

# Akrasia and Epistemic Impurism

**ABSTRACT:** *This essay provides a novel argument for impurism, the view that certain non-truth-relevant factors can make a difference to a belief's epistemic standing. I argue that purists, unlike impurists, are forced to claim that certain 'high-stakes' cases rationally require agents to be akratic. Akrasia is one of the paradigmatic forms of irrationality. So purists, in virtue of calling akrasia rationally mandatory in a range of cases with no obvious precedent, take on a serious theoretical cost. By focusing on akrasia, and on the nature of the normative judgments involved therein, impurists gain a powerful new way to frame a core challenge for purism. They also gain insight about the way in which impurism is true: my argument motivates the claim that there is moral encroachment in epistemology.*

**KEYWORDS:** akrasia, impurism, moral encroachment, pragmatic encroachment, practical rationality

A great deal of recent work in epistemology concerns the following claim:

Impurism: Some paradigmatically epistemic properties of a belief that  $p$  (like whether  $p$  is epistemically rational, or whether it is knowledge) depend on factors that are not relevant to the truth of  $p$ .

At first glance, impurism may seem unattractive. If a factor has nothing to do with the truth of  $p$ , how could it make a difference to the epistemic rationality of belief in  $p$ ? How could it make a difference to whether she knows that  $p$ ?

In order to illustrate how impurism might be true, many theorists follow in the footsteps of Keith DeRose (1992: 913); they offer case pairs like the following one.

**Parked Car Low Stakes.** Ava parked her car four hours ago, and she cannot see it from where she is currently sitting. Ava's friend Emil, a reliable testifier, points out that, if her car is parked illegally, she will almost certainly get a written warning. Ava thinks back, and she seems to remember (although not too vividly) that she parked it legally. She forms the belief that her car is currently parked legally, and she remains sitting in her easy chair.

For helpful discussion, I am grateful to Mike Ashfield, Ethan Brauer, Patrick Croskery, Justin D'Arms, Jenni Ernst, Brian McLean, Julia Jorati, Matthew Shields, Keshav Singh, several anonymous referees, and audiences at meetings of the Eastern APA and the Ohio Philosophical Association. Special thanks to Tristram McPherson and Declan Smithies, who provided invaluable help at every stage of the drafting process.

Parked Car High Stakes. César parked his car four hours ago, and he cannot see it from where he is currently sitting. César's friend Maryam, a reliable testifier, tells him that, if his car is parked illegally, his car will almost certainly be towed. César knows that, if his car is towed, he will be late to an extremely important event. César thinks back, and he seems to remember (although not too vividly) that he parked it legally. He forms the belief that his car is currently parked legally, and he remains sitting in his easy chair.

The only significant difference between Ava's case and César's case is the severity of the penalty for illegal parking. And the severity of the penalty for illegal parking is not *relevant to the truth* of the proposition that a car is parked legally. In other words, it makes no difference to the probability of that proposition, either from the believer's point of view or from any more objective point of view. (This gloss on *truth-relevance* follows DeRose [2009: 25] and Roeber [2020: 2650n1].) But if impurism is true, then this non-truth-relevant factor might make a difference to an epistemic property. Perhaps, for instance, Ava's belief amounts to knowledge while César's does not.

Case pairs like the one above illustrate how impurism might be true. But should we believe that impurism is true? What are the best arguments in favor of and against impurism?

Blake Roeber (2020: 2650, 2018: 173) sorts existing arguments for impurism, into 'intuition-based arguments' and 'principle-based arguments'. The most prominent intuition-based argument for impurism is offered by Jason Stanley (2005). Stanley suggests that our intuitions about the aptness of knowledge attribution in a variety of example cases can best be accommodated by impurism. Intuition-based arguments can be (and have been) challenged in two ways: one can attempt to debunk the claim that we have the relevant intuitions (see Rose et al 2017), or one can argue that another view provides a better fit with those intuitions than impurism does (see DeRose 2009: ch. 7).

Principle-based arguments in favor of impurism, unlike intuition-based arguments, begin by making a case for exceptionless principles connecting knowledge to certain truth-irrelevant factors. John Hawthorne and Jason Stanley (2008: 578), for instance, defend the principle that 'Where one's choice is *p*-dependent, it is appropriate to treat the proposition that *p* as a reason for acting iff you know that *p*.' Jeremy Fantl and Matthew McGrath (2009: 66), on the other hand, defend the principle that 'If you know that *p*, then *p* is warranted enough to justify you in  $\phi$ -ing, for any  $\phi$ .' Principle-based arguments go on to claim that, if the principle in question is true, impurism follows. Purists like Jessica Brown (2008) and Baron Reed (2010) push back on these arguments by proposing putative counterexamples. Whether these examples show the principles to be hopeless, or instead simply point to interesting, defensible results, is a matter of some controversy.

One can expand Roeber's taxonomy by calling attention to a class of arguments for impurism that appeal to the *function* of knowledge attributions in social practice. Call these *function-based arguments*. Matthew McGrath (2015: 150) for instance,

argues for impurism from the insight that knowledge attribution allows us to communicate our evaluations of agents' actions (see also Fantl and McGrath 2007: 561–64). And both Stephen Grimm (2015) and Michael Hannon (2017) support their impurist views by appealing to the notion, inspired by Edward Craig (1999), that a primary function of knowledge discourse is to allow participants to flag reliable informants.<sup>1</sup> (Note that some function-based arguments are also principle-based arguments: Fantl and McGrath [2007] and Hawthorne and Stanley [2008], for instance, use the social role of knowledge ascriptions as evidence for knowledge-action principles.) Purists might resist function-based arguments by offering alternative stories about the function of knowledge discourse; Mikkel Gerken, for instance, argues that knowledge discourse is at most a 'reasonably accurate communicative heuristic' (2015: 156).

In this essay, I offer a novel argument for impurism. This argument, unlike those just surveyed, does not rely on much-contested intuitions about cases, an exceptionless principle connecting knowledge to action, or an account of the function of knowledge discourse. My argument instead supports impurism by drawing attention to an underappreciated connection between purism and *akrasia*: certain high-stakes cases force us to choose between impurism and rationally required *akrasia*.<sup>2</sup> Since the only way to defend purism is to embrace rationally required *akrasia*, purism comes along with a serious theoretical cost.

One of the upshots of my argument, then, is that a focus on the irrationality of *akrasia* can help impurists to paint a uniquely compelling picture of the costs of purism. Another upshot is that there is *moral encroachment* in epistemology.

