# Applied Epistemology: What is It? Why Do It?
Alex Worsnip

The remaining seven papers (eight, if you count this introductory piece) in this volume of *Oxford Studies in Epistemology* constitute a special issue on applied epistemology, an exciting, novel, and currently burgeoning subfield of epistemology. The term 'applied epistemology' is a relatively recent one, however, and anecdotally, many people I've encountered are not quite sure what it denotes, or what different works within the field have in common. In this introductory piece, I'll venture some views about these questions, and about why applied epistemology is worth doing, as well as about its dangers. Doing so will set the state for me to situate the papers in this volume within the subfield.

## I.     What is applied epistemology?

As a rough starting point, we can conceive of applied epistemology as standing to more "traditional" or "abstract" epistemology as applied ethics stands to normative ethics.[1] Whereas normative ethics looks for answers to the most *general* questions about which actions are (among other things) morally required, morally permissible, or morally prohibited, applied ethics considers which actions are required, permissible, and prohibited in more *specific* contexts. Moreover, these specific contexts tend to be ones that most arise not infrequently in real life, and that are invested with special significance to many real people – contexts such as those of medical decision-making, war, business practices, and intimate relationships. (There's of course a sense in which considering moral questions about any specific situation, real or imaginary, involving applying ethical principles and concepts; but we don't usually think of, say, "trollyology" as applied ethics, presumably because trolly cases are infrequently encountered in the real world and, relatedly, occupy little space in the collective socio-political imagination of non-philosophers.) That said, the boundaries between normative ethics and applied ethics are deeply fuzzy, since both involve asking *first-order* (rather than metaethical) questions about, *inter alia*, moral right and wrong, and because many of the things that distinguish them (generality, importance to ordinary people, socio-political significance) come in degrees. The attempt to aggressively police those boundaries seems both pointless and, in a good sense, doomed to failure. Still, some work in ethics is clearly *more* applied than other work is.

As in ethics, so in epistemology. Whereas abstract epistemology looks for answers to the most general questions about when beliefs and belief-forming processes are (among other things) justified or rational or knowledge-constituting, applied epistemology considers which beliefs and belief-forming processes are justified or rational or knowledge-constituting in specific contexts, particularly contexts of that have socio-political importance or are otherwise of great significance to us. Indeed, the traditional classification of epistemology as one of the "theoretical" or "abstract" subfields of philosophy—in supposedly sharp contrast to ethics—belies the way that ordinary discourse regularly invokes and appeals to distinctively epistemic norms, even if not under that description. We care not just about the actions that people perform, but about what they believe. Indeed, epistemic

---

[1] Cf. Coady & Chase (2019).

considerations are now more culturally salient than ever, what with contemporary concerns—especially, but far from exclusively, in the United States—about echo chambers, distrust of establishment experts, conspiracy theories, gaslighting, racial prejudice, polarization, partisanship, and much more. Applied epistemologists, then, have been thinking about questions like (among many others): can we *know* that prominent conspiracy theories are false? Can it be *rational* to hold on to beliefs that you realize have been formed in an echo chamber? *Under what conditions* is deference to experts required?

With regard to many of these questions, there is a received wisdom—one that, at least to the chattering classes, approaches appearance as a truism. Polarization is an unalloyed bad, echo chambers are always and everywhere pernicious, people should defer to the experts, conspiracy theories are blatantly irrational, and so on. Perhaps in many cases this received wisdom is correct. But here as so often, philosophy is here to ask whether it is, and if so, why. And as so often, when we dig just a little, the truth is more complex than it might first seem; or, at the very least, it's less straightforward to explain *why* the received wisdom is true than it might appear.

Again as in ethics, the boundaries between abstract and applied epistemology are fuzzy. Is work on epistemic partiality in friendship applied? Is (all) work on pragmatic and moral encroachment applied? (What if it uses some examples involving racism?) As with applied ethics, I'm not sure it's either possible or desirable to precisely delineate the boundaries of applied epistemology in a way that would decisively answer these questions. But again, it's clear that some work is more applied than other work.

