic norms to be exploitable in this way.

The defender of Anti-Significance (and IT₂) might respond here in one (or both) of two ways. First, they might say that in cases where a belief *is* in fact biased or more generally ill-supported, we *can* show that people ought to give up the belief by simply doing what was demanded: that is, by supplying a first-order argument that the belief is false or ill-supported by the first-order evidence. If that's so, we don't need to engage in critique of the believer's motivations, and accepting IT₂ and Anti-Significance need not preclude us from engaging in critique of biased, ill-supported beliefs.

However, in many of the cases at hand, the first-order ill-supportedness of the belief is extremely hard to demonstrate conclusively and in a dialectically effective way. First, the rationalizations that people cook up for their beliefs via motivated reasoning are often highly sophisticated, and can't just be shown to be fallacious in some straightforward and obvious way. They may involve simply putting too much weight on some evidential factors and not enough weight on others, but it's exceptionally tricky to be precise about the exact amount of weight that different items of evidence have, and so it's hard to just "point out" that some factor was overweighted or underweighted. Moreover, political disputes will often come down to disagreements about foundational premises, often about questions of value, rather than involving one side engaging in mistaken reasoning *from* those premises. Again, mistaken foundational premises, especially those about value, are hard to conclusively demonstrate to be false. Yet when these foundational premises have themselves plausibly been adopted merely because they lead to convenient conclusions, I submit that the difficulty of conclusively showing them to be false should not stop us from being suspicious of them, and the difficulty of showing the arguments for these conclusions to be fallacious should not stop us from thinking those arguments are smokescreens.

A second response for the defender of Anti-Significance (and IT₂) holds that considerations about the bad effects of Anti-Significance in public discourse, and about how things will go when people try to follow it, are just reasons not to *propound* the norm of Anti-Significance in public discourse, rather than reasons not to accept it, or that show it to be false.⁴² This is reminiscent of some ideas in ethics: the "government house utilitarianism" often attributed to Sidgwick (1907), according to which utilitarianism is true but we shouldn't tell this to people (and should instead propound norms of common sense morality) because they'll do a terrible job of maximizing utility.

The proponent of the view that Anti-Significance is true but we shouldn't propound it is forced into a rather difficult corner practically, for even expressing this view seems to count as

⁴⁰ In fact—to be a little controversial—I suspect that IT₂ and Anti-Significance are sometimes adopted by philosophers who *themselves want to use them to defend their own suspicious views*, precisely (albeit not consciously) for the reason that they will allow them to do so. (Consider again the case whereby the religious philosophers defends the view that irrelevant influences are epistemically irrelevant.) If this is so, IT₂ and Anti-Significance don't just serve an ideological function, but are adopted precisely because they do so.

⁴¹ Indeed, the fact that it doesn't legitimate them *de jure* is part of the rhetorical strategy that enables it to legitimate them *de facto*. For it allows the defender of Anti-Significance to protest, as we just saw, that her view *doesn't* legitimate genuinely biased beliefs (*de jure*), and so to obscure the ways in which it does so *de facto*; that is, to obscure its own ideological function.

⁴² Thanks to Roger White for this suggestion. Compare also Weatherston's (2019: 131) suggestion that we sharply separate what an agent is rational to believe and what we should *advise* them to believe.

propounding Anti-Significance in at least one good sense. Perhaps the idea is that we shouldn't *publicly* propound it, ensuring that our epistemological articles and books stay within the cloistered walls of the academy, and don't make it outside to the unwashed masses, lest we count as having propounded the norms we defend. I don't think that is a very attractive vision of the role of epistemology and philosophy more broadly in public life, and in my view, if accepting IT_2 forces us to endorse it, so much the worse for IT_2 . And moreover, since philosophers and other academics are also vulnerable to motivated reasoning, even propounding Anti-Significance only to them carries at least some of the risks of propounding it more widely.

Moving beyond this practical issue for the proponent of the view that Anti-Significance is true but we shouldn't propound it, there is, I think, a deep metaepistemological disagreement between myself and those who push this line that cannot be fully resolved here. Specifically, I accept a mildly pragmatic metaepistemological view according to which the epistemic norms that we endorse (and, perhaps even more so, the methodological positions we accept, such as IT_2 or its negation) should be partly shaped by what we want (or should want) epistemological theory to *do* for us—what use we want to put it to. Moreover, the *truth* of such norms is not fully independent of this: part of what *makes* an epistemic norm true is its serving the purposes we should want epistemic norms to serve.⁴³ As such, there is less potential for a gap to open up between the epistemological norms that are true and those we ought to propound, and this makes IT_2 very hard to sustain.

