

Does *Ought* Imply *Can* in Epistemology?*

Tim Smartt

timjsmartt@gmail.com

Forthcoming in the *American Philosophical Quarterly*

Abstract

Can epistemology require the impossible? Many philosophers believe the answer is ‘no’; if an epistemological theory generates a requirement that it would be impossible to satisfy, then, no matter how attractive the theory is, it is false. This view seems to depend on the idea that *ought* implies *can* in epistemology, at least in some sense. How should we understand this principle? In this paper, I introduce and explore a strong version of an epistemic *ought* implies *can* principle whereby an epistemic *ought* implies a psychological *can*. I outline a novel way to understand this principle—one which takes it in a contextualist spirit—and argue that it can meet two serious challenges that often lead philosophers to reject similar principles. The challenges are making the notion of ‘psychological ability’ sufficiently clear and avoiding an excessive permissiveness which diminishes the normative force of epistemic requirements. The upshot is that a strong epistemic *ought* implies *can* principle is more plausible than many philosophers think.

Keywords: Epistemic normativity · *ought* implies *can* · demandingness objections · ability · contextualism

Can epistemology require the impossible? Many philosophers believe the answer is ‘no’; if an epistemological theory generates a requirement that it would be impossible to satisfy, then, no matter how attractive the theory is, it is false.

*Earlier versions of this paper were presented at the CUNY Graduate Centre, King’s College London, Goethe University Frankfurt, the University of Vienna, the University of Sydney, and the AAP Annual Conference at the University of Wollongong. I’m grateful to these audiences for very helpful discussion. For comments and discussion that improved this paper, I would like to thank Jessica Brown, Mark Colyvan, Daniel Greco, Grace Helton, Jenann Ismael, Maria Lasonen-Aarnio, Andy Mueller, Philip Pettit, Nicholas J. J. Smith, and especially Brian Hedden.

Just like moral or political theories that require the impossible are sometimes resisted on the grounds that *ought* implies *can*—that is, that one can be required to do something in these domains only if one is able to do it—analogous objections are often raised against epistemic requirements. But how should we understand the principle at work in these objections? Some objections only require a weak principle, whereby an epistemic *ought* implies a sense of ‘can’ that denotes logical, metaphysical, or nomological possibility. Others require a much stronger principle.

In this paper, I introduce and explore a strong version of an epistemic *ought* implies *can* principle, whereby an epistemic *ought* implies a psychological *can*. This version of the principle is intuitively appealing; if epistemology cannot require the impossible, it is natural to think that is because it cannot require what is impossible for creatures like us. But stronger principles face two serious challenges that have led many epistemologists who accept that epistemology cannot require the impossible to settle for weaker principles. The first challenge is making the notion of ‘psychological ability’ sufficiently clear (Hedden 2017). The second is avoiding an excessive permissiveness that diminishes the normative force of epistemic requirements (Greco 2012).¹ My goal is to take up both challenges. I outline a version of an epistemic *ought* implies a psychological *can* principle and argue that it meets these challenges. However, although I am attracted to this principle in several respects, my overall verdict about it is cautious. I hope to show that a strong epistemic *ought* implies *can* principle is more plausible than many people think, while acknowledging that the defence I explore uncovers a couple of issues that still remain to be addressed.

This paper is in five parts. In [Section 1](#), I present examples of prominent demadingness objections in epistemology—objections that claim that an epistemological view is false because it requires the impossible. In [Section 2](#), I map several ways of interpreting an epistemic *ought* implies *can* principle that might vindicate objections such as these. In [Section 3](#), I demonstrate that a strong version of the principle is required to vindicate *all* the objections discussed in [Section 1](#), and I argue that we should prefer one of the stronger principles over the other; namely, an epistemic *ought* implies a psychological *can* principle over an epistemic *ought* implies a volitional *can* principle. In [Section 4](#), I outline and defend a novel version of the principle. In [Section 5](#), I take up two objections that one might have to the view I introduce, and I suggest that the view receives

¹Feldman (1988, 2000, 2008), Ryan (2003, 2013), Lycan (1985), and Mizrahi (2012) have also rejected strong epistemic *ought* implies *can* principles on the basis of these challenges.

a mixed scorecard: one of these objections can be overcome, and one seems much more vexing. I end by noting that this reveals an interesting choice point for how we theorise about demandingness in epistemology.

1 Demandingness objections in epistemology

Objections that claim that a view is false because it requires the impossible are raised against all sorts of epistemic requirements. To appreciate how widespread these kinds of demandingness objections are, let's consider five examples from debates about Bayesianism, the norm of belief, epistemic dilemmas, time-slice rationality, and justification.

First, some object that Bayesianism demands the impossible. Take this humorous remark from Richard Holton:

However much we might want consciously to entertain sets of credences that conform to the axioms of the probability calculus, and manipulate them by methods like conditionalization, it looks as though we simply cannot do it, any more than wanting to swim like a dolphin enables us to do that. (Holton 2008, 34).

Holton claims that Bayesianism requires us to have credences and update them in technical ways. He thinks that we simply cannot do this, so Bayesianism is false, at least when understood as a normative epistemological theory. Another common objection to Bayesianism and other probabilistic theories of rationality is that they require logical omniscience, whereby the rational credence one ought to have in any logical truth is always 1. For example, Ralph Wedgwood (2017, 37) writes, 'We might well wonder whether such logical infallibility is even possible for normal human thinkers like you and me.'

Second, demandingness objections have been raised against different accounts of the norm of belief. According to the knowledge norm of belief, one should believe only what one knows (Williamson 2000). According to a version of the truth norm of belief, one should believe all and only the truths.² Both norms have been objected to on the grounds that—to stick with Holton's phrase—they require what we simply cannot do. Regarding the knowledge

²Different versions of a truth norm of belief have been defended by a number of philosophers. See, for example: Shah (2003), Velleman and Shah (2005), Gibbard (2000, 2005), and Wedgwood (2002). For a defence of the view that the truth norm of belief requires one to believe all the truths, see Greenberg (2020).

norm, you might think that we can fail to know p , without being in a position to know that we fail to know p .³ In this case, the knowledge norm says that we should not believe p , since we do not know p , even though we are not in a position to know that we fail to know p . When we are not in a position to know what we do not know, this norm will require of us what we simply cannot do.⁴ Regarding the truth norm, it seems like there are far too many true propositions to be believed, that some true propositions will be far too complex to be believed, and that some true propositions would be too costly to try to learn.⁵

Third, demandingness objections have been raised against epistemic dilemmas. Some philosophers defend the view that there are numerous epistemic requirements that cannot always be jointly satisfied.⁶ Just like we can find ourselves in moral dilemmas where it is impossible to perform an action that satisfies all the relevant requirements, so we can find ourselves in epistemic dilemmas where no matter what one believes, one will end up violating an epistemic requirement. Susanna Rinard objects to the very idea of epistemic dilemmas on the grounds that they make epistemic requirements too demanding:

It is central to our concept of rationality that rationality constitute an ideal to which one could coherently aspire, and by which one could be guided. But if doxastic dilemmas were possible, rationality could not play this role in those cases. We cannot be guided by the voice of rationality if rationality tells us to neither believe, nor disbelieve, nor suspend judgment on p . (Rinard 2018, 259).