In section 1 of this essay, I provide a picture of the sort of normative judgment involved in *akrasia*. In section 2, I argue that, a few interesting exceptions aside, *akrasia* generally involves a problematic sort of irrationality. When a theory implies that *akrasia* is rationally required in a range of cases that have no obvious precedent, then, the theory takes on a serious theoretical cost. In section 3, I show that purists, unlike impurists, take on that serious theoretical cost. In section 4, I respond to objections. Finally, in section 5, I explain why my argument provides a case for moral encroachment, and I explain some advantages of my argument over related arguments for impurism.

## 1. Objective and Rational Obligations

*Akrasia*, paradigmatically, involves a person who simultaneously believes that she ought to take some action and fails to take that action. There are at least two sorts of 'ought' belief that can feature in *akratic* action: beliefs about *objective* obligation and beliefs about *rational* obligation. Rational obligation, unlike objective obligation, is sensitive to the ethically significant properties of merely

<sup>1</sup> Thanks to an anonymous referee for the suggestion to expand this taxonomy.

<sup>2</sup> For doubts as to whether the notion of 'stakes' can be made usefully precise, see Worsnip (2015) and Anderson and Hawthorne (2019). I remain neutral on this dispute; with 'stakes,' I mean to refer to the non-truth-relevant factor, whatever it is, that impurists should say makes a difference for knowledge.

epistemically possible circumstances. More loosely put, rational obligation, unlike objective obligation, is sensitive to risk.

Cases of incomplete or misleading evidence provide support for the distinction between objective and rational obligations. Some ‘ought’ claims about such cases take into account epistemic limitations. But some other ‘ought’ claims do not.

For example, imagine that you must choose to take home one of two gift bags. Your evidence suggests that bag A contains the better gift, but your evidence is misleading; the better gift is in bag B. Which bag ought you choose? In one sense, you ought to choose bag B, since it contains the better gift. This, in my terminology, is a claim about what you *objectively* ought to do. But it is very plausible that there is another sense in which you ought to pick bag A. This sort of ‘ought’ claim, unlike a claim about what you objectively ought to do, is sensitive to your epistemic position. I call the latter sort of ‘ought’ claim a claim about what you *rationally* ought to do. In adopting this terminology, I am not taking a stand about the uniquely best way to use the term *rational*. I use the term *rational*, rather, to pick out a notion that is both clearly of theoretical importance and familiar from philosophical discussion of appropriate action.

What a person rationally ought to do, then, is sensitive to mere epistemic possibilities; what she objectively ought to do is not. But this does not settle the question of what flavor of normativity is in play. There are many norms on action that take epistemic position into account. There is a norm that says what would be most *prudent*, given your epistemic position; there is a norm that says what would be *morally* best, given your epistemic position; perhaps there’s even a norm that says which move would best promote victory in chess, given your epistemic position. For the purposes of this essay, I set these norms on action aside. The ‘ought’ that concerns me weighs *all* these considerations—prudential, moral, and so on—together. This is sometimes called the all-things-considered ‘ought,’ ‘ought’ *simpliciter*, or, in Philippa Foot’s memorable turn of phrase, the ‘free floating and unsubscripted’ ‘ought’ (1997: 320n15). (For doubts about the notion of such an ‘ought’, see Tiffany [2007] and Baker [2018]. For defense, see Thomson [2001: 46] and McPherson [2018].)

Now, distinguish between two sorts of ‘all-things-considered’ obligations. One sort of obligation is not relativized to the agent’s epistemic position; it takes into account *all* of one’s reasons for action (prudential, moral, and otherwise). This is what I call *objective* obligation. The other sort is sensitive to one’s epistemic position; though it weighs up different sorts of reasons, including prudential reasons, moral reasons, and so on, it only takes into account reasons that are, in some sense, within the subject’s ken. This is what I call *rational* obligation. (I do not claim that the reasons relevant to rational obligation are a subset of the objective reasons. They might be related to objective reasons in a looser way; see Sylvan [2015] and Wodak [2019] for useful discussion.)

To make this distinction more concrete, consider the following ‘high-stakes’ case:

Naomi’s Medical Supplies. It is Friday afternoon, and Naomi is on her way to a dinner party. As she left work, her boss gave her a package of medical supplies and asked her to drop it off at the post office.

Based on her memory of the interaction, she is very confident—suppose she has a rational credence of 0.95—that the package she is carrying is Package A. Package A contains some cough suppressant, and no one is counting on receiving it particularly soon. But Naomi's memory leaves open a slim possibility—suppose she has a rational credence of 0.05—that the package she is carrying is Package B. Package B contains life-saving medicine, and if it is not delivered today, five innocents will soon die for lack of the medicine. Naomi has no way to get more evidence about which package she is carrying.

Naomi sees a long line at the post office. If she waits in line, it will make her late to her dinner party. She has promised to be at the party on time, and breaking that promise would upset her and several of her closest friends. But if she goes straight to the dinner party, she will not be able to drop off her package until tomorrow. Naomi chooses to go straight to the dinner party.

As it turns out, Naomi is carrying Package A. She attends the dinner party on time, and she returns to the post office to mail her medical supplies later on, in plenty of time to meet her boss's expectations.

When she chooses to pass the post office by, does Naomi do what she ought to do? In one sense, she does: she performs the action that is most choiceworthy, given *all* the facts. Since she is carrying Package A, she can best meet her multiple duties (that is, her duties of promise keeping, her professional duties, and her duties of aid to others) by passing the post office by. In other words, she does what she *objectively* ought to do. But, of course, what Naomi does is also unacceptably risky. The responsible course of action for her, given the possibility that innocent lives will be lost unless she drops off the package today, is to wait at the post office. In other words, she does not do what she *rationally* ought to do.

As Naomi's case illustrates, a person's rational obligations can depend on the ethically significant properties of *merely epistemically possible* circumstances. Naomi's carrying Package B is a merely possible circumstance. But in that possible circumstance, Naomi's choice to pass by the post office would have extremely serious consequences; five innocents would die. Plausibly, this explains in part why Naomi rationally ought to wait in line; in a 'lower-stakes' version of the case, where nothing of ethical importance even possibly hangs on her sending the supplies, she would be rationally permitted (indeed, rationally *required*) to drive straight to the dinner party.

Naomi's case also illustrates a further point: by and large, objective obligations are *not* sensitive to ethical facts about merely epistemically possible circumstances. The best thing for Naomi to do, given the way the world *actually* stands, is to go to the dinner party on time. Varying the ethical properties of merely possible circumstances does nothing to change this; even if Package B were medicine needed to save hundreds of innocents, rather than only five, the action best supported by her *actual* choice situation would still be to pass the post office by. Rational obligations, then, are sensitive to ethical features of merely possible circumstances. But objective obligations need not be.