One of the things we should want epistemological theory to do for us, I suggest, is to make critique possible: in this instance, to make possible the critique of arguments that are smokescreens for truth-irrelevant interests and desires. If what I've argued is right and it is often hard to compellingly critique these arguments on first-order grounds, then we need to accept a norm like (Specific) Significance to make this critique possible, while accepting Anti-Significance precludes us from doing so. On the mildly pragmatist metaepistemological picture I favor, this does have some bearing on whether (Specific) Significance or Anti-Significance is true, and not merely on whether we ought to propound them as norms.

A final point, which is independent of these (controversial) metaepistemological claims. It's plausible that the fact that people will do a bad job of following Anti-Significance—of hanging onto their suspicious beliefs just if those beliefs are supported by the first-order evidence—actually *makes it the case* they ought instead to follow the policy of downgrading their confidence in their suspicious beliefs. But if they ought to do this, then (Specific) Significance just is true, and Anti-Significance is false (at least in the sense of 'ought' that should actually guide people's beliefs and belief-forming procedures). Again, this puts pressure on IT_2 , the view that what will happen when people actually try to follow a norm is irrelevant for assessment of this norm, as well as the view that these considerations bear only on whether we should *propound* the norm in question.

4.5 *A related distorting maneuver*

I now want to explore a related maneuver that is often made by proponents of Anti-Significance, which is not a direct product of IT_1 or IT_2 , but is broadly in the spirit of ideal theory, and is, I think,

⁴³ See also Srinivasan (2020: 415), and the references therein.

similarly distorting. Anti-Significance says that whether one should reduce confidence in one's suspiciously convenient belief depends simply on whether it is supported by the first-order arguments and evidence. So it's natural for proponents of Anti-Significance to subdivide cases of suspiciously convenient beliefs (and of higher-order evidence more broadly) into two categories: (i) those where the arguments and (first-order) evidence do not in fact support the belief much; and (ii) those where the arguments and (first-order) evidence do in fact support the belief. Proponents of Anti-Significance then often make the following maneuver, or something like it. First, they note that everyone—fans of Significance and Anti-Significance alike—agrees that you should reduce confidence in your belief in cases of type (i).⁴⁴ Then, they say that because of this, the “interesting” cases are those of type (ii), and so we should focus the debate exclusively on those cases. Finally, and consequently, they stipulate, in every concrete case of suspiciously convenient beliefs (or irrelevant influences, or higher-order evidence more broadly) that they discuss, that the belief in question is in fact supported by the first-order evidence and arguments.⁴⁵

I want to suggest that this maneuver is distorting in two ways. First, notice that (at least when we're dealing with higher-order evidence that suggests that one's belief is irrational, such as suspicious convenience), the type-(ii) cases are ones in which the higher-order evidence is misleading, and the type-(i) cases are ones in which it is not misleading. But it's distorting to our intuitions about higher-order evidence to focus exclusively on cases in which higher-order evidence is misleading! Doing so leads to disproportionate focus on the cases where higher-order evidence leads us astray, and insufficient focus on the cases where it does us a service by drawing attention to our genuine failings—and this skews our intuitions toward Anti-Significance. The distortion is particularly severe because in reality, the majority of cases of higher-order evidence are in fact of type (i). This is so because higher-order evidence, like all evidence, is usually not misleading: arguably, it just follows from E's being good evidence for p that E is at least a fairly reliable indicator that p , and this is just as so when the p in question has to do with one's own irrationality.

Second, when features of imaginary cases are stipulated, we will plausibly tend to *project* those stipulations onto the epistemic position of the subjects that the cases concern. So if it's stipulated that a subject (say, someone like Connor, in our opening examples) is in a type-(ii) case, we may (perhaps tacitly) imagine the case as if Connor himself has access to this fact, and can tell for sure that his case really is one where his belief really is supported by the first-order evidence, its suspicious convenience notwithstanding. If that were so, then of course it would then follow that Connor should not lose confidence in his belief. But it is not in fact so. Because motivated reasoning is largely first-personally undetectable, Connor is not in a position to know that he has not succumbed to it (even if he actually hasn't). Indeed, it's part of the setup of the case that he has strong evidence that he has succumbed to it. And in evaluating what it's rational for *Connor* to believe, we ought to take seriously the limitations of Connor's epistemic position.