In other words, Rinard rejects epistemic dilemmas because she thinks that we cannot be rationally required to do what we simply cannot do.

Fourth, demandingness objections have been raised against time-slice rationality.⁷ According to time-slicers, all fundamental epistemic requirements are synchronic; they target how one's cognitive attitudes ought to be at a par-

³Philosophers who accept the knowledge norm of belief tend to also accept that knowledge is not 'luminous'—that is, roughly, that being in the state of knowing p does not guarantee that one is also in the state of knowing that one knows p . See Williamson (2000) Chapter 4.

⁴For an argument to this effect, see Mueller and Ross (2017).

⁵This objection has been forcefully pressed by Anandi Hattiangadi and Krister Bykvist. See Bykvist and Hattiangadi (2007, 2013) and Hattiangadi (2010, 2019).

⁶For a general defence of epistemic dilemmas, see Hughes (2019). Other accounts of the possibility of genuine epistemic dilemmas include Pryor (2018), Christensen (2010a, 2010b, 2016), Ross (2010) and Williamson (forthcoming).

⁷Time-slice rationality has been most prominently defended by Hedden (2015a, 2015b) and Moss (2015a, 2015b).

ticular time rather than how they ought to behave over time. One objection to this view is that it is much too demanding—it requires that we update instantaneously, something we simply cannot do since our cognitive processes take time.⁸

Fifth, demandingness objections feature in work on justification. William Alston (1988) argued that the deontological conception of epistemic justification—roughly, the view that we ought to adopt justified beliefs and jettison unjustified beliefs—assumes that we have control over our beliefs. Alston argued that it is implausible that we have any sort of meaningful control over our beliefs, so deontological conceptions of epistemic justification are not apt.

All these objections reject an epistemic requirement on the grounds that it is excessively demanding. For these objections to be legitimate, it must be the case that *ought* implies *can* in epistemology, at least in some sense.⁹

2 Five epistemic *ought* implies *can* principles

Although demandingness objections are common in epistemology, it is not clear how we should understand the principle they depend on.¹⁰ In this section, I set out five ways to interpret an epistemic *ought* implies *can* principle based on more or less stringent modalities associated with different senses of

⁸This objection has been pressed in Podorski (2016, 2017), Snedgar (2017), and Lenman (2017).

⁹A couple of caveats: First, there may be *other* reasons to object to an epistemic requirement targeted by a demandingness objection apart from the fact that it asks of us what we simply cannot do. You might even think that these other reasons are often more important. But I take it that if demandingness does provide an appropriate objection to an epistemic requirement, then *ought* implies *can* in epistemology, in some sense. Second, it is worth mentioning that there are other kinds of demandingness objections that do not depend on an *ought* implies *can* principle. To take an example from ethics, Bernard Williams' (1973) integrity objection to act utilitarianism is a kind of demandingness objection. Williams objects that act utilitarianism is excessively demanding, not because it requires of agents what they simply cannot do, but because it requires agents to set aside personal projects and act in ways that conflict with their deepest values and commitments; a situation Williams thought amounts to an erosion of their personal integrity. Williams' objection does not rely on an *ought* implies *can* principle. Were it possible for an agent to satisfy all the requirements of act utilitarianism, Williams would *still* hold that act utilitarianism demands too much. So although demandingness objections *in general* do not depend on an *ought* implies *can* principle, many of them do, including those that claim that an agent cannot be required to do what they simply cannot do. All of the demandingness objections in epistemology that I am aware of are of this latter kind.

¹⁰Some philosophers have directly discussed an epistemic *ought* implies *can* principle, offering accounts of how the principle might be understood and either defending or rejecting it. This includes Greco (2012), Helton (2018), Wedgwood (2013), Neta (2014), Mizrahi (2012), Smithies (2015), Christensen (2004), and Chuard and Southwood (2009). However, the majority of writers assert or reject demandingness objections without discussing how they understand the principle.

‘can.’¹¹ Here are the principles, based on a logical, metaphysical, nomological, psychological, and volitional sense of ‘can.’

(1) Epistemic Ought Implies Can_{LOG}

If S ought to believe p , then it is logically possible for S to believe p .

(2) Epistemic Ought Implies Can_{MET}

If S ought to believe p , then it is metaphysically possible for S to believe p .

(3) Epistemic Ought Implies Can_{NOM}

If S ought to believe p , then it is nomologically possible for S to believe p .

(4) Epistemic Ought Implies Can_{PSY}

If S ought to believe p , then it is psychologically possible for S to believe p .

(5) Epistemic Ought Implies Can_{VOL}

If S ought to believe p , then S has control over whether they believe p .

The principles naturally divide into a weaker class—those with a logical, metaphysical, and nomological sense of ‘can’—and a stronger class—those with a psychological and volitional sense of ‘can.’ I will briefly discuss how we might understand each principle, starting with the weaker principles.

Let’s start with the weakest version of the principle. On a logical sense of ‘can’, an epistemic *ought* implies *can* principle holds that epistemic requirements are constrained by logical possibility. It cannot be the case that an agent

¹¹All three notions in this principle have received plenty of philosophical attention: *ought*, *implies*, and *can*. I focus on ‘can’ since it does most of the work in determining the strength of the principle. Other philosophers have discussed which notions of ‘ought’ and ‘implies’ make for the most plausible *ought* implies *can* principle. For discussion of what sort of ‘oughts’ imply ‘can’ see Southwood (2016). For discussion of the nature of the implication relation between ‘ought’ and ‘can’ see Sinott-Armstrong (1984).

ought to be in an arrangement of doxastic states that violates the rules of a given logic.

On a metaphysical sense of ‘can’, an epistemic *ought* implies *can* principle holds that epistemic requirements are constrained by metaphysical possibility. Exactly what it takes for a belief to be metaphysically possible will depend on one’s metaphysics of belief. Daniel Greco (2012) accepts an epistemic *ought* implies *can* principle with this sense of ‘can’, arguing that it is not the case that one ought to be a near-global sceptic because it is metaphysically impossible—on a number of different metaphysics of belief—to be agnostic between all hypotheses that are consistent with how things seem to be.¹²

On a nomological sense of ‘can’, an epistemic *ought* implies *can* principle holds that epistemic requirements are constrained by nomological possibility. It cannot be the case that an agent ought to be in a doxastic state that it is impossible to be in given the actual laws of nature.

The next two senses of ‘can’ are stronger; both hold fixed a wider range of facts about specific agents. These two senses are sometimes run together, so I would like to take the time to clearly separate them.¹³

On a psychological sense of ‘can’, an epistemic *ought* implies *can* principle holds that epistemic requirements are constrained by some set of facts about the psychological abilities and limitations of actual agents. Actual agents have all kinds of psychological abilities—they can recognize faces, crunch numbers, use languages, and so on. And they have all kinds of psychological limitations—they can only discriminate so much detail, only memorize so much information, only direct their attention towards so many stimuli, and so on. On this sense of ‘can,’ psychological facts place limits on what is epistemically required of an agent.