## 2. The Presumption against Rationally Required Akrasia

Having discussed the rationality of action, I turn to questions about the epistemic rationality of *belief*. I claim that, a few exceptions aside, akrasia involves either epistemically irrational belief or irrational action.

Epistemic rationality, on my usage, is more intimately tied to knowledge than are other available notions of rationality. Suppose, for instance, that I can earn a huge monetary reward if I believe that Madrid is the capital of Australia. Though the reward makes it *desirable* for me to form the false belief, there is also an important norm on belief formation that forbids my doing so. When I make claims about rational belief, I mean to be picking out this restricted norm on belief formation. Importantly, it is a norm that is not sensitive to considerations like direct threats or bribes for belief. (For fuller defense of this distinction, see Kelly [2002]; for worries about it, see Rinard [2017].) In what follows, I argue for a surprising conclusion: that even this more austere notion of epistemic rationality is subtly sensitive to certain non-truth-relevant ethical facts. This is not, importantly, to collapse the distinction between *epistemically rational belief* and *belief that it is best to have*.

My argument for impurism relies on the claim that, a few interesting exceptions aside, akrasia involves irrationality. More precisely, the requirements of epistemically rational belief and the requirement to act rationally do not usually conspire to require akrasia. Specifically, my focus is on a particular sort of akrasia: first-personally believing that one objectively ought to  $\phi$ , while failing to  $\phi$ .

It is tempting, at first, to think that rationality never requires this sort of akrasia. In fact, some argue that rationality never *permits* akrasia; for an example, see Declan Smithies (2019: ch. 8). To see the force of this idea, consider an example. Imagine that you are trying to figure out what you ought to do this afternoon. After some reflection, you conclude that you objectively ought to go to the grocery store. Further, imagine that your belief is rationally appropriate. Now try to imagine that, in the same case, it is rationally impermissible for you to go to the grocery store. Something seems to have gone wrong; if your belief is appropriate given your epistemic position, then surely your epistemic position does not also make it inappropriate to act as your belief suggests you should!

Reflections like these provide *prima facie* support for the following principle:

Belief-action link. If epistemic rationality requires you to believe that you objectively ought to  $\phi$ , then it is rationally permissible for you to  $\phi$ .

The belief-action link is tempting. But, as written, it is too strong. In what follows, I offer three reasons to think that the belief-action link does not hold in full generality. Though it is important to note these possible exceptions, it is also crucial to see why they cause no trouble for my argument: none of them casts doubt on my claim that, *a few interesting exceptions aside*, rationality does not require akrasia.

In the first class of exceptions to the belief-action link, rational beliefs about one's objective obligations do not settle the question of what to do. The much-discussed

case of the miners (Kolodny and MacFarlane 2010; Parfit 2011) is a paradigmatic example: ‘Ten miners are trapped either in shaft A or in shaft B, but we do not know which. Flood waters threaten to flood the shafts. We have enough sandbags to block one shaft, but not both. If we block one shaft, all the water will go into the other shaft, killing any miners inside it. If we block neither shaft, both shafts will fill halfway with water, and just one miner, the lowest in the shaft, will be killed’ (Kolodny and MacFarlane 2010: 115).

In this case, what should we believe about our obligations? Well, it is epistemically rational for us to believe that we *rationaly* ought to block neither shaft. But what does epistemic rationality tell us to believe about our *objective* obligations?

At a first pass, rationality recommends suspending judgment about our objective obligations. At a second pass, things are more complicated. Our epistemic position surely supports believing that we objectively ought to *block whichever shaft contains the miners*, thereby saving all of them. The epistemic rationality of this belief makes problems for the belief-action link. Even if epistemic rationality requires me to believe that I objectively ought to block whichever shaft contains the miners, it is clearly not rationally permissible for me to block the shaft that in fact contains the miners; indeed, the only rationally permissible option for me is to block neither shaft. This example, then, shows that we can sometimes act against our beliefs about what we objectively ought to do without the irrationality characteristic of akrasia.

A defender of the belief-action link might respond by revising her principle, perhaps as follows:

Belief-action link\*. If epistemic rationality requires you to believe of a particular action,  $\phi$ , that you objectively ought to perform it, and  $\phi$ ing is a live option for you, then it is rationally permissible for you to  $\phi$ .<sup>3</sup>

There are questions about how to understand the notion of a ‘live option’. But we can set these questions aside, because the revision does not solve the problem raised by the case of the miners. Granted, this revised belief-action link\* does not threaten the result that it is rationally permissible for me to *block the shaft that in fact contains the miners*; it is plausible that, in some important sense, blocking that shaft is not a live option for me. But consider a different action: *blocking one of the two shafts*. My epistemic position supports believing that I objectively ought to take this action, and on any immediately attractive way of construing ‘live option’, it is a live option for me. Nevertheless, I am not rationally permitted to take the action.

The case of the miners, then, shows that the belief-action link does not hold in full generality. Some might be tempted to draw a bolder conclusion from the case: beliefs about objective obligations are *never* sufficient to generate *akratic* tension. Perhaps akrasia, rightly considered, is always a tension between action and belief about *rational* obligation—and never a tension between action and belief about *objective* obligation. But this line of thought moves too quickly. It is true that we can hold

<sup>3</sup> Thanks to an anonymous referee for suggesting discussion of this complication.

*certain* beliefs about objective obligations without thereby settling on a course of action. But other beliefs about objective obligations *do* fully settle the question of what to do, and do so in a way that can give rise to *akratic* tension. Suppose, for instance, that I form the belief that I objectively ought to block shaft A, but fail to block it. This seems like a paradigmatic instance of akrasia, even if we stipulate that I lack any beliefs about my rational obligations. So even though there are *some* examples of belief about objective obligations that do not seem apt to give rise to akratic tension, there are others that do. Loosely speaking: whenever I adopt a sufficiently specific, rich set of beliefs about what I objectively ought to do, under a sufficiently clear, informative description, I see myself as obligated to follow a particular course of action. Failures to follow that course of action involve problematic akrasia.

The second class of putative counterexamples to the belief-action link involves cases of misleading higher-order evidence. Some, including Allen Coates (2012) and Maria Lasonen-Aarnio (2014), hold that misleading higher-order evidence can make akrasia rational. Suppose, for instance, that some expert mathematicians tell you that you have done a math problem incorrectly. Further suppose that your first-order mathematical belief was formed rationally. In this case, some think, it is both rational for you to believe that you ought to abandon your first-order mathematical belief *and* rational for you to retain that belief. This would (on the assumption, which I grant for the sake of argument, that the variable  $\varphi$  in the belief-action link can be filled by a belief) constitute a counterexample to the belief-action link.