## II.    Applied epistemology and "social epistemology"

The term 'social epistemology' is much longer-established than the term 'applied epistemology' is. A reasonable question is what the relationship between the two is.[2]

The answer, of course, depends on what 'social epistemology' denotes. On a very broad conception of social epistemology that one sometimes encounters, social epistemology is any kind of epistemology that acknowledges that we are socially situated beings that encounter and interact with each other, and not just isolated individual inquirers. (The common stereotype of the alternative, supposedly "traditional" conception of epistemology is Descartes alone in his study trying to establish a foundation for certain knowledge of the external world, *ex nihilo* and with the aid of no-one else.[3]) On this conception of social epistemology, it's clear that much social epistemology is not particularly applied. For example,[4] any work on peer disagreement is clearly *ipso facto* social epistemology in this sense: there obviously cannot be peer disagreement without other inquirers, and it is incoherent to ask what people should do upon encountering disagreement while also thinking of them as never encountering other rational beings. For all that, though, a great deal (albeit not all) of the peer

---

[2] Toole (ms.) has a nice discussion of this question.

[3] Though for what it's worth, it's not obvious to me that many contemporary epistemologists, including those who do very abstract work, presuppose this vision of the inquirer.

[4] Toole (ms.) also mentions disagreement as a social but non-applied topic.

disagreement literature is steadfastly abstract and un-applied. It might begin with lip service to the socio-political resonance of disagreement, yet it often quickly moves to simplifying assumptions that almost never hold in reality (such as disputants having perfectly identical evidence) and less "noisy"— thus maximally *a*political and *a*resonant—cases involving (say) the splitting of restaurant checks.[5]

On the other hand, on a very narrow conception of social epistemology that one also sometimes encounters, social epistemology is epistemology that evaluates only the beliefs (and other doxastic states) of *group* agents—or the knowledge-producing procedures of *institutions*, and doesn't concern itself with individual agents' beliefs at all. On this conception of social epistemology, it's clear that much applied epistemology is not social. Many of the questions of applied epistemology concern what it's rational for individuals to believe—socially-situated individuals, to be sure, but individuals all the same. For example, we can ask whether it's rational for individuals (given their social environments) to believe conspiracy theories or distrust experts.

Occasionally one hears the provocative opinion that work about the rationality of individual belief, even considered as socially situated or in very applied contexts, isn't "interesting" or "important". Groups and institutions, it's said, are where the action is. I incline toward a much more pluralistic approach. We are members of groups, but we are also individuals, and as individuals we face questions about what to believe in an increasingly confusing epistemic environment. I, at least, would like to know what I should believe in this confusing environment. I can't force others to be interested in that question, but nor can they force me to cease finding it urgent and gripping.

One thing that may encourage the attitude I'm setting myself against here is the often-expressed (and considerably more reasonable) opinion that the way to *solve* our various epistemic shortcomings as a society is not to try to get individuals to adopt for rational belief-forming procedures but to reform the institutions that structure their epistemic environments. That may well be right (whether it is, obviously, is in large part an empirical question). But the tacit assumption in moving from this claim to the claim that we should stop theorizing about the rationality of individual belief is that the (sole) function of epistemological theorizing is to solve society's epistemic problems. That seems to me both unrealistically utopian and excessively instrumental. Philosophers and non-philosophers alike may want to know what it's rational to believe in our baffling world precisely because that world baffles us, and not out (solely) of a desire to engineer a world that ceases to baffle.

### III.    Why do we need applied epistemology?

Why do we *need* applied epistemology—or, for that matter, applied ethics? If the answers to the questions of normative ethics or abstract epistemology are supposed to be perfectly general, shouldn't they be mechanically applicable to *any* case? Won't the only questions about *how* they apply be empirical

---

[5] Nothing against those cases (at least, not on these grounds of their apoliticality alone). The point is just that this work isn't particularly *applied* in the sense identified in the previous section, despite fitting the present characterization of social epistemology.

ones about whether the various general conditions for, say rational belief (or moral rightness, or whatever) are, in fact, satisfied?