A few philosophers have endorsed an epistemic *ought* implies *can* principle with this sense of ‘can.’ However, they disagree on what psychological facts are salient. Grace Helton (2018) endorses an epistemic *ought* implies *can* principle with quite a strong reading of the psychological ‘can’—she argues that epistemic requirements are constrained by a specific agent’s repertoire of ‘psychological skills and mechanisms.’ For example, someone who is fluent in English but not Mandarin cannot be required to revise their inaccurate belief

¹²It is also accepted in Hedden (2017).

¹³For instance, Greco (2012), Mizrahi (2012), Feldman (1988, 2000, 2008) and Ryan (2003, 2013) criticise and reject an epistemic *ought* implies *can* principle where the modality of ‘can’ is either psychological or volitional. But it is not clear to me which sense they have in mind.

about a text just by looking at its Mandarin translation. Nor can a tone-deaf agent be rationally required to revise their belief about which key a particular song is in. Whereas someone whose current repertoire of psychological skills and mechanisms includes proficiency in Mandarin, or perfect pitch, might be required to.¹⁴ Fred Dretske (2000), Ram Neta (2014), Mona Simion (2024), and Alvin Goldman (1978) endorse a weaker epistemic *ought* implies psychologically *can* principle whereby epistemic requirements are constrained by general facts about normal human psychological abilities.¹⁵ Ralph Wedgwood (2013) endorses a version of the principle whereby epistemic requirements are constrained by just *one* psychological mechanism; namely, reasoning. On this view, if it is the case that *S* ought to believe *p*, then *S* is able to reason in such a way that they believe *p*.

Last, on a volitional sense of ‘can’, an epistemic *ought* implies *can* principle holds that if *S* ought to believe *p*, then *S* has control over whether they believe *p*. On this view, if there is nothing one can do to get oneself to believe *p*, then one cannot be required to believe *p*. Alston seemed to accept this principle. He understands control on the model of intentional action—whereby one has control over whether one believes *p* only if one can form and execute an intention to believe *p*—and argued that we do not have any meaningful control over what we believe.¹⁶

Is there really a difference between the psychological and volitional senses of ‘can’? I think there is, so I would like to clarify how I understand it. I agree with a number of writers (Helton 2018, Chuard and Southwood 2009, and McHugh 2012) who argue that it is a mistake to think that if an agent can ϕ then they have control over whether they ϕ . Agents can do plenty of things that

¹⁴Helton acknowledges that one might accept a competing explanation of each example whereby the salient epistemic difference between the agent who can read Mandarin and the agent who cannot is that they are in possession of different evidence. However, she introduces a third case that she claims breaks the tie between evidence-based and ability-based explanations of the differences. She writes: ‘Suppose a subject is in the middle of a long mathematical proof. She has, in some sense, all of the evidence she needs to reach some very difficult conclusion. However, she is working at the absolute limit of her cognitive processing capacity and cannot, psychologically speaking, draw the relevant conclusion. It seems to me that this subject is not rationally obligated to draw this conclusion, even though she possesses the relevant evidence. The claim that rational obligation is constrained by psychological ability can explain this result, whereas evidence-based considerations cannot.’ (Helton 2018, 15).

¹⁵For instance, Goldman (1978, 510) claims that epistemic requirements ‘must take account of the powers and limits of the human cognitive system.’ Neta (2014, 52) claims that epistemic requirements must be sensitive to the normal capacities of rational agents. And Simion (2024, 208) claims that an agent’s epistemic obligations ‘imply capacities in the kind of cognizer that she is—e.g. cognitive capacities that adult cognizers have.’

¹⁶Or, as he puts it, control requires that one has it in their power to carry out an intentional action to take up a certain propositional attitude. (Alston 1988, 277).

they do not have control over. They can sneeze and can hiccup, they can fall asleep and can dream (McHugh 2012, 85, 87). They can distinguish cows from sheep, and empathise with others (Chuard and Southwood 2009, 618). They can discriminate between pitches and understand language (Helton 2018, 12–13). A number of these are psychological abilities associated with cognition, cognitive architecture, mental processes, and so on. The psychological ‘can’ distinguishes these abilities from the ability to execute an intention. The distinction also holds in the opposite direction. An agent might have control over whether they ϕ even though they lack the psychological ability to ϕ . For example, an agent might have control over whether they perform a quick act of senseless killing—in the sense that were they to intend to perform the act, then they would do it—despite being psychologically incapable of forming an intention to do so. I emphasize all this since the distinction between the psychological and volitional senses of ‘can’ is relevant when considering which principle can plausibly motivate a range of demandingness objections.

3 Demandingness and *ought* implies *can*

We now have five ways to understand an epistemic *ought* implies *can* principle. Do any allow us to raise all the demandingness objections in [Section 1](#)? It would be nice if one did. At their core, demandingness objections hold that epistemology cannot require of us what we simply cannot do. If these objections could be motivated by one principle, then we would be in much better position to evaluate the whole class of objections. For instance, if a version of the principle were able to play this role *and* it could be defended, then the class of objections would be vindicated. We would also be in a better position to understand what constitutes a legitimate demandingness objection. I think this is a promising strategy for epistemologists sympathetic to demandingness objections. I will call this the unity desideratum.

The weaker principles—those with a logical, metaphysical, and nomological sense of ‘can’—are not strong enough to do so. A number of requirements targeted by demandingness objections ask of agents what is logically, metaphysically, and nomologically possible. For instance, I take it that it is logically, metaphysically, and nomologically possible (at least on some metaphysics of belief) to be logically omniscient, but that is just the kind of requirement that is often targeted by demandingness objections. Likewise, I take it that it is log-

ically, metaphysically, and nomologically possible to have precise credences and update them by conditionalization, to only believe what one knows, and to adopt justified beliefs and jettison unjustified beliefs.

It is worth emphasizing that my claim is *not* that weaker principles cannot motivate token demandingness objections. I think they can; these principles are not toothless. Each one motivates at least one demandingness objection from [Section 1](#). To rule out dilemmic requirements as Rinard wants to, all that is required is an epistemic *ought* implies logically *can* principle.¹⁷ To rule out the truth norm of belief, all that is required is an epistemic *ought* implies metaphysically *can* principle, since it is plausible that due to the nature of *de se* and *de re* content it is metaphysically impossible for any agent to believe all the truths.¹⁸ And to rule out time-slice rationality, all that is required is an epistemic *ought* implies nomologically *can* principle—causation takes time, so it is not possible, in worlds like ours, for an agent to update instantaneously. But none of these are strong enough to motivate all the objections we have considered.

The principles with a psychological and volitional sense of ‘can’ fare better in this regard. Both are strong enough to motivate each demandingness objection from [Section 1](#). One could plausibly object that each requirement—that is, Bayesian requirements, requirements to believe only what one knows, requirements to believe all the truths, dilemmic requirements, synchronic requirements, and requirements to adopt justified beliefs and jettison unjustified beliefs—asks of an agent what they simply cannot do in light of their psychological abilities. Likewise, one could plausibly object that each requirement asks of an agent what they simply cannot do in light of their powers of control.