This ‘level-splitting’ approach to higher-order evidence is controversial (for criticism of the view, see Titelbaum [2015]; Smithies [2019: ch. 9]). But for the sake of argument I grant that it is on the right track. Even if examples involving misleading higher-order evidence provide examples of rational akrasia, everyone should agree that they do nothing to cast doubt on irrationality of akrasia more generally.

Finally, there might be counterexamples to the belief-action link in cases that involve misleading normative evidence. Some, including Coates (2012), Elizabeth Harman (2015), and Brian Weatherson (2019: ch. 3), have argued that misleading normative evidence does not make a difference to a person’s rational obligations. To see why, first note that misleading *non*normative evidence can render wrongful action blameless. For instance, suppose that I feed you poison, but I had excellent evidence that it was in fact medicine. In this case, though my action is unfortunate, I do precisely what I *rationally* ought to have done, and my action is exculpated—blameless.

Next, note that it is less obvious that misleading *normative* evidence can render wrongful action blameless in the same way. Harman (2015: 58–62), for instance, discusses Bob, who chooses not to teach his daughter to drive on the basis of misleading testimonial evidence about the appropriate place of women in society. Harman claims that, unlike the unintentional poisoner, Bob’s wrongful action is also blameworthy. Harman further suggests that there is a tight connection between rational obligation and blameworthiness: an agent is blameworthy only if she violates rational obligations (2015: 56–57). So Bob is rationally required to



teach his daughter to drive. Harman's approach to misleading normative evidence, in short, creates room for required akrasia in cases like Bob's: perhaps, though Bob has an epistemically impeccable belief about what he ought to do, he is rationally forbidden from acting on those beliefs.

Again, we may have a counterexample to the belief-action link; but, again, we have no reason to think that the counterexample will generalize. All parties to this debate, including Harman (2015: 58) and Weatherston (2019: ch. 5), acknowledge that akrasia is nowhere near as respectable in cases of mixed or misleading evidence about *nonnormative* matters as it is in cases of mixed or misleading evidence about *purely normative* matters. My argument relies only on the former sort of mixed evidence.

To this point I have offered several reasons (some more controversial than others) to suspect that the belief-action link does not hold in full generality. But each of these reasons seemed contained to a class of cases with a particular character. None seemed to cast doubt on the compelling idea that, a few interesting exceptions aside, akrasia involves irrationality.

Even those who defend *some* instances of rational akrasia, then, should be wary of theories that allow for new cases of rational akrasia, especially when those cases do not appear to have precedents in other, more familiar forms of rational akrasia. A theory that requires akrasia in unprecedented cases, in other words, thereby takes on a significant theoretical cost. In what follows, I argue that purism does just that.

### 3. An Argument for Impurism

Cases like Naomi's Medical Supplies present a problem for purist treatments of akrasia. Impurists can explain why cases like Naomi's do not rationally require akrasia, while purists cannot. In fact, purists must allow that some cases like Naomi's *do* rationally require akrasia. This is a significant theoretical cost for purism.

Recall the basics of the case: Naomi is carrying Package A, which does not contain life-saving medical supplies. But there is a slim epistemic possibility for her that she is instead carrying Package B, which contains life-saving medical supplies that urgently need to be dropped off at the post office. She cannot both drop the medical supplies off and keep her promise to attend a dinner party on time.

There is a strong *prima facie* case to be made for each of the following three claims about Naomi's case:

1. Naomi rationally ought to wait at the post office.
2. Epistemic rationality requires Naomi to believe *I objectively ought to drive straight to the dinner party*.
3. The rational requirements on Naomi's thought and action do not conspire to require *akrasia*.

But these three claims are in tension. Suppose that Naomi meets both of the requirements named in the first two claims: she waits at the post office while believing *objectively, I ought not do this; objectively, I ought to drive straight to the dinner party*. This is a paradigmatic instance of akrasia. So the first and

second claims jointly imply that the third claim is false; if the first two claims are true, Naomi is required to be akratic.

Why think, as I have claimed, that there is a strong prima facie case to be made for each of the three claims? Start with the first. Above (in section 1), I explained why this claim is so plausible, in part by distinguishing between objective and rational obligations. Rational obligations, I argued, are sensitive to ethically important error possibilities. Given the ethical importance that attaches to Naomi's sliver of doubt, she is rationally obliged to wait in line. (If you suspect that she is not, feel free to imagine a nearby case in which the stakes are higher.)

There is also a strong case to be made for the second claim. Naomi is rationally very confident that she is carrying Package A. We can stipulate that her epistemic position makes the proposition that she is carrying Package A probable in just the same way, and to just the same degree, that would usually be sufficient to rationalize beliefs about the items one is carrying. But the question of what Naomi is objectively required to do, here, hangs entirely on whether she is carrying Package A. Given that she is carrying Package A, she objectively ought to drive straight to the dinner party. (Compare: in the Gift Bag case from section 1, given that the better prize is in Gift Bag B, you objectively ought to select Gift Bag B.) So Naomi has very strong epistemic support, of an entirely banal kind, for the true proposition that she objectively ought to drive straight to the dinner party.

Above (section 2), I provided a prima facie case in favor of the third claim. Even if there are some cases in which akrasia is rationally required, there is a strong default presumption against rationally required akrasia. Moreover, Naomi's case does not seem to fit any of the three familiar case types presented above (section 2), in which (some have thought that) rationality requires akrasia. Her belief about what she objectively ought to do would not provide woefully incomplete guidance for her action (as would, for instance, the belief *I objectively ought to block whichever shaft the miners are in*). To the contrary, it would entirely settle the question of what to do in her choice situation. Nor is it afflicted with the peculiar force of higher-order evidence. Finally, Naomi's uncertainty about what she objectively ought to do derives entirely from a mixed body of *nonnormative* evidence about which package she is carrying, not from a mixed body of *normative* evidence. So the attractive rule of thumb that rationality does not require akrasia seems undefeated in this case.

Impurists can avoid the tension between the first and third claims by denying that the second is true. On an impurist view, the non-truth-relevant features of Naomi's choice situation can make a difference to epistemic standards. Suppose, for instance, that Naomi forms the belief *I objectively ought to drive straight to the dinner party*. Impurists can grant that this belief has epistemic support of just the sort that often suffices to make belief rational. But they can also claim that, in virtue of certain (non-truth-relevant) features of Naomi's choice situation, epistemic standards are unusually high for her in this case. So there is no epistemic requirement for her to believe that she objectively ought to drive straight to the dinner party. Rationality does not require her to be akratic. Impurists can rest assured that this move—

protecting against rationally required akrasia by noting variance in epistemic standards—is open to them in any case similar to Naomi’s.