⁴⁴ As Jim Pryor pointed out to me, this actually fails to show that proponents of Significance and proponents of Anti-Significance agree on cases of type (i). For fans of Significance, it might be that in cases of type (i), the higher-order evidence gives you a special and *additional* reason to lose confidence, over and above that provided by the first-order evidence, that you ought to lose *more* confidence than the first-order evidence alone would warrant your losing, or that remaining equally confident would be *more* unreasonable than if you lacked the higher-order evidence in question.

⁴⁵ For maneuvers of this kind, see, e.g., White (2010: 600); Begby (forthcoming: 9-10, n. 25).

In light of these two points, there are good grounds for resisting the proponent of Anti-Significance’s tendency to subdivide cases into those of type (i) and those of type (ii) and then to stipulate, in giving examples, that the case is of type (ii). Rather, perhaps we should simply lay out the evidence they have – evidence that suggests that there’s a strong possibility that their beliefs are due to motivated reasoning – and then ask how they should react to this evidence, without stipulating whether these belief are really due to motivated reasoning or not. In doing so, we are asking the question about how the agent should adjust their beliefs *from the point of view of that agent*, who is (or should be) *uncertain* about whether their beliefs are well-supported by the evidence, or whether they are due to motivated reasoning—not from the perspective of an omniscient third party who has infallible access to the facts about what the first-order evidence really supports.⁴⁶ I suggest that once we ask the question from the perspective of the agent, Anti-Significance is much harder to sustain.

4.6 *A leftover thread: evidentialist arguments for Anti-Significance*

There is one kind of argument for Anti-Significance that I have not commented on here, and that at least seems not to turn on any kind of ideal-theoretic maneuvers. This argument proceeds from the two premises: (1) the evidentialist norm that you should hold on to a belief iff it is supported by your total evidence, and (2) the claim that suspicious convenience (and, indeed, higher-order evidence more generally) does not affect whether your belief is supported by your total evidence. (2) is typically supported by the view that higher-order evidence (at least of the kind that concerns the malfunctioning of one’s own cognitive processes) is not itself evidence for or against one’s belief, *plus* the contention that it can’t be assimilated to the phenomenon of undercutting defeat,⁴⁷ whereby one gets evidence that affects what one’s total evidence supports by undercut the force of one’s existing evidence for one’s belief.⁴⁸

There is a flourishing literature on this topic, and there are a number of ways of responding to this argument for Anti-Significance. Some, in effect, resist (1) by arguing that there are elements of our total evidence that we are required to “bracket” when we get higher-order evidence,⁴⁹ or that there are special non-evidential reasons to suspend judgment.⁵⁰ Others, in effect, resist (2) by arguing that

⁴⁶ Here I’m chiming with similar points made by Johnson King (2022) in defense of epistemic internalism against the arguments of Srinivasan (2020). As she aptly writes: “there are no real cases in which it is helpfully stipulated to us that our hunches and gut feelings are or are not reliable” (49). Thus, “it’s no good [...] to discuss vignettes in which we stipulate information about the reliability of subjects’ belief-forming mechanisms that ordinary people don’t possess about themselves or one another, calling the subjects justified or unjustified according to this god’s-eye-view stipulated information. These are unhelpful responses to the realities of epistemic oppression. As with other forms of oppression, the only acceptable response to epistemic oppression is to fight it as best we can. So, those of us with radical worldviews should be interested in what we—actual people, with the information that we have—can do to resist bad ideology. [...] For that project, we will need some of the more action-guiding principles that are the internalist’s stock-in-trade” (61).

⁴⁷ For the classic discussion of undercutting defeat, see Pollock (1986: 39).

⁴⁸ The claim that higher-order evidence cannot be assimilated to undercutting defeat is accepted by many philosophers, including both proponents and opponents of Anti-Significance. See e.g. Feldman (2005: 111-3); Christensen (2010: 193-5); White (2010: 585); Lasonen-Aarnio (2014: 317-8); González de Prado (2020: 327-8); and Lord & Sylvan (2021: 124).