Is there any reason to prefer either principle? I think there is, since the volitional principle forces one to also accept doxastic voluntarism in order to avoid an implausible minimalism according to which *all* epistemic requirements are excessively demanding. If the volitional principle is *true*—that is, epistemic requirements cannot ask of us what is beyond our powers of doxastic control—and doxastic voluntarism is *false*—that is, we have no substantial powers of doxastic control—then *all* epistemic requirements are excessively demanding. To avoid this result, one must accept doxastic voluntarism. Voluntarism is, of course, a controversial thesis, many find it implausible that we have any mean-

¹⁷Although, for an argument that dilemmic requirements are not in fact logically contradictory, see Hughes (2019), especially §4.

¹⁸For an argument to this effect, see Grim (1985).

ingful powers of doxastic control. So, the volitional version of the principle comes with a costly commitment.

4 Epistemic *ought* implies psychologically *can*

Let's take stock. In [Section 1](#), I reviewed a number of prominent demandingness objections in epistemology, all of which depend on some sort of epistemic *ought* implies *can* principle. In [Section 2](#), I mapped five ways to interpret this principle. In [Section 3](#), I argued that the psychological version of the principle is best suited to the role of motivating all the objections discussed in [Section 1](#). I also suggested that those sympathetic to demandingness objections have good reason to seek out *one* principle that can motivate the class of objections.

One might worry that even if this version of the principle is fit for purpose, it nevertheless has other problems which means, all things considered, it is not plausible. In this section, I argue that the psychological version of the principle can be understood in a contextualist spirit, which avoids two serious challenges. I also argue that this version of the principle is independently motivated from theoretical considerations about the semantics of ability ascriptions. The upshot of this section is that a strong epistemic *ought* implies *can* principle can be more thoroughly motivated than it has been thus far.

4.1 Two problems

A number of epistemologists, including those sympathetic to the view that epistemology cannot require the impossible, reject an epistemic *ought* implies psychologically *can* principle due to two challenges. I will call these the permissiveness problem and the specification problem.

First, here is the permissiveness problem. You might worry that although this principle provides a nice way to object to excessively demanding epistemic requirements, it does so at the cost of making normative epistemology excessively permissive. Surely epistemic requirements are demanding in some sense; *ought* does not imply *actual*. There are many things that it might be psychologically impossible for an agent to believe that we may want to insist they ought to believe. But an epistemic *ought* implies psychologically *can* principle says that if an agent cannot revise a belief, then they are not required to. This lets people off the hook much too easily. This objection is pressed by Greco (2012), and he presents a case that nicely captures the problem:

A hopeless paranoid schizophrenic may be unable to give up the belief that he is the target of a Martian conspiracy. That he is psychologically incapable of giving up this belief does not mean that he's not irrational for holding it; in the relevant sense of "ought," it is still the case that he ought to give it up. Being unable to appreciate the force of the evidence that Martians neither exist nor have infiltrated Earthly governments does not change the force of that evidence. (Greco 2012, 350).¹⁹

Second, here is the specification problem. The claim that epistemic requirements are constrained by psychological ability is unclear (Hedden 2017, 616-617). Which psychological abilities determine whether an agent is capable of satisfying an epistemic requirement? Advocates of the principle fill in the notion of 'psychological ability' in different ways. Helton (2018) holds that the salient abilities are the psychological skills and mechanisms of a specific agent, Wedgwood (2013) holds that just *one* mechanism is salient—namely, an agent's powers of reasoning—and Dretske (2000), Neta (2014), Simion (2024), and Goldman (1978) hold that the salient abilities are the general psychological abilities that humans tend to have. There are other ways to clarify this idea. You might hold that the salient psychological abilities are those that an agent has at a specific time, or tends to have across time (regardless of whether they are temporarily diminished or disabled), or that tend to be found in members of a specific reference class, or that are paradigmatically associated with cognition.²⁰ The specification problem is that without a detailed account of how this part of the principle is to be filled in, the principle is unclear.

¹⁹Similar cases are presented in Mizrahi (2012), Ryan (2003), and Feldman (1988), as counterexamples to an epistemic *ought* implies *can* principle in general. The cases all involve agents who cannot get themselves to believe that a close friend or family member has done something terrible (for instance, that their spouse committed a murder) despite possessing strong evidence. As I mentioned earlier (fn. 13), it is not clear to me whether each of these cases is intended to target the psychological version of the principle, the volitional version, or both. In any case, these cases do raise the permissiveness problem for the psychological version of the principle.

²⁰Further possibilities can be found in the literature on the metaphysics of ability. For example, Kenny (1975, 131) claims that 'can' is often best understood as signifying having an ability *and* having the opportunity to exercise it. And many writers (e.g., Mele 2002, Fara 2008, Maier 2015) emphasise the distinction between general and specific abilities. These kinds of distinctions identify additional ways that one might fill in the notion of 'psychological ability' more concretely.

4.2 Psychological ability in context

One possible solution to both problems is to understand the principle in a contextualist spirit. In a nutshell, the context of epistemic evaluation supplies the salient psychological facts in view of which it is either possible or impossible for an agent to satisfy a requirement.

An epistemic *ought* implies psychologically *can* principle, where one is contextualist about the salient psychological facts, holds that it is true that epistemic requirements are constrained by facts about an agent's psychological ability, but these facts are not held fixed across all contexts. Context determines which psychological facts determine whether an agent can satisfy a requirement. This way of understanding the principle differs from existing accounts which all hold that, across all contexts, a fixed notion of psychological ability determines whether an agent can satisfy a requirement.

On this view, sometimes in one context *S* ought to believe *p*, and in a different context it is not the case that *S* ought to believe *p*. This is because the different contexts of epistemic evaluation make different psychological abilities salient in each case, which generates a pair of epistemic *ought* implies psychologically *can* principles with different contents relative to each case. For instance, in contexts where the salient facts are about general psychological abilities that tend to be distributed across the whole species, we get the principle:

(4)' Epistemic Ought Implies Can_{PSY}

If *S* ought to believe *p*, then it is possible *in view of general facts about human psychology* for *S* to believe *p*.

Whereas, in contexts where the salient facts are about the psychological abilities that tend to be possessed by agents who are members of a specific reference class, we get the principle:

(4)'' Epistemic Ought Implies Can_{PSY}

If *S* ought to believe *p*, then it is possible *in view of facts about the psychological abilities of the members of reference class R (of which S is a member)* for *S* to believe *p*.

These principles allow that, sometimes, it can be true that in one context *S* ought to believe *p* and in another it is not the case that *S* ought to believe *p*.