Purists, on the other hand, face a challenge with respect to cases like Naomi’s. In order to avoid requiring akrasia in a wide range of ‘high-stakes’ cases, they must find a principled way to rule out the possibility that epistemic rationality requires Naomi, or anyone in a *structurally similar case*, to form the well-supported belief about her objective obligations.

Moreover, they cannot appeal to ethical features of the believer’s choice situation to provide this explanation. To see why, first note that the truth of the belief in question—Naomi’s belief about what she *objectively* ought to do—is settled entirely by the question of which package she is carrying. Since she is carrying Package A, she objectively ought to keep her promise and go to the dinner party on time; if she were carrying Package B, she would instead be objectively obligated to wait in line. Next, note that the ethical risks associated with the possibility that Naomi is carrying Package B are *not relevant to the truth* of the proposition that she is carrying Package A. In other words, the possible scenario in which Naomi is carrying Package B is a *risky* one, but that does not make it a more *probable* one (from Naomi’s perspective or from any more objective perspective). This means that the ethical risks in question also do not make more or less probable (from Naomi’s perspective or from any more objective one) the proposition *Naomi objectively ought to drive straight to the dinner party*. So, when it comes to this belief about objective obligations, the possible risk to the lives of five innocents is not truth-relevant. But purists, of course, hold that epistemic facts do not depend on non-truth-relevant factors. So purists, unlike impurists, cannot appeal to the risk to the lives of five innocents to justify rejecting the second claim I have made about Naomi’s case, much less the analogous claim about all relevantly similar cases.

Is there another way for purists to argue that epistemic rationality, in Naomi’s case, does not require outright belief? Well, they could argue that the *truth-relevant* features of Naomi’s case fail to make belief rational; perhaps, for instance, a rational credence of 0.95 in *p* never comes along with a requirement to believe *p*. But this simply reorients the purist’s explanatory task. We can stipulate an alternative version of Naomi’s case in which her epistemic position with respect to the proposition *I objectively ought to drive straight to the dinner party* makes appropriate a credence much higher than 0.95. Even in this new case, as long as Naomi’s sliver of doubt is sufficiently ethically fraught, rationality can still require her to play it safe by waiting at the post office. The threat of rationally required akrasia will still loom large.

Across a wide variety of ‘high-stakes’ cases, then, the purist faces uniform pressure to say that rationality requires a troubling sort of akrasia. The impurist, by contrast, faces no such pressure. This is a serious theoretical cost for purism.

#### 4. Objections

There are at least two strategies that purists are likely to adopt when resisting my argument. One strategy appeals to infallibilism about rational belief; another

stresses the rationality of following one's beliefs in rational obligations, rather than one's beliefs in objective obligations.

#### 4.1 Infallibilism about Rational Belief

One way for purists to avoid embracing akrasia is to argue that in any case like Naomi's epistemic rationality fails to require outright belief about one's objective obligations. Some purists might take up this strategy by pointing to Naomi's uncertainty. Perhaps, whenever a person is rationally required to form some belief, there is no epistemic possibility for her that the belief is false. Call this principle *infallibilism about rational belief*.

If infallibilism about rational belief is true, then whenever epistemic possibilities about some ethical proposition are divided for an agent, she is not rationally required to form a belief about that proposition's truth or falsehood. As a result, the question of whether an agent is required to believe that  $p$  is never affected by the sort of non-truth-relevant factors that arise in Naomi's case. If  $p$  is certain for her, then there is no chance that, as in Naomi's case, an ethically significant epistemic possibility that not- $p$  will have implications for her actions. If, on the other hand,  $p$  is *not* certain for her, then it is not the case that she rationally ought to believe that  $p$ .

There are two importantly different ways of developing infallibilism about rational belief. Neither offers the purist a promising way to reject the argument above. The first variety of infallibilism threatens to diminish our rational obligations in an implausible way. The second, when developed to avoid this problem, is entirely compatible with the spirit of the argument above.

First, consider an infallibilist view on which everyday, prosaic propositions are generally not certain for us. Take, for instance, my belief that Madrid is not the capital of Australia. If this proposition is not certain for me, then infallibilism says that I am not required to believe it. Similarly, the everyday proposition that Naomi is carrying Package A is not certain for her, and she is therefore not rationally required to believe that she is.

This sort of infallibilism rejects my argument for impurism at an unacceptable cost. Surely, there are plenty of facts about what a rational person would have to believe. The version of infallibilism we are currently considering seems unable to account for those facts. The conclusion that the rational requirements on our beliefs are so sparse seems, in fact, far less plausible than the conclusion that those requirements are sensitive to the choices we face.

Second, consider an infallibilist view on which many everyday propositions *are* certain for us. On this view, it is usually *not* possible for me that Madrid is the capital of Australia; in ordinary cases, that proposition is certainly false for me, and I am required to believe that it is false. This sort of infallibilism does not provide a promising way for the purist to avoid a commitment to rationally required akrasia in a troubling range of cases without precedent (see above, section 3). In order to avoid that result, this sort of infallibilism must allow that, in a wide range of 'high-stakes' cases, certain ordinary propositions are *not* certain for us. But, to avoid vitiating our epistemic obligations quite generally, this version

of infallibilism must also acknowledge that, outside this range of ‘high-stakes’ cases, many ordinary propositions *are* certain for us. What could guarantee that the relevant propositions are never certain in cases where akrasia looms? The impurist has an answer here: whether a proposition is certain for a person can depend on non-truth-relevant factors about her choice scenario. Perhaps the importance of hedging bets, in Naomi’s case, explains why it is not certain for her that she is carrying Package A. This is an attractive way for impurists to use infallibilist machinery to explain just why Naomi is not required to believe that she should drive straight to the party. The purist has no obvious route forward here.

Infallibilism, then, does not provide an attractive way to avoid the force of my argument. Some ways of developing the view are objectionably deflationary about rational obligation, and others simply relocate the problem for the purist.

#### 4.2 Following Beliefs in Rational Obligations

Above (section 2), I argue that beliefs about rational obligations need not be involved in akrasia; acting against one’s beliefs about objective obligations can be enough. But even those who grant this point might argue that, when one has *conflicting* beliefs about one’s rational obligations and objective obligations, there is nothing problematically akratic about following the former instead of the latter. Suppose, for instance, that Naomi believes that she objectively ought to drive to the party, but *also* believes that she rationally ought to wait in line at the post office. And suppose she chooses to guide her action in accordance with this latter belief. If so, she does not fail altogether to conform her actions to her believed obligations; she simply aims to conform her actions to her believed *rational* obligations instead of her believed *objective* obligations.