Here different theorists might give different answers. Some—the particularists—will be skeptical that perfectly general (and informative, non-trivial) answers to epistemological or ethical questions are there to be found, and *a fortiori* about the possibility of mechanically applying such principles to applied contexts. Others will think that while there are general principles to be found, figuring out how they apply to applied contexts is a partly philosophical, and not purely empirical, task. For example, maybe it's generally true that we ought to show respect for others' humanity, or that we ought to take account of our own epistemic limitations. But seeing exactly what that requires in some particular situation seems to require more than just empirical data.[6] Others may accept that (the true) general principles plus empirical data mechanically yield verdicts about cases, but simply maintain that those verdicts about cases are the endgame of the project of formulating those general principles in the first place: by applying our general principles to cases, we can determine what they recommend in the cases people care about, and issue guidance on what to do (and believe).

Regardless of one's view about how mechanically general principles (plus empirical data) yield verdicts in particular cases, though, there is a different reason to do applied epistemology (and applied ethics). We can begin to discern it by drawing a familiar distinction between *explanatory* and *epistemic* priority. It's often true that A is more explanatory fundamental than B, but nevertheless that judgments about B are more epistemically secure than (direct) judgments about A and so constitute a better starting-point for inquiry. (Consider the rules of grammar and judgments about the grammaticality of individual sentences, for example.) So too with epistemological (and ethical) principles and judgments about cases. If a theory yields unintuitive results in individual cases, so much the worse for the theory. This is, of course, a commonplace, recognized even in the most abstract epistemology through the use of thought-experiments. But recently, some epistemologists have made the intriguing claim that it's especially important that our theories make predictions that we can live with *in the cases that are of deep significance to us*, and in cases with socio-political resonance.[7] Epistemic concepts are, after all, *for us*, and we should want them to do particular kinds of work: to enable us to diagnose clear epistemic pathologies, to cut through ideological obfuscations, and to help us expose pseudo-justificatory practices that simply retrench established power. Whether the received abstract epistemological theories are up to that task is an open question. In this way, doing applied epistemology is a way to reexamine and potentially rethink our most general epistemological ideas.

## IV.     Dangers in applied epistemology

---

[6] Here, though, things depend on what our answers to the most general answers are. If we're strict hedonic utilitarians in ethics, it's more plausible that the application of our principles to some case does just involve discovering, empirically, what will maximize pleasure in that case. An example of an epistemological theory with that same result might be reliabilism. If justified belief is just belief formed by a reliable process, it's more plausible that the application of our principles to some case does just involve discovering, empirically, how reliable some process of belief-formation is.

[7] E.g., Srinivasan (2020), Johnson King (2022).

Applied epistemology is a relatively new, but also rapidly expanding, field. Just a few years ago I would have said that the literature was relatively sparse, and now it is (like so many philosophical fields) expanding well beyond the bounds of what any one individual could possibly read. In a recent bibliography of applied epistemology work produced by the Applied Epistemology Project that I direct at UNC,[8] the *vast* majority of items are from 2018 or later. This no doubt reflects many factors, including an "applied turn" in philosophy more generally that has also taken place in other formerly "abstract" subfields such as philosophy of language, mind, and metaphysics; as well as the sudden prominence of epistemological ideas (such as "epistemic crisis", "alternative facts", "post-truth" and the like) in the wake of (among other things) Brexit and Trump. But the pace at which the field has expanded also means that, even if the literature is now much larger than it was, there is something of a lack of clarity about which are the more promising questions for applied epistemology to investigate, and which are the best methods for investigating them. A minority of work in the field, to be blunt, seems like it has found a *topic* of great interest without identifying a *question* about the topic that philosophy can productively illuminate.