I can spell out this account further with a couple of pairs of cases.²¹ First, take Greco's character: an agent with a psychological compulsion arising from a serious mental illness who cannot give up an irrationally held belief. Consider two specifications of the case that feature different contexts of evaluation: a *clinical* context and a *lab* context. In a *clinical* context—say, while determining care plans and treatment options for this particular person—the context of evaluation makes the salient facts about this agent's psychological abilities at a time. In this case, we get the principle:

(4)* Epistemic Ought Implies Can_{PSY}

If *S* ought to believe *p* at *t*₁, then it is possible *in view of facts about S's psychological abilities at t*₁ for *S* to believe *p* at *t*₁.

Other contexts might make the salient facts about the psychological abilities of people with the same diagnosis. For example, consider a *lab* context, where researchers are working on better understanding the cognition of people with paranoid schizophrenia. In this case, we get the principle:

(4) Epistemic Ought Implies Can_{PSY}**

If *S* ought to believe *p*, then it is possible *in view of facts about the psychological abilities of people with paranoid schizophrenia* for *S* to believe *p*.

(4)* rules out the claim that Greco's character ought to revise his belief, but (4)** does not. If we fix the facts that determine what is psychologically possible for this agent as facts about his individual psychological abilities at a time, then we get a principle like (4)*, which generates the verdict that the character cannot revise his belief. However, if we fix the facts that determine what is psychologically possible for this agent as facts about the psychological abilities of agents in a reference class of which he is a member—in particular, people with the same diagnosis—then we get a principle like (4)**. This principle generates the verdict that the character *can* revise his belief, since although people with the same diagnosis do struggle with delusions, compulsions, and similar afflictions, it is possible (and indeed, actual) that they revise an affected belief, although it may be very difficult. The pair of *ought* implies *can* principles provide the result that in some contexts the character *can* satisfy the requirement,

²¹This follows a standard strategy for providing a contextualist account of an epistemological term. For example, DeRose (1992) and Cohen (1999).

so it remains the case that he ought to, and in other contexts he *cannot* satisfy the requirement, so it is not the case that he ought to.

Second, consider the following case, which does not involve pathology.

JET LAG. Amy has just flown from Sydney to New York. She is highly numerate, and ordinarily has no problem accurately performing mental arithmetic and similar cognitive tasks. At breakfast the next morning, she knows that she should leave a 15% tip. She has not yet converted any money from Australian to American dollars. To transfer the correct funds, she needs to calculate both 15% of her bill and the amount that this comes to in Australian dollars. She knows the bill's total and the relevant conversion rate. But after a long-haul flight, hours spent shuffling through airport terminals, and a broken night's sleep, she simply cannot crunch the numbers and arrive at a true belief about the amount to transfer.

On the context of evaluation that is naturally suggested by the case—say, evaluating Amy as you sit across the breakfast table from her—the salient facts are about the agent's psychological abilities at a time, and so we get the principle (4)* outlined above. But consider a different context for evaluating Amy's performance in this case. Suppose Amy is being interviewed for a job as a high school maths teacher and this incident is brought before the attention of the panel. In this context of evaluation the salient facts are about the agent's psychological abilities across time, regardless of whether they are temporarily diminished or disabled. That is, this context makes salient facts about Amy's general numeracy. In this context, we get the principle:

(4)[†] **Epistemic Ought Implies Can_{PSY}**

If *S* ought to believe *p* at *t*₁, then it is possible *in view of facts about S's psychological abilities (regardless of whether S has these abilities at t₁)* for *S* to believe *p* at *t*₁.

(4)* rules out the claim that Amy ought to form a true belief about the tip, but (4)[†] does not. If we fix the facts that determine what is psychologically possible for Amy as facts about her psychological abilities at breakfast the morning after a long-haul flight, then we get a principle like (4)*, which generates the verdict that she cannot form a true belief about the tip. However, if we fix the facts that determine what is psychologically possible for Amy as facts about

her psychological abilities across time, then we get a principle like (4)[†]. Since Amy is capable of forming a true belief about the tip under normal circumstances, (4)[†] generates the verdict that she *can* believe as she ought to. Again, the pair of principles provide the result that in some contexts Amy *can* believe as she ought to, so it remains the case that she ought to, and in other contexts she *cannot*, so it is not the case that she ought to.

It is worth emphasising that this view is supported by a couple of semantic features of ‘can’ and other modal auxiliary verbs. First, it is widely accepted that one might fix the sense of a term like ‘can’ and maintain that token instances of the term are sensitive to context.²² So it is natural that a version of the principle might involve just one of the five senses of ‘can’ discussed in [Section 2](#) and that the content of this sense might vary across contexts. Second, Angelika Kratzer (1977, 1981, 1991) convincingly argues—and most philosophers take it as semantic orthodoxy—that ability attributions are context-sensitive. On this picture, although ‘can’ (and similar terms like ‘must’ and ‘ought’) has an invariant core logical structure across contexts, an essential part of its meaning is contributed by the context in which the attribution is made. In particular, Kratzer holds that the meaning of ‘can’ is sensitive to two parameters: a set of circumstances that is taken as given, and a way of ranking various aspects of those circumstances, or, in her terms, a ‘modal base’ and an ‘ordering source.’ In everyday uses of ‘can,’ speakers do not articulate exactly those circumstance that they are taking for granted or how they are to be ranked. But Kratzer holds that context supplies these details, and that both parameters are essential for determining the meaning of any token use of ‘can.’ The upshot is that the view I have been exploring fits naturally with prominent views about the semantics of ‘can.’

A contextualist construal of an epistemic *ought* implies a psychological *can* principle can respond to both the specification and permissiveness problem. The specification problem is avoided since the principle includes an account of how ‘psychological ability’ is to be understood: the salient psychological

²²For instance, Jason Stanley (2005, 53) writes: ‘Modal expressions are context-sensitive. In the first instance, there are what some philosophers have thought of as different ‘senses’ of possibility, such as physical possibility, logical possibility, epistemic possibility, and metaphysical possibility. But, even fixing upon one sense of modality, there are different readings of a modal term such as ‘could’, depending upon the context of use. For example, where ‘could’ is interpreted as physical possibility, one might mean physical possibility in a more or less restricted sense.’ And Hawthorne and Srinivasan (2013, 10), when discussing senses of ‘ought,’ write: ‘Even within any one of these broad categories, there is considerable scope for context-dependence. For example, the truth conditions of a deontic ought-claim will also be sensitive to which facts are held fixed in the conversational context.’

abilities are fixed by the context of epistemic evaluation and they vary across contexts. Is it plausible that different contexts of epistemic evaluation make salient different facts about agents' psychological abilities? I think it is. Many epistemologists have observed that we should not treat all contexts of epistemic evaluation the same way. Sometimes a feature of a context makes a significant difference regarding how to understand an epistemological term. Some features that have received attention include the practical stakes associated with a context, the moral significance of a context, and the kind of alternative possibilities a context makes relevant. Since these accounts emphasise the non-epistemological nature of these features, it is natural to expect a variety of features to have epistemological significance. My account holds that contexts of epistemic evaluation can contain information that selects different facts about psychological ability as a way of expanding or contracting what is psychologically possible for an agent. Some contexts will fix the relevant kind of ability as what is normal across the species, others as what is normal within a subset of the species of which the agent is a member, and so on.²³ The result is that the specification problem is avoided.