This point does not suffice to do away with the appearance of problematic akratic tension. The tension still arises because, in the words of Errol Lord, ‘deliberation aims at what’s best’ (2015: 44). If I have a view of precisely what I ought to do given *all* the facts, and that view differs from my view of precisely what I ought to do *in the sense that only takes into account a limited set of considerations*, I should care more about conforming my action to the former sort of obligation. At a first pass, acting against one’s beliefs in objective obligations is problematic even when one chooses to follow beliefs in rational obligations instead.

A fallibilist purist might attempt to resist this line of reasoning by pointing out that belief does not require certainty. The fact that Naomi has a *belief* about her objective obligations, on a fallibilist view, does not mean that she is *certain* about them. And indeed, the case as described stipulates that it is rational for Naomi *not* to be certain; she appropriately has credence 0.05 that she is not objectively obligated to go straight to the dinner party. Further, even if ‘deliberation aims at what’s best’, a rational agent in Naomi’s place would surely be sensitive to the possibility of error about what is best. The objector might lean on this insight to argue that, in cases like Naomi’s, there is nothing problematic or perverse about following one’s beliefs about rational obligations instead of one’s beliefs about objective obligations.

This objection gets a lot right. It is true that a rational agent is sensitive to the possibility of error, and it is also true that Naomi rationally ought to wait in line

precisely because of the risks attached to the possibility of error in her case. I also grant the fallibilist point that a person can believe that  $p$  without being certain that  $p$ . But it is far from clear that these points do anything to weaken the presumption against theories that would rationally require Naomi to wait in line *while believing she objectively ought not do so*.

To see why, compare Naomi's case to a paradigmatic case of akrasia. Suppose that Wayne believes he objectively ought to go to the gym, but instead stays at home and watches television. Even if we grant that Wayne holds his belief without certainty, we should accept that he is problematically akratic. We should accept, in other words, that beliefs about one's objective obligations are, even when held without certainty, the sort of mental state that can stand in relationships of problematic incoherence with action. (Some, like Dorst [2019: 200–1], hold that *certainty*, not belief, is the mental state that stands in these relationships of problematic incoherence. But purists who adopt this view would thereby abandon the datum that acting against one's believed obligations is, in general, akratic. This is no small cost.) The question at hand is how far this phenomenon generalizes. Is Naomi's belief like Wayne's, in that it generates coherence requirements on action? Or is it unlike Wayne's, in that it does not generate those coherence requirements? The fact that Naomi's belief is held without certainty does not favor one of these views over the other. (This, after all, is a point of similarity between Naomi's and Wayne's cases.) Importantly, the fact that it is rational for Naomi to be responsive to the risk of error also fails to favor either of these views over the other. It *does* favor the conclusion that Naomi rationally ought to wait in line at the post office. But it does not favor the view that she rationally ought to wait in line *while retaining her belief* (because, in her case, belief need not cohere with action) over the view that she rationally ought to wait in line *while suspending judgment* (because, in her case, belief should cohere with action). The considerations brought up by our fallibilist objector, in short, leave open two apparently consistent pictures of the extent to which beliefs about objective obligations rationally must cohere with action.

Which of these pictures is more attractive? The irrationality of akrasia is precisely the sort of consideration that seems apt to guide us in answering this question. The purist could *say*, of course, that coherence requirements involving beliefs about objective obligations vanish in high-stakes cases. But the fact that this claim is available and apparently consistent does not mean that it is attractive or well motivated. And it is an unattractive feature of the picture at hand that it requires agents like Naomi to act against their beliefs about what is best. All else being equal, we should prefer the view on which Naomi is not so required.

Some might suspect that this consideration only has force against the background of a picture on which belief plays a 'settling' role, perhaps by providing 'fixed points' in deliberation. (See especially Fantl and McGrath [2009: ch. 5]; Wedgwood [2012]; and Ross and Schroeder [2014].) I locate the burden of proof differently. Even for theorists who reject all precisifications of the claim that belief plays a 'settling' role, it should be uncontroversial that akrasia usually involves irrationality. These anti-'settling' views, then, owe us an alternative story about the source of the rational demand for coherence between actions and beliefs about one's own

objective obligations. I am happy to grant, for the purposes of this essay, that this story can be told successfully. The question is whether the theorist who takes up that task can explain why demands of coherence do arise in paradigmatic ‘low-stakes’ cases of akrasia, but not in certain ‘high-stakes’ cases of akrasia. What’s more, discharging this burden is no trivial task. Take a tempting proposal: the defender of nonsettling belief might claim that Wayne (but not Naomi) is problematically incoherent because Wayne (but not Naomi) fails to act in the way that he takes to maximize expected value. Proposals in this vein are non-starters; since a coherent agent can judge that he ought not maximize expected value, this cannot explain the distinctive badness of akrasia like Wayne’s.

I have argued that the irrationality of akrasia provides a significant mark in favor of impurist theories of rational belief, and against purist ones. All else being equal, we should prefer impurist views, precisely because they allow us to avoid positing rationally required akrasia in a wide range of cases without precedent. Of course, all else may not be equal. There may be independent reasons to accept a picture of rationality on which certain coherence requirements vanish in cases like Naomi’s. (Perhaps, for instance, this is the only way to retain a picture on which rational belief is as *stable* as we might pretheoretically like.) In other words, there may be considerations that will persuade some theorists to treat this essay’s *modus ponens* as a *modus tollens* and to accept the conclusion that rationally required akrasia is much more common than generally thought. But this is not to say that the purist has debunked the presumption against rationally required akrasia, or that she has shown why it should not be extended to cases like Naomi’s. It simply shows that there may well be considerations that will persuade purists to swallow an unattractive theoretical cost.

This objection is an important one, and my response has been in part concessive. Though I argue that the presumption against rationally required akrasia remains in force in cases like Naomi’s and thereby provides support for impurism, I grant that it could in principle be outweighed. But even for the purist who considers the cost of requiring akrasia to be bearable, it is crucial to see that it is indeed a cost. All purists should acknowledge that their view has striking implications about the extent of rationally required akrasia.

## 5. Upshots

What sort of impurism is supported by the argument I have offered? And in what respects does that argument make progress beyond other prominent arguments for impurism?

### 5.1 What Sort of Impurism?

Two features of the impurism that emerges from my argument are particularly important to note.