There are some particular dangers that work in applied epistemology can risk, and needs to navigate past. One such risk is that it merely develops a more technical and inaccessible way to couch ideas and claims that are antecedently obvious to all, and that no-one needs an epistemologist to inform them of. Some, but not all, of the literature introducing innumerable concepts with names of the form "epistemic[ally] X," where X is some noun or verb, sails perilously close to this danger. Some such work does not address any question about which one could reasonably feel pulled in more than one direction; the answers are obvious, and what is contributed is only fancier ways to express them. For instance, it's obvious that claiming absolute certainty about something that is very uncertain on the evidence is irrational, or that exposing people to deliberately misleading evidence can undermine their knowledge and is thereby harmful. It's not clear we need a special technical vocabulary to express these commonplace thoughts, or that introducing such vocabulary achieves anything much—other than to generate publications.[9]

A second, quite different, danger is that work in applied epistemology can risk identifying questions that *are* interesting and not antecedently obvious, but that are essentially empirical—and then investigating them largely from the armchair, perhaps idealizing away from the messy empirical details. For example, the questions of *why* people believe conspiracy theories, distrust scientific authorities, or are seduced by misinformation and propaganda clearly can't be addressed by purely *a priori* methods. The same is true of questions of which institutional arrangements best promote the acquisition and spread of knowledge. It may be well that philosophers have something important to contribute to the resolution of these questions, in clarifying their conceptual underpinnings or generating hypotheses about their answers. But any work seeking to offer complete answers to these questions simply must

---

[8] See https://aep.unc.edu/resources-for-research/. The fantastic Genae Matthews led the work on this bibliography.
[9] There are other cases where novel terminology really does help us, because it allows us to recognize a unified class of phenomena that were previously hard to see (hence remedying "hermeneutical injustice" (Fricker 2007: ch. 7). But sometimes the phenomenon is plain to see without any special terminology.

be done in rich dialogue with social-scientific work about these questions (of which there is a great deal) to be intellectually responsible.

## V.      Overview of the volume

As I've just said, there are dangers that work in applied epistemology has to navigate. As can be expected, some work in the field falls prey to these dangers. But it is also possible to avoid them. In my opinion, the contributors to this volume, among numerous other scholars, have distinguished themselves in doing work that demonstrates the fecundity of applied epistemology and its capacity to yield deep insights about difficult questions. Their contributions to this volume are no exception, and represent some of the best work at the cutting edge of the field, across a very broad range of topics. Given the aforementioned lack of clarity about what the most fruitful research questions in the field are, my hope is that these papers taken as a whole will make a contribution toward identifying some particularly promising directions for future research, and crystallizing an agenda for the field.

As I've already mentioned, the topic of "echo chambers"—very roughly, epistemic environments (on-line and off-) that simply reinforce our existing beliefs while excluding or discrediting dissenting voices—has been a central topic both in popular discourse about the pathologies of the contemporary epistemic landscape, and in the applied epistemology literature. The received wisdom tells us that echo chambers are deeply harmful both to the rationality of our resulting beliefs and to social harmony and mutual understanding. In previous work (Avnur 2020), Yuval Avnur developed a sophisticated explanation of why echo chambers are harmful in the former way, appealing to the idea that they exacerbate our already-latent tendency toward what psychologists call "motivated reasoning," the tendency to process evidence in a biased way in order to reach a desired conclusion. His work on the topic expertly brought together the philosophical and psychological literature, drawing attention to empirical phenomena of which many philosophers have been ignorant while also drawing out philosophical and normative nuances of these phenomena that empirical researchers have passed over. Since the publication of Avnur's original work on this topic, however, a number of philosophers have argued, against the received wisdom, that echo chambers are *not* inherently epistemically problematic.[10] Avnur's contribution to this volume simultaneously represents a novel development and evolution of his previous view and a response to those who have challenged its conclusion. He finds inspiration in a perhaps surprising source: the 17th century French philosopher Blaise Pascal. While Pascal is certainly best-known for his famous "wager" concerning the existence of God, Avnur shows that he has much broader epistemological contributions to make, with a special emphasis on the role of the *heart* in our belief-formation process. This fits with the motivated reasoning account, since our heart generates motivations that shape, in predictable and systematically biased ways, our reasoning. But in the present paper, Avnur shows us that, for Pascal, the heart also affects the way we *perceive* or *see* the world in the first instance. Avnur explores the ways in which echo chambers can exacerbate this phenomenon. This complicates the picture about the rationality of the beliefs that result from being in an echo chamber at least somewhat, since the way we *perceive* the world

---

[10] Lackey (2021), Fantl (2021), Levy (2023), Begby (forthcoming).

affects the nature and content of the evidence we have, and hence, one might think, what it's rational to believe on the basis of that evidence. Yet Avnur argues, drawing on the work of Susanna Siegel (2017)—whose work is also represented in this volume; see below—that these perceptions themselves, and even the affective and motivational states that influence them, can be irrational, and that the resulting beliefs can inherit this irrational status. The result is a sophisticated and nuanced defense of the view that beliefs resulting from echo chambers are, at least by tendency, irrational.