The permissiveness problem is not outright avoided, but it does lose much of its force. A contextualist version of an epistemic *ought* implies psychologically *can* principle does not *generally* let agents off the hook too easily. There will be contexts where the principle applies *and* an agent still ought to satisfy the requirement. It is possible that, in view of how the salient psychological abilities are specified in a particular context, the agent can satisfy the requirement. In some contexts, Greco's character is still required to give up his belief, and in some contexts, Amy is still required to form a true belief about the amount of her tip. The permissiveness problem holds that an epistemic *ought* implies psychologically *can* principle delivers the verdict that if an agent cannot revise a belief then they are not required to. On a contextualist understanding of the principle, there are legitimate applications of the principle that do not lead to this verdict.

However, this response only goes so far. Although a contextualist version of the principle does not face the permissiveness problem in general, it does face particular instances of it. For example, in those contexts where the salient facts are about an individual's current psychological abilities, we get the re-

²³It is worth flagging that in some contexts, the principle will not impose very severe restrictions, namely where certain ideal agents (or other kinds of agents that are cognitively superior to humans) are salient.

sult that if an agent cannot satisfy a requirement then it is not the case that they ought to. Consider Greco's paranoid schizophrenic character. When the salient facts are about his current psychological abilities, a contextualist epistemic *ought* implies psychologically *can* principle provides the verdict that it is not the case that he ought to revise his belief in a Martian conspiracy. This will strike many as unacceptably permissive; even though he cannot, in view of his current psychological abilities, give up this belief, it still seems that he really ought to.

In cases like this, I think the account must bite a bullet and accept this result. However, this is a defensible position in view of some broader consideration about epistemological evaluation. As Alston (1995) argues, there are numerous dimensions of epistemic evaluation—we can still say that Greco's character's beliefs are not supported by the evidence, are not rational, and that he is epistemologically worse off than if he revised his belief.

It might seem unnatural to separate these evaluative notions from what an agent is required to do, but it is common to hold that sometimes it is inappropriate to evaluate an agent along a particular axis. For instance, a number of existing views about epistemic evaluation tease out different ways we might criticise an agent and hold that only some of these are appropriate. Some (Armendt 1993, Kaplan 1996) argue that we may only criticise an agent's doxastic states, not the agent themselves. Others (Chuard and Southwood 2009, 621-623) hold that we may criticise an agent for failing to satisfy an epistemic requirement, but we may not blame them. Others (Hughes 2019, Southwood 2016, and Srinivasan 2015) hold that we may blame an agent for failing to satisfy some epistemic requirements, but we may only criticise them for failing to satisfy others.

I will not defend any of these views here; I just want to emphasise that sometimes one kind of epistemological evaluation is appropriate and another isn't. There is a more specific lesson here, too, in that each distinction illustrates my claim that although in some contexts it is not the case that Greco's character ought to revise his belief, we are still left with meaningful ways to evaluate the case. On the first view, we may hold that his *state* is rationally criticisable. On the second and third views, we may hold that the agent *themselves* is criticisable, although he is probably not blameworthy. So, although a contextualist version of an epistemic *ought* implies a psychological *can* principle will sometimes lead to the verdict that an agent is not required to satisfy a requirement because they are unable to, this does not imply that this is the final word of

evaluation.²⁴

5 Scorecard

So far, I have introduced a new view and emphasised several positive things that can be said in its defence. However, I think that the scorecard for this view is ultimately more mixed. In this section, I discuss two challenges that might be raised against the arguments of [Section 4](#). Both challenges pick out issues associated with the contextualist framework surrounding the principle. The first challenge can be met, but I think the second is more vexing. But it is also illuminating: it reveals a choice point that future work on demandingness in epistemology ought to be aware of. Let's look at the challenges.

Q1. Which features of a context determine which psychological facts are salient?

Contextualist views claim that the semantic content of a substantial term—in our case, 'can'—varies depending on the value of some parameters that are determined by conversational context. We might want this picture to be filled in a little further—what are the relevant parameters, and how are they determined? These are difficult questions to answer with any precision, and we ought to have modest expectations about how completely these details can be specified.²⁵ In general, contextualists hold that the values of the parameters

²⁴Don't these considerations also *support* demanding requirements? For example, someone who accepts that requirements can ask the impossible might claim that one can nevertheless be blameless in failing to satisfy them. I think this is right. For instance, advocates of knowledge-first epistemology (Williamson forthcoming, Hawthorne and Stanley 2008, and Srinivasan 2015) standardly claim that an agent might violate a knowledge norm but have a good excuse and so be blameless. One point I would like to emphasise here is that this distinction comes with a cost. Greco (2021, 3-4) notes that to judge an agent's faultless norm violation as 'merely excused is to damn them with faint praise; excused beliefs include beliefs you form because you're a child, or insane, or drugged.' (A similar point is made in Schechter 2017, 150). Being excused is not as good as being rational. Although I avail myself of this distinction, I do so much less often than views on which *ought* does not imply *can* in epistemology. For instance, the view holds that a character like Amy in *JET LAG*, on some specifications of the case, fails to do what she ought to, but may be blameless. But on plenty of other specifications of the case, the view holds that it is not the case that Amy fails to do what she ought to. Whereas a view like the truth norm of belief holds that on *any* specification of this case, the best Amy can do is be merely excused. Fans of demanding requirements will avail themselves of the distinction more frequently than me, resulting in a picture of normative epistemology where agents might be blameless, but are rarely successful.

²⁵A number of philosophers have emphasised this point. For example, John Hawthorne (2004, 61, 68) writing about contextualism about knowledge attributions, writes: 'Ideally, one would wish for some kind of epistemic recipe book that specified exactly how features of a context would suffice to make a certain possible mistake relevant to a particular knowledge attribution. Nothing like that recipe books is currently in our possession, nor are we close to possessing one. Perhaps

are determined by a complicated function involving variables such as the interests, purposes, and presuppositions of the participants in the conversation, and the moral and practical stakes associated with the conversation. But I can make this picture clearer in a few respects.

First, it bears emphasising that it is uncontroversial that ‘can’ is a context-sensitive term. I do not need to introduce any novel mechanisms to capture the parameters and their dynamics, I can work within Kratzer’s prominent framework for ability ascriptions.

Second, I can clarify how contextual features determine one of Kratzer’s parameters; namely, the modal base. Kratzer holds that the meaning of any token utterance of ‘can’ is sensitive to a set of circumstances that is taken as given. This set of circumstances is often left implicit, but can be derived from the conversational background. One way of doing so is by adding *in view of* clauses to ability ascriptions. For example, here are three sentences involving ability ascriptions:

- A. Hydrangeas can grow here.
- B. Smokers cannot smoke here.
- C. Jude cannot attend the birthday party.

On Kratzer’s semantics, the conversational background of each claim can fix the modal base in such a way that the claims express propositions like:

- A.* *In view of climate and soil conditions*, hydrangeas can grow here.
- B.* *In view of what the law provides*, smokers cannot smoke here.
- C.* *In view of his commitments*, Jude cannot attend the birthday party.