First, although my discussion has been primarily concerned with beliefs about objective obligations, there are good reasons to think that impurism extends to a wider range of beliefs. The question of whether Naomi objectively ought to drive

straight to the dinner party hangs entirely on the question of whether she is carrying Package A. Given this, it would be very odd for her to form the outright belief that she is carrying Package A while suspending judgment about whether she objectively ought to drive straight to the party. Since epistemic rationality requires us to make our beliefs coherent, then, impurism will likely spread beyond beliefs about our obligations. This means that even the rationality of beliefs about prosaic matters of fact (for example, *that I am carrying Package A*) will depend on non-truth-relevant facts.

Second, my argument supports a claim that some have called *moral encroachment*. (For defenses of moral encroachment, see Pace [2011]; Fritz [2017]; Moss [2018]; Basu and Schroeder [2019]; Bolinger [2020].) If there is moral encroachment in epistemology, then some paradigmatically epistemic facts (like whether a belief is knowledge) depend on non-truth-relevant *moral* considerations. To see why my argument provides a case in favor of moral encroachment, recall that the objective and rational obligations that I have been discussing are both species of *all-things-considered* obligations. Since moral considerations can make a difference to a person's all-things-considered obligations, some rational obligations (like Naomi's) are determined in part by moral considerations. Moreover, moral considerations can make a difference to what a person rationally ought to do even if they are of no importance *for the person*—that is, even if they have no bearing on her projects, goals, or well-being. From this conclusion, it is a short step to moral encroachment: since Naomi's rational obligations are shaped by moral considerations, and rational obligations constrain what we are epistemically rational to believe about our objective obligations, the rationality of belief is sensitive to non-truth-relevant *moral* considerations.

My argument's focus on all-things-considered obligations, then, leads to moral encroachment. And there are principled reasons for framing my argument in the way I have: only beliefs about all-things-considered obligations give rise to problematic akrasia. There need not be anything problematic about an agent who believes *I prudentially ought to  $\varphi$*  while failing to  $\varphi$ , or one who believes *I morally ought to  $\varphi$*  while failing to  $\varphi$ . Such an agent might simply fail to judge that prudence, or morality, is most important in her current choice situation. Understanding the sort of normative judgment involved in akrasia, then, helps to show that moral encroachment is just as well motivated as pragmatic encroachment more generally (see also Fritz 2017).

## 5.2 Other Arguments for Impurism

In the introduction to this essay, I draw on Roeber (2018) to sort existing arguments for impurism into intuition-based arguments, principle-based arguments, and function-based arguments. My argument has noteworthy advantages over these other forms of argument.

Unlike intuition-based arguments, my argument does not depend on much-contested intuitions about epistemic properties in high- and low-stakes cases. Indeed, I have simply used a high-stakes case to illustrate the



uncontroversial point that rational obligations depend on ethical facts about mere possibilities. Unlike function-based arguments, my argument does not rely on any sweeping claims about the function of epistemic evaluation or epistemic discourse.

My argument is importantly similar to at least one existing principle-based argument: that found in Fantl and McGrath (2007). This argument turns on the observation that, according to purism, certain agents must both (a) know which action maximizes *actual* utility, and (b) rationally must take a less risky action—one that, it might be tempting to think, maximizes *expected* utility. Fantl and McGrath claim that this is impossible: ‘if you know that A will have the highest *actual* utility of the options available to you, then A has the highest expected utility of the options available to you, assuming of course that what one is rational to do is the available act with the highest expected utility’ (2007: 568). Their argument, like mine, centers on the observation that purism rationally requires some agents to believe an apparently action-guiding proposition, while failing to be guided by that proposition.

There are at least two respects in which my line of argument goes beyond Fantl and McGrath’s—and does so in ways that may help clarify what is at stake in, and what follows from, this approach to impurism. First, as I argue above (section 5.1), it is important to prosecute this sort of argument with a focus on all-things-considered ‘ought’ judgments. Having taken up that focus, I (unlike Fantl and McGrath) am in a position to draw out the lesson that there is moral encroachment in epistemology. (This is a conclusion for which Fantl and McGrath later tentatively express sympathy, in the context of pursuing an entirely different line of argument for impurism [2009: 76n21].) Second, unlike Fantl and McGrath, I have drawn attention to the point that my argument against purism does not depend on the successful defense of an exceptionless principle. My argument depends, instead, on the nearly platitudinous claim that akrasia is a paradigmatic form of irrationality. Importantly, this is not to say that *all* instances of akrasia involve irrationality; indeed, section 2 was largely devoted to acknowledging possible exceptions to that general rule. My argument, then, may help clarify where the problem for purism lies: purists face problems concerning coherence between belief and action *even if we grant* that there are exceptions to any simple, general principle connecting the rationality of belief and action (and, therefore, *prima facie* problems for all principle-based arguments).

My argument also has advantages over a similar approach recently offered by Roeber (2020). Roeber argues against purism on the grounds that the purist cannot explain why, in certain paradigmatic ‘high-stakes’ cases, an agent should take a less risky option rather than pursuing known actual utility (2020: 2655–60). Unfortunately, the purist has a ready response to this explanatory challenge. A fallibilist purist can claim that we are often called upon to aim at maximizing expected utility, rather than known actual utility, precisely because knowledge of an action’s actual utility does not tell the full, nuanced story about my epistemic position with respect to that action’s actual utility. It does not follow from the fact that I know the actual utility of an action that my best guide to pursuit of actual utility is simply to take what I know for granted. (Roeber may take himself to have ruled out this approach by noting that the fallibilist cannot embrace a

Ramseyan view, on which ‘we’re always in effect “guessing” what the world is like’ [2020: 2657]. But this would be too quick; fallibilist purists can both accept that knowledge goes beyond merely guessing and also claim that, in pursuit of actual utility, I should sometimes consider possibilities of error.)

My argument, by contrast, does not rest on the putative difficulty of explaining how a sensible agent with knowledge could avoid running risks. (I grant, above in section 4.2, that fallibilist purists can explain in a satisfying way why Naomi rationally ought to wait in line.) It rests, instead, on the point that there is a theoretical cost associated with sanctioning an unprecedented form of rationally required akrasia. I have also argued (in section 4.2), that this cost does not go away when we note that there is an available and consistent picture on which high-stakes cases do in fact rationally require akrasia. All else equal, we should prefer a theory that does not require akrasia in high-stakes cases *at all* over one that says some things to explain why it is intelligible.

My argument from the irrationality of akrasia gives the impurist a unique and forceful way to make clear just why considerations of coherence between belief and action count against purism. Even those who cannot bring themselves to give up purism, then, stand to gain something important from this essay: they should acknowledge that their view has striking, underappreciated implications about the extent of rationally required akrasia.