Another central applied epistemology topic is disinformation. A cluster of recent literature in both social science and philosophy has examined a particularly scary and epistemically pernicious form of disinformation: the deepfake, whereby a video is produced in which some person (typically a public figure) is made to appear to be saying or doing something that they never, in fact, said or did.[11] In her contribution to this volume, Megan Hyska diagnoses what makes deepfakes so epistemically pernicious and why it matters. On the view she articulates and argues for, deepfakes undermine the acquisition of *collective* knowledge, which in turn undermines the possibility of collective *action*, and the potential for videography to contribute political mobilization. Hyska draws on a deep knowledge of the history of video media to articulate a rich, nuanced and informative account of the ways in which deepfake technology has these effects.

Various prominent domains of applied epistemology—for example, the epistemology of law and the epistemology of prejudice—have to grapple with thorny issues about statistical evidence and its epistemic import (or lack of). Of particular interest have been inferences from an individual's membership of a demographic group plus evidence about the proportion of members of that demographic group who bear a certain property, to the conclusion that the individual has the property. Here we can think about cases from the relatively innocuous—such as predicting an individual's music taste—to the extremely high-stakes—such as trying to predict an individual's likelihood of defaulting on a loan for lending purposes, or even trying to predict the likelihood that they committed a crime of which they are accused—on the basis of their demographic information plus macro-level demographically-sorted statistics. In many such cases, making such predictions on the basis of race, gender, or class strikes us as at best distasteful, and at worse deeply morally wrong. Yet there is a challenge to say what, if anything, is epistemically wrong about making use of such information.

In their contribution to this volume, Zoë Johnson King and Boris Babic continue, generalize, and extend an ongoing collaboration in which they attempt to find a purely epistemic flaw with such statistically-based inferences. Though other philosophers have also attempted this,[12] Johnson King and Babic challenge an assumption that is widely taken for granted throughout the literature—namely that statistical evidence can easily be highly probabilifying, and that any attempt to say it can't be would have to defy Bayesian orthodoxy. The complication for this assumption, as they point out, is that in real-world, non-stipulated cases, our statistical information may be flawed in various ways. Not only can it be incorrect in the sense that is simply an inaccurate report of the existing sample; it can also be correct as a report of the existing sample but unrepresentative of broader trends, either for systematic

---

[11] See e.g., Fallis (2020), Rini (2020), Harris (2021), Habgood-Coote (2023).
[12] E.g. Munton (2019), Gardiner (2020).

reasons or by chance. It might be naïvely thought that this is not important because statistical information is no more likely to be mistaken in one direction than in another, and so the risks of error "cancel out." But Johnson King and Babic rigorously and ingeniously show through formal and computational modeling that this is not necessarily so. The upshot is as follows. Suppose that across 100 observed cases of some crime C, 91 of them were perpetuated by a member of some group G. The probability that in a new instance of C, C was perpetuated by a member of G may nevertheless be significantly less than 0.91—perhaps only negligibly higher than 0.5—*even in the absence of any additional (non-statistical) evidence* individualized to the particular case. As well as having deep relevance for debates about the import of statistical evidence, Johnson King and Babic's paper exposes a case study in how (as I've already mentioned) idealizing away from messy, real-world circumstances and then applying our findings back to those same circumstances is fraught with dangers.