⁴⁹ Christensen 2010.

⁵⁰ Lord & Sylvan 2021.

although higher-order evidence cannot be assimilated to classic undercutting defeat, it defeats one's justification for one's belief in some other way.⁵¹

My own opinion is that the view that higher-order evidence *can* be seen as giving effect to classic undercutting defeat has been dismissed too quickly,⁵² and that it can be defended once we appreciate the holistic character of evidential support.⁵³ If this is so, it is (2) that should be rejected, but not for the reasons others have given. Due to constraints of space—and because this takes me too far away from my main topic of the relationship between Anti-Significance and ideal theory—defense of this view will have to wait for another occasion. However, it's worth noting that if what I've argued in §4.4 above is correct—that is, if I'm right both that IT₂ should be rejected and that once this is done, Anti-Significance must be rejected too, then it follows that the evidentialist argument for Anti-Significance must be unsound. If we can explain why suspicious convenience (and higher-order evidence more broadly) *does* affect what one's total evidence supports, and thus preserve the evidentialist norm that the only thing affecting what you should believe is what your total evidence supports—rejecting (2) and preserving (1)—great. But ultimately, if the cost of accounting for the way in which suspicious convenience matters were to reject the evidentialist norm (1), this would be a price I'd be willing to pay, in light of the considerations surveyed in the foregoing sections.

5. Conclusion: Anti-Significance norms in public discourse

As I mentioned in §4.4, one sometimes sees the idea that we should not cast doubt on others' motivations for believing as they do—and should instead focus only on evaluating their arguments, and offering counterarguments—cropping up in public discourse. Sometimes this is abetted by philosophers (and in particular by philosophers who, ironically enough, I suspect have their own ideological motivations for abetting it).⁵⁴ If what I've argued in this paper is correct, however, there is nothing *epistemically* irrelevant or fallacious about seeking to question the rationality or well-supportedness of others' beliefs by considering their motivations, and the extent to which those beliefs are suspiciously convenient.⁵⁵ These facts are higher-order evidence, and higher-order evidence does affect the rationality of belief. This leaves open the possibility that there is, as a couple of philosophers have recently intriguingly suggested, something *ethically* problematic about trying to cast doubt on others' beliefs by questioning their motivations.⁵⁶ I am skeptical of this too, and hope to address it in future work. But for now, the point is that the battle will have to be joined at the ethical level; our best

⁵¹ González De Prado 2020.

⁵² To my knowledge, the only authors who suggest that higher-order evidence does effect undercutting defeat are Avnur & Scott-Kakures (2015: 22-24). However, even they concede (too quickly, in my view) that higher-order evidence doesn't undercut *propositional* justification, only doxastic justification. And they do not explicitly respond to the arguments that have been given as to why higher-order evidence does not undercut justification.

⁵³ Interestingly, the person who comes closest to my view about the way in which higher-order evidence can affect what one's total evidence supports is Weatherson (2020: 134, 137). But confusingly (to me), Weatherson takes his explanation to *rival* explanations in terms of "level-crossing principles," according to which there is a necessary connection between a first-order and a second-order epistemic claim. By contrast, I take it to *vindicate* General Significance, which *just is* a level-crossing principle (see n. 27). Again, I hope to adjudicate this messy debate in future work.

⁵⁴ E.g.: <https://twitter.com/TomasBogardus/status/1665216560082870272>

⁵⁵ Contrast the widespread view that to do so is to commit something called the "genetic fallacy" (Cohen & Nagel 1934).

⁵⁶ Smyth 2021; Flowerree forthcoming.

epistemic norms will not condemn us either for our suspicion about others' motivations for their beliefs, nor our bringing that suspicion to bear on our evaluations of those beliefs' rationality.