JAMES FRITZ 

VIRGINIA COMMONWEALTH UNIVERSITY

[jamie.c.fritz@gmail.com](mailto:jamie.c.fritz@gmail.com)

## References

- Anderson, Charity, and John Hawthorne. (2019) ‘Knowledge, Practical Adequacy, and Stakes’. In Tamar Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 6 (Oxford: Oxford University Press), 234–57.
- Baker, Derek. (2018) ‘Skepticism about Ought *Simpliciter*’. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 13 (Oxford: Oxford University Press), 230–52.
- Basu, Rima, and Mark Schroeder. (2019) ‘Doxastic Wronging’. In Brian Kim and Matthew McGrath (eds.), *Pragmatic Encroachment in Epistemology* (New York: Routledge), 181–205.
- Bolinger, Renée Jorgensen. (2020). ‘The Rational Impermissibility of Accepting (Some) Racial Generalizations’. *Synthese* 197, 2415–31.
- Brown, Jessica. (2008) ‘Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning’. *Noûs*, 42, 167–89.
- Coates, Allen. (2012) ‘Rational Epistemic Akrasia’. *American Philosophical Quarterly*, 49, 113–24.
- Craig, Edward. (1999) *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- DeRose, Keith. (1992) ‘Contextualism and Knowledge Attributions’. *Philosophy and Phenomenological Research*, 52, 913–29.
- DeRose, Keith. (2009) *The Case for Contextualism*, vol. 1. Oxford: Oxford University Press.
- Dorst, Kevin. (2019) ‘Lockeans Maximize Expected Accuracy’. *Mind*, 128, 175–211.
- Fantl, Jeremy, and Matthew McGrath. (2007) ‘On Pragmatic Encroachment in Epistemology’. *Philosophy and Phenomenological Research*, 75, 558–89.
- Fantl, Jeremy, and Matthew McGrath. (2009) *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

- Foot, Philippa. (1997) 'Morality as a System of Hypothetical Imperatives'. In Stephen Darwall, Allan Gibbard, and Peter Railton (eds.), *Moral Discourse and Practice* (Oxford: Oxford University Press), 313–22.
- Fritz, James. (2017) 'Pragmatic Encroachment and Moral Encroachment'. *Pacific Philosophical Quarterly*, 98, 643–61.
- Gerken, Mikkel. (2015) 'The Roles of Knowledge Ascriptions in Epistemic Assessment'. *European Journal of Philosophy*, 23, 141–61.
- Grimm, Stephen. (2015) 'Knowledge, Practical Interests, and Rising Tides'. In David K. Henderson and John Greco (eds.), *Epistemic Evaluation: Purposeful Epistemology* (Oxford: Oxford University Press), 117–37.
- Hannon, Michael. (2017) 'Knowledge's Threshold Problem'. *Philosophical Studies*, 174, 607–29.
- Harman, Elizabeth. (2015) 'The Irrelevance of Moral Uncertainty'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 10 (Oxford: Oxford University Press), 53–79.
- Hawthorne, John, and Jason Stanley. (2008) 'Knowledge and Action'. *Journal of Philosophy*, 105, 571–90.
- Kelly, Thomas. (2002) 'The Rationality of Belief and Some Other Propositional Attitudes'. *Philosophical Studies*, 110, 163–96.
- Kolodny, Niko, and John MacFarlane. (2010) 'Ifs and Oughts'. *Journal of Philosophy*, 107, 115–43.
- Lasonen-Aarnio, Maria. (2014) 'Higher-Order Evidence and the Limits of Defeat'. *Philosophy and Phenomenological Research*, 88, 314–45.
- Lord, Errol. (2015) 'Acting for the Right Reasons, Abilities, and Obligation'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 10 (Oxford: Oxford University Press), 26–51.
- McGrath, Matthew. (2015) 'Two Purposes of Knowledge Attribution and the Contextualism Debate.' In David K. Henderson and John Greco (eds.), *Epistemic Evaluation: Purposeful Epistemology* (Oxford: Oxford University Press), 138–60.
- McPherson, Tristram. (2018) 'Authoritatively Normative Concepts'. In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 13 (Oxford: Oxford University Press), 253–77.
- Moss, Sarah. (2018) 'Moral Encroachment'. *Proceedings of the Aristotelian Society*, 118, 177–205.
- Pace, Michael. (2011) 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' "Will to Believe"'. *Noûs*, 45, 239–68.
- Parfit, Derek. (2011) *On What Matters*. Oxford: Oxford University Press.
- Reed, Baron. (2010) 'A Defense of Stable Invariantism'. *Noûs*, 44, 224–44.
- Rinard, Susanna. (2017) 'No Exception for Belief'. *Philosophy and Phenomenological Research*, 94, 121–43.
- Roeber, Blake. (2018) 'The Pragmatic Encroachment Debate'. *Noûs*, 52, 171–95.
- Roeber, Blake. (2020) 'How To Argue for Pragmatic Encroachment'. *Synthese*, 197, 2649–64.
- Rose, David, Edouard Machery, Stephen Stich, Mario Alai, Adriano Angelucci, Renatas Berniūnas, Emma E. Buchtel, et al. (2017) 'Nothing at Stake in Knowledge'. *Noûs* 53, 224–47.
- Ross, Jacob, and Mark Schroeder. (2014) 'Belief, Credence, and Pragmatic Encroachment'. *Philosophy and Phenomenological Research*, 88, 259–88.
- Smithies, Declan. (2019) *The Epistemic Role of Consciousness*. Oxford: Oxford University Press.
- Stanley, Jason. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Sylvan, Kurt. (2015) 'What Apparent Reasons Appear to Be'. *Philosophical Studies*, 172, 587–606.
- Thomson, Judith Jarvis. (2001) *Goodness and Advice*. Edited by Amy Gutmann. Princeton: Princeton University Press.
- Tiffany, Evan. (2007) 'Deflationary Normative Pluralism'. *Canadian Journal of Philosophy*, 37 (supplement), 231–62.
- Titelbaum, Michael G. (2015) 'Rationality's Fixed Point (Or: In Defense of Right Reason)'. In Tamar Szabó Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 5 (Oxford: Oxford University Press), 253–94.
- Weatherston, Brian. (2019) *Normative Externalism*. Oxford: Oxford University Press.
- Wedgwood, Ralph. (2012) 'Outright Belief.' *Dialectica*, 66, 309–29.
- Wodak, Daniel. (2019) 'An Objectivist's Guide to Subjective Reasons'. *Res Philosophica*, 96, 229–44.
- Worsnip, Alex. (2015) 'Two Kinds of Stakes'. *Pacific Philosophical Quarterly*, 96, 307–324.