What determines the modal base? Two general principles of pragmatics provide a sketch of its dynamics: the principle of charity and David Lewis’s rule of accommodation.²⁶ First, charity holds that we ought to interpret speakers such that their speech constitutes a sensible contribution to a conversation. Lewis (1976, 150) provides a helpful example of how charity is required to derive a modal base for an ability ascription from the conversational background.

the mechanisms by virtue of which context-dependent predicates get their extension is neither readily accessible to *a priori* reflection nor fully amenable to empirical investigation. . . Whether or not we are contextualists, we should admit that we have at best a dim appreciation of how any such confluence determines the extension in question. So it is no objection to the contextualist if she cannot give the full story here.’ See also DeRose (1995, 10), and Cohen (1999, 61).

²⁶Both principles play a significant role in Lewis’s (1996, 1979) contextualism about knowledge attributions, a view that has much in common with Kratzer’s semantics of ability ascriptions.

Suppose, when discussing who I will bring to Helsinki as my diplomatic attaché, I assert, ‘Amy cannot speak Finnish.’ Now suppose, when conducting a biological comparison of apes and humans, I assert, ‘Amy can speak Finnish.’ In the diplomat case, I am charitably interpreted as asserting something like, ‘*In view of facts about Amy’s proficiency with languages, Amy cannot speak Finnish,*’ and in the scientist case, I am charitably interpreted as asserting something like, ‘*In view of facts about her larynx and anatomy and nervous system, Amy can speak Finnish.*’

Second, the rule of accommodation holds that when a speaker makes a claim involving a context-sensitive term, the ‘conversational score’ tends to change to make the sentence true.²⁷ Suppose Amy is in a metaphysics seminar and she says, ‘I cannot divide by zero but I can run a marathon.’ Jude sceptically replies, ‘*Pft!* There is no way *Amy* can run a marathon!’ The rule of accommodation holds that Jude’s reply changes the modal base of ‘can’ from metaphysical possibility to what is possible given Amy’s actual physical traits.

This framework provides a little more detail about how the context of epistemic evaluation can make different psychological facts salient. Roughly, the meaning of any token ability attribution is partially determined by its modal base, the value of this parameter is usually supplied by conversational background, and the mechanics of the parameter’s contribution are governed by—among other things—the principle of charity and the rule of accommodation. I take it all this will be quite subtle. Nevertheless, the upshot is that someone sympathetic to a contextualist version of an epistemic *ought* implies a psychological *can* principle can respond to questions about the mechanics of contextual variation by drawing on an elegant and powerful semantics for modal auxiliaries and some plausible principles of pragmatics.

Q2. *Why not embrace a more general contextualism about ‘can’?*

On this view, the sense *itself* shifts between contexts; in one context an epistemic *ought* implies a metaphysical ‘can’, in another a volitional ‘can’, and so on. Kratzer’s framework supports this view just as much as it supports the more restricted contextualism of the view I have explored. The resulting picture is that an epistemic *ought* implies whatever sense of ‘can’ is salient in a

²⁷Lewis (1979, 354–355) was explicit that the modal base associated with ‘can’—or, ‘the boundary between the relevant possibilities and the ignored ones’—was one feature of the conversational score.

context. Is there reason to fix the contextual variance to just one sense of ‘can’? In the absence of an answer to this challenge, the semantics I have drawn on also supports a very different version of an epistemic *ought* implies *can* principle to the one I have introduced and explored.

I think this is a challenging objection. I can develop a partial response, drawing on two considerations discussed earlier. A principle that is more generally contextualist about the meaning of ‘can’ will be both stronger and weaker than the principle I introduce, and these principles face problems of their own. First, I reject voluntarism (see [Section 3](#)), so I do not think that an epistemic *ought* implies a volitional ‘can’, even in contexts where that sense is salient. Second, the weaker principles (those with a logical, metaphysical, or nomological sense of ‘can’) run afoul of the unity desideratum: none of them can motivate the range of demandingness objections that I discussed in [Section 1](#). So, insofar as someone sympathetic to vindicating demandingness objections in epistemology wishes to hold onto that methodological commitment, these principles will not be fit for purpose. I have made use of Kratzer’s semantics to argue that attributions of psychological ability are context-sensitive, but the positive case for an epistemic *ought* implies a psychological *can* principle does not rest on semantics alone. These other considerations provide some support here.

But this response only gets one so far. In [Section 4](#), I explored the case for thinking that two problems that target a psychological version of the principle can be overcome, so perhaps the problems I mention above, which target other versions of the principle, could be overcome, too. And in any case, the original complaint still stands: the machinery that allows one to be contextualist about the psychological sense of ‘can’ also supports a broader contextualism about the meaning of ‘can’ that undercuts the view I introduce.

I think this is ultimately a mark against the view I have explored. But I also think that it contains a useful lesson for future work on demandingness in epistemology. The view I have introduced is supported by three considerations. First, it provides a principle that is fit for purpose in light of a methodological commitment I termed the unity desideratum. That is, it selects a single principle that can motivate a range of demandingness objections (and can provide a blueprint for these types of considerations). Second, it can respond to two problems that have often led philosophers to think that stronger *ought* implies *can* principles are implausible in epistemology. And, third, its contextualism is supported by the standard semantics of ‘can.’

I think this challenge reveals that there is a tension between the first and the third lines of argument. Kratzer’s semantics supports a more general contextualism about the meaning of ‘can’, that produces epistemic *ought* implies *can* principles that do not allow one to vindicate a whole set of demandingness objections. The semantics of the third argument motivates a range of principles that are not consistent with the methodological commitment of the first argument. Which horn should we embrace? We might accept the unity desideratum and go looking for some *other* way to motivate and defend the kind of view that I have introduced and explored in [Section 4](#). This strategy might be pursued by developing a way to set boundaries on the acceptable variance of ‘can,’ or it might be pursued by making much less use of contextualist semantics of ability attributions. Or we could opt for the other horn. We might embrace a wider contextual variance of ‘can’ and abandon the unity desideratum—that is, abandon the project of finding a single principle that can carry the weight of every legitimate demandingness objection. This seems to be the strategy pursued recently by Jennifer Carr. She takes it that the defining feature of non-ideal epistemology is that it accepts that *ought* implies *can* in epistemology in ‘some substantive sense’—meaning that what we epistemically ought to do is ‘*somehow or other* sensitive to our cognitive limitations.’ (Carr 2022, 1136–1138. Italics hers) But, she holds, ‘can’ is so ‘wildly flexible in its interpretation,’ that the landscape of demandingness objections in epistemology will be more complicated than suggested by the unity desideratum.

6 Conclusion

I have introduced and explored a strong epistemic *ought* implies *can* principle, whereby an agent is not epistemically required to do what they are psychologically unable to. I have proposed that the principle be understood in a contextualist spirit and argued that it can meet two serious challenges that lead philosophers to reject similar principles. Strong epistemic *ought* implies *can* principles are more contentious than weaker principles. I hope to have provided a more thoroughgoing rationale for these kinds of views than they have previously received. There are positive benefits to the view I explore, and I hope to have shown that stronger principles are more plausible than many philosophers think. But the view also faces several challenges, some of which can be met, while others strike me as more vexing. This leaves the view with

a mixed scorecard. But I think it also helps us better understand the way forward in the larger project of determining whether epistemology can require the impossible.