As I've also already mentioned, putative epistemic norms appear not just in the writings of philosophers but also in ordinary discourse. One of the most striking recent examples of this is provided by the #Metoo movement and its slogan "Believe Women." Some philosophers have explored the possibility of trying to vindicate this slogan under the framework of moral encroachment, according to which moral considerations can make a difference to what we are *epistemically* justified in believing.[13] In her contribution to this volume, Stephanie Leary argues that the moral encroachment framework is poorly positioned to do this: first, it's doubtful that moral encroachment can explain why we're justified in believing sexual assault allegations across the full range of plausible cases; and second, even if it does, it does so only on grounds that are demeaning toward victims. Leary's method is an excellent example of the approach that I mentioned at the end of §III, on which an important test for the adequacy of general abstract theories like moral encroachment is whether they yield acceptable results in cases that matter to us a great deal, like beliefs about sexual assault allegations. If what she argues is right—and she makes a very strong case—moral encroachment fares poorly on this score. Those who wish to advance an encroachment diagnosis of the cases at hand will be required to reckon with her lucid, compelling arguments.

Over the last several decades, numerous epistemologists—particularly but not exclusively feminist epistemologists—have drawn our attention to the importance of focusing our epistemological theorizing not just on knowledge but also on ignorance.[14] Often we want to say that people are criticizable—perhaps distinctively epistemically criticizable—for their ignorance. But there is a very difficult question about *when* we are criticizable for our ignorance. We cannot know or learn about everything, and often, that seems just fine. But other times, it doesn't. Moreover, this can happen even in cases where one does not yet possess the evidence that would put one in a position to justifiably form a belief—one can be criticizable for having failed to gather the evidence that would put one in such a position; for example, in many cases of prejudicial ignorance or insensitivity.

In her contribution to this volume, Jessie Munton deftly takes on the important and difficult question of when ignorance is epistemically criticizable. On the view she offers, roughly, ignorance is

---

[13] E.g. Crewe & Ichikawa (2021), Lloyd (2022).
[14] For just a small sampling, see e.g. Bailey (2007), Mills (2007), Pohlhaus (2012), and Goldberg (2017).

criticizable when it is *resilient* in the sense that it is maintained even when the cost of remedying it is low and the epistemic reward of remedying it is (obviously) high. Munton argues that resilient ignorance is bound to violate important norms on attention. And she argues that prejudicial ignorance tends to be resilient, so that her account explains its epistemic criticizability. Her closely and persuasively argued paper makes real progress on what might seem to be an intractable question.

In touching on the question of what we ought to gather evidence about, Munton's paper makes contact with the recent "zetetic turn" in epistemology, which urges us to evaluate and theorize about not just doxastic states like belief and credence but also processes of inquiry ('zetetic' being the adjective for that which pertains to inquiry).[15] Susanna Siegel's contribution to this volume more directly and explicitly deals with zetetic norms. Though there has been a lot of theorizing about zetetic norms of late, Siegel's paper is unique—and uniquely applied—among such paper in discussing the practice of *journalistic* inquiry (and the way that journalism *directs* inquiry) to yield broader insights about zetetic norms. Siegel commandingly discusses several richly detailed case studies, including coverage of the breakdowns in water treatment facilities in Jackson, Mississippi in 2022 and the *Kansas City Star*'s racist crime coverage over a period of 140 years (for which it issued a public apology in 2020). Her analysis of these cases and others illustrates two central claims that run through her paper: first, that we can understand the way that inquiry unfolds in terms of a practice of "treating things as zetetically relevant to questions," and that this practice can be normatively evaluated in a distinctively zetetic way. The genius of Siegel's attention to journalistic practice is that it is one of the arenas in which the process of inquiry has, outside philosophy, been most richly detailed and theorized; it's thus not surprising that it yields such rich and important insights for zetetic epistemology.

As I've already alluded to, deference to experts and partisanship are each central topics both in applied epistemology and in broader social discourse about our social condition. In the final contribution to this volume, Elise Woodard examines a question at the intersection of these two topics: is it epistemically permissible to (disproportionately) defer to co-partisans? While several prominent epistemologists have recently argued that it is,[16] Woodard argues that partisan deference is both morally and epistemically problematic. She argues compellingly that shared partisan allegiances do not always reflect shared values, and that even if they did, the fact that someone shares my values is not a strong ground for thinking them a reliable testifier. If Woodard is right, her conclusion raises interesting and difficult questions about deference more broadly. As several philosophers have argued, it is very hard for ordinary people to figure out whom to defer to when there is disagreement between experts. As such, co-partisanship might seem to be one feasible ground on which to choose between competing experts. Yet if Woodard is right, it is not an acceptable ground. In light of that, we are back with the question of how we *should* choose between disagreeing experts, if we should at all.