References

- Avnur, Y. & Scott-Kakures, D. (2015). "How Irrelevant Influences Bias Belief," *Philosophical Perspectives*, 29, 7-39.
- Baron, J. (2009). "Belief-Overkill in Political Judgments," *Informal Logic*, 29(4), 368-378.
- Begby, E. (forthcoming). "From Belief Polarization to Echo Chambers: A Rationalizing Account," *Episteme*. Page references are to the Online Early View version (doi: 10.1017/epi.2022.14).
- Carr, J.R. (2021). "Why Ideal Epistemology?," *Mind*, 131(524), 1131-1162.
- Carter, J.A. & McKenna, R. 2020. "Skepticism Motivated: On the Skeptical Import of Motivated Reasoning." *Canadian Journal of Philosophy*, 50(6), 702-718.
- Chen, Y. & Worsnip, A. (forthcoming). "Disagreement and Higher-Order Evidence," in Baghramian, Carter & Rowland (eds.), *The Routledge Handbook of the Philosophy of Disagreement*. Routledge.
- Christensen, D. (2007). "Epistemology of Disagreement: The Good News," *Philosophical Review*, 116(2), 187-217.
- (2010). "Higher Order Evidence," *Philosophy and Phenomenological Research*, 81(1), 185-215.
- Cohen, G.A. (2000). *If You're an Egalitarian, How Come You're So Rich?* Harvard University Press.
- Cohen, M.R. & Nagel, E. (1934). *An Introduction to Logic and Scientific Method*.
- Ditto, P.H., Jemmott, J.B. & Darley, J.M. (1988). "Appraising the Threat of Illness: A Mental Representation Approach," *Health Psychology*, 7, 183-201.
- Feldman, R. (2005). "Respecting the Evidence," *Philosophical Perspectives*, 19, 95-119.
- Flores, E. & Woodard, E. (2023). "Epistemic Norms on Evidence-Gathering," *Philosophical Studies*, 180(9), 2547-2571.
- Flowerree, A.K. (forthcoming). "When to Psychologize," *Australasian Journal of Philosophy*.
- Gilovich, T. (1983). "Biased Evaluation and Persistence in Gambling," *Journal of Personality and Social Psychology*, 44, 1110-1126.
- González de Prado, J. (2020). "Dispossessing Defeat," *Philosophy & Phenomenological Research*, 101(2), 323-340.
- Greco, D. (2023). *Idealization in Epistemology: A Modest Modeling Approach*. Oxford University Press.
- Hieronymi, P. (2006). "Controlling Attitudes," *Pacific Philosophical Quarterly*, 87(1), 45-74.
- Horowitz, S. (2014). "Epistemic Akrasia," *Noûs*, 48(4), 718-744.
- Jervis, R. (1976). *Perception and Misperception in International Politics*. Princeton University Press.
- Johnson King, Z. (2022). "Radical Internalism," *Philosophical Issues*, 32, 46-64.
- Kahan, D.M. (2016). "The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It," *Emerging Trends in the Social and Behavioral Sciences*.
- Kahan, D.M., Jenkins-Smith, H. & Braman, D. (2011). "Cultural Cognition of Scientific Consensus," *Journal of Risk Research*, 14(2), 147-174.
- Kassarjian, H.H. & Cohen, J.B. (1965). "Cognitive Dissonance and Consumer Behavior," *California Management Review*, 8, 55-64.