Taken together, then, these papers span a wide range of different *topics*—echo chambers, disinformation, statistical evidence, moral encroachment, ignorance, inquiry, and partisan deference—

---

[15] For an overview, see Friedman (forthcoming).
[16] E.g. Rini (2017), Levy (2022).

in connection with a wide range of different applied *domains*—politics, law, sexual assault allegations, racial prejudice, and journalism. They illustrate the potential for epistemology to illuminate some of the topics that we care about most, and that are most culturally salient in the present moment. Yet they provide only a small sample of the range of applications to which epistemological theorizing can be put. If applied epistemology continues to grow and develop at its current rapid pace, we can expect many more such productive and illuminating applications.

## References

Avnur, Y. (2020). "What Wrong with the Online Echo Chamber: A Motivated Reasoning Account," *Journal of Applied Philosophy*, 37(4), 578-593.

Bailey, A. (2007). "Strategic Ignorance," in Sullivan and Tuana (eds.), *Race and Epistemologies of Ignorance*. SUNY Press.

Begby, E. (forthcoming). "From Belief Polarization to Echo Chambers: A Rationalizing Account," *Episteme*.

Coady, D. & Chase, J. (2021). "The Return of Applied Epistemology," in Chase & Coady (eds.), *The Routledge Handbook of Applied Epistemology*. Routledge.

Crewe, B. & Ichikawa, J.J. (2021). "Rape Culture and Epistemology," in Lackey (ed.), *Applied Epistemology*. Oxford University Press.

Fallis, D. (2020). "The Epistemic Threat of Deepfakes," *Philosophy and Technology*, 34(4), 623-643.

Fantl, J. (2021). "Fake News vs. Echo Chambers," *Social Epistemology*, 35(6), 645-659.

Fricker, M. (2007). *Epistemic Injustice*. Oxford University Press.

Friedman, J. (forthcoming). "Zetetic Epistemology," in Flowerree & Reed (eds.), *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge.

Gardiner, G. (2020). "Profiling and Proof: Are Statistics Safe?," *Philosophy* 95, 161-183.

Goldberg, S. (2017). "Should Have Known," *Synthese* 194(8), 2863-2894.

Habgood-Coote, J. (2023). "Deepfakes and the Epistemic Apocalypse," *Synthese* 201(103).

Harris, K.R. (2021). "Video on Demand: What Deepfakes Do and How They Harm," *Synthese* 199(5-6), 13373-13391.

Johnson King, Z. (2022). "Radical Internalism," *Philosophical Issues*, 32, 46-64.

Lackey, J. (2021). "Echo Chambers, Fake News, and Social Epistemology," in Bernecker, Flowerree & Grundmann (eds.), *The Epistemology of Fake News*. Oxford University Press.

Levy, N. (2022). *Bad Beliefs: Why They Happen to Good People*. Oxford University Press

------ (2023). "Echoes of Covid Misinformation," *Philosophical Psychology*, 36(5), 931-948.

Lloyd, A. (2022). "#MeToo and the Role of Outright Belief," *Ethical Theory and Moral Practice*, 25(2), 181-197.

Mills, C. (2007). "White Ignorance," in Sullivan and Tuana (eds.), *Race and Epistemologies of Ignorance*. SUNY Press.

Munton, J. (2019). "Beyond Accuracy: Epistemic Flaws with Accurate Statistical Generalizations," *Philosophical Issues*, 29(1), 228-240.

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

Rini, R. (2017). "Fake News and Partisan Epistemology," *Kennedy Institute of Ethics Journal*, 27(S2), 43-64.

------ (2020). "Deepfakes and the Epistemic Backstop," *Philosophers' Imprint*, 20(24).

Siegel, S. (2017). *The Rationality of Perception*. Oxford University Press.

Srinivasan, A. (2020). "Radical Externalism," Philosophical Review, 129(3), 395-431.

Toole, B. (ms.) "Standpoint Epistemology – Social or Applied?"