- Kelly, T. (2002). "The Rationality of Belief and Some Other Propositional Attitudes," *Philosophical Studies*, 110(2), 163-196.
- (2005). "The Epistemic Significance of Disagreement," *Oxford Studies in Epistemology*, 1, 167-196.
- (2008). "Disagreement, Dogmatism, and Belief Polarization," *Journal of Philosophy*, 105(10), 611-633.
- (2010). "Peer Disagreement and Higher Order Evidence," in Feldman & Warfield (eds.), *Disagreement*. Oxford University Press.
- Kunda, Z. (1987). "Motivation and Inference: Self-Serving Generation and Evaluation of Evidence," *Journal of Personality and Social Psychology*, 53, 636-647.
- (1990) "The Case for Motivated Reasoning," *Psychological Bulletin*, 108(3), 480-498.
- Kunda, Z. & Sanitioso, R. (1989). "Motivated Changes in the Self-Concept," *Journal of Experimental Social Psychology*, 25, 272-285.
- Lackey, J. (2021). "Echo Chambers, Fake News, and Social Epistemology," in Bernecker, Flowerree & Grundmann (eds.), *The Epistemology of Fake News*. Oxford University Press.
- Lasonen-Aarnio, M. (2014). "Higher-Order Evidence and the Limits of Defeat," *Philosophy and Phenomenological Research*, 88(2), 314-345.
- Leiter, B. (2004). "The Hermeneutics of Suspicion: Recovering Marx, Nietzsche, and Freud," in Leiter (ed.), *The Future for Philosophy*. Oxford University Press.
- Levy, N. (2020). "Suspiciously Convenient Belief," *Ethical Theory and Moral Practice*, 23(5), 899-913.
- Lord, C.G., Ross, L. & Lepper, M.R. (1979). "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence," *Journal of Personality and Social Psychology*, 37, 2098-2109.
- Lord, E. & Sylvan, K. (2021). "Suspension, Higher-Order Evidence, and Defeat," in Brown & Simion (eds.), *Reasons, Justification, and Defeat*. Oxford University Press.
- McKenna, R. (2023). *Non-Ideal Epistemology*. Oxford University Press.
- McWilliams, E. (2021). "Evidentialism and Belief Polarization," *Synthese*, 198(8), 7165-7196.
- (forthcoming). "Evidentialism and Duties to Inquire," *Philosophical Quarterly*.
- Mills, C. (2005). "'Ideal Theory' as Ideology," *Hypatia*, 20(3), 165-184.
- Murphy, L. (2003). *Moral Demands in Nonideal Theory*. Oxford University Press.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Rowman and Littlefield.
- Ricoeur, P. (1970). *Freud and Philosophy*. Yale University Press.
- Ross, M., McFarland, C. & Fletcher, G.J.O. (1981). "The Effect of Attitude on Recall of Past Histories," *Journal of Personality and Social Psychology*, 32, 829-840.
- Ryan, T.J. & Aziz, A.R. (2021). "Is the Political Right More Credulous? Experimental Evidence against Asymmetric Motivations to Believe False Political Information," *Journal of Politics*, 83(3), 1168-1172.
- Saint-Croix, C. (2018). *Non-Ideal Epistemology in a Social World*. Ph.D. dissertation, University of Michigan.
- Sanitioso, R., Kunda, Z. & Fong, G.T. (1990). "Motivated Recruitment of Autobiographical Memory," *Journal of Personality and Social Psychology*, 59, 229-241.

- Schoenfield, M. (2014). "Permission to Believe: Why Permissivism is True and What it Tells Us About Irrelevant Influences on Belief," *Noûs*, 48(2), 193-218.
- (2018). "An Accuracy Based Approach to Higher Order Evidence," *Philosophy and Phenomenological Research*, 96(3), 690-715.
- Sheeks, M. (forthcoming). "The Myth of the Good Epistemic Bubble," *Episteme*.
- Sidgwick, H. (1907). *The Methods of Ethics* (7th edition). Macmillan.
- Smyth, N. (2021). "A Moral Critique of Psychological Debunking," *Journal of Social Philosophy*, 53(2), 255-272.
- Srinivasan, A. (2020). "Radical Externalism," *Philosophical Review*, 129(3), 395-431.
- Tal, E. (2020). "Is Higher-Order Evidence Evidence?," *Philosophical Studies*, 178(10), 3157-3175.
- Titelbaum, M.G. (2015). "Rationality's Fixed Point (or: In Defense of Right Reason)," *Oxford Studies in Epistemology*, 5, 253-294.
- Valentini, L. (2012). "Ideal vs. Non-Ideal Theory: A Conceptual Map," *Philosophy Compass*, 7(9), 654-664.
- Vavova, K. (2018). "Irrelevant Influences," *Philosophy and Phenomenological Research*, 96(1), 134-152.
- Weatherson, B. (2019). *Normative Externalism*. Oxford University Press.
- White, R. (2010). "You Just Believe that Because..." *Philosophical Perspectives*,
- Williams B. (1973). "Deciding to Believe," in his *Problems of the Self*, 136–51. Cambridge University Press.
- Williams, D. (forthcoming). "The Case for Partisan Motivated Reasoning," *Synthese*.
- Wilson, T.D. (2002). *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Harvard University Press.
- Winters, B. (1979). "Believing at Will," *Journal of Philosophy*, 76(5), 243-256.
- Worsnip, A. (2018). "The Conflict of Evidence and Coherence," *Philosophy and Phenomenological Research*, 96(1), 3-44.
- (2019). "The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media," in Fox & Saunders (eds.), *Media Ethics: Free Speech and the Requirements of Democracy*. Routledge.