References

- Alston, William P. 1988. 'The Deontological Conception of Epistemic Justification.' *Philosophical Perspectives* vol. 2: 257-299.
- . 1989. *Epistemic Justification: Essays in the Theory of Knowledge*. Cornell University Press.
- . 1995. *Beyond "Justification": Dimensions of Epistemic Evaluation*. Cornell University Press.
- Armendt, Brad. 1993. 'Dutch Books, Additivity, and Utility Theory.' *Philosophical Topics* vol. 21, no. 1: 1-20.
- Bykvist, Krister and Anandi Hattiangadi. 2007. 'Does Thought Imply Ought?' *Analysis* vol. 67, no. 4: 277-285.
- . 2013. 'Belief, Truth and Blindspots.' In *The Aim of Belief*, ed., Timothy Chan. Oxford University Press, 100-122.
- Carr, Jennifer. 2022. 'Why Ideal Epistemology?' *Mind* vol. 131, no. 524: 1131-1162.
- Christensen, David. 2004. *Putting Logic in Its Place: Formal Constraints on Rational Belief*. Oxford University Press.
- . 2010a. 'Higher-Order Evidence.' *Philosophy and Phenomenological Research* vol. 81, no. 1: 185-215.
- . 2010b. 'Rational Reflection.' *Philosophical Perspectives* vol. 24, no. 1: 121-140.
- . 2016. 'Conciliationism, Uniqueness, and Rational Toxicity.' *Noûs* vol. 50, no. 3: 584-603.
- Chuard, Philippe and Nicholas Southwood. 2009. 'Epistemic Norms Without Voluntary Control.' *Noûs* vol. 43, no 4: 599-632.
- Cohen, Stewart. 1999. 'Contextualism, Skepticism, and the Structure of Reasons.' *Philosophical Perspectives* vol. 13: 57-89.
- DeRose, Keith. 1992. 'Contextualism and Knowledge Attributions.' *Philosophical and Phenomenological Research* vol. 52, no. 4: 913-929.
- . 1995. 'Solving the Skeptical Problem.' *The Philosophical Review* vol. 104,

no. 1: 1-52.

Dretske, Fred. 2000. 'Entitlement: Epistemic Rights without Epistemic Duties?' *Philosophy and Phenomenological Research* vol. 60, no. 3: 591-606.

Fara, Michael. 2008. 'Masked Abilities and Compatibilism.' *Mind* vol. 117, no. 468: 843-865.

Feldman, Richard. 1988. 'Epistemic Obligations.' *Philosophical Perspectives* vol. 2: 235-256.

———. 2000. 'The Ethics of Belief.' *Philosophy and Phenomenological Research* vol. 60, no. 3: 667-695.

———. 2008. 'Modest Deontology in Epistemology.' *Synthese* vol. 161, no. 3: 339-355.

Gibbard, Allan. 2003. 'Thoughts and Norms.' *Philosophical Issues* vol. 13, no. 1: 83-98.

———. 2005. 'Truth and Correct Belief.' *Philosophical Issues* vol. 15: 338-350.

Goldman, Alvin. 1978. 'Epistemics: The Regulative Theory of Cognition.' *The Journal of Philosophy* vol. 75, no. 10: 509-523.

Greco, Daniel. 2012. 'The Impossibility of Skepticism.' *The Philosophical Review* vol. 121, no. 3: 317-358.

———. 2021. 'Justification and Excuses in Epistemology.' *Noûs* vol. 55, no. 3: 517-537.

Greenberg, Alexander. 2020. 'Should I Believe All the Truths?' *Synthese* vol. 197, no. 1: 3279-3303.

Grim, Patrick. 1985. 'Against Omniscience: The Case from Essential Indexicals.' *Noûs* vol. 19, no. 2: 151-180.

Hattiangadi, Anandi. 2010. 'The Love of Truth.' *Studies in History and Philosophy of Science* vol. 41, no. 4: 422-432.

———. 2019. 'No, One Should Not Believe All the Truths.' *Inquiry* vol. 62, no. 9-10: 1091-1103.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford University Press.

Hawthorne, John and Amia Srinivasan. 2013. 'Disagreement Without Trans-

- parency: Some Bleak Thoughts.' In *The Epistemology of Disagreement: New Essays*, eds., Jennifer Lackey and David Christensen. Oxford University Press, 9-30.
- Hawthorne, John and Jason Stanley. 2008. 'Knowledge and Action.' *The Journal of Philosophy* vol. 105, no. 10: 571-590.
- Hedden, Brian. 2015a. *Reasons Without Persons: Rationality, Identity, and Time*. Oxford University Press.
- . 2015b. 'Time-Slice Rationality.' *Mind* vol. 124, no. 494: 449-491.
- . 2017. 'Reply to Critics.' *Analysis* vol. 77, no. 3: 607-618.
- Helton, Grace. 2018. 'If You Can't Change What You Believe, You Don't Believe It.' *Noûs* vol. 17, no. 3: 1-26.
- Holton, Richard. 2008. 'Partial Belief, Partial Intention.' *Mind* vol. 117, no. 465: 27-58.
- Hughes, Nick. 2019. 'Dilemmic Epistemology.' *Synthese* vol. 196, no. 10: 4059-4090.
- Kaplan, Mark. 1996. *Decision Theory as Philosophy*. Cambridge University Press.
- Kenny, Anthony. 1975. *Will, Freedom, and Power*. Blackwell.
- Kratzer, Angelika. 1977. 'What "Must" and "Can" Must and Can Mean.' *Linguistics and Philosophy* vol. 1, no. 3: 337-355.
- . 1981. 'The Notional Category of Modality.' In *Words, Worlds, and Contexts: New Approaches in Word Semantics*, eds., H. Eikmeyer and H. Rieser. De Gruyter, 38-74.
- . 1991. 'Modality.' In *Hanbuch Semantik*, eds., A. Von Stechow and D. Wunderlick. De Gruyter, 639-650.
- Lenman, James. 2017. 'Reasons Without Humans.' *Analysis* vol. 77, no. 3: 586-595.
- Lewis, David. 1976. 'The Paradoxes of Time Travel.' *American Philosophical Quarterly* vol. 13, no. 2: 145-152.
- . 1979. 'Scorekeeping in a Language Game.' *Journal of Philosophical Logic* vol. 8, no. 1: 339-359.

- . 1996. 'Elusive Knowledge.' *Australasian Journal of Philosophy* vol. 74, no. 4: 549-567.
- Lycan, William. 1985. 'Epistemic Value.' *Synthese* vol. 64, no. 2: 137-164.
- Maier, John. 2015. 'The Agentive Modalities.' *Philosophy and Phenomenological Research* vol. 90, no. 1: 113-134.
- McHugh, Connor. 2012. 'Epistemic Deontology and Voluntariness.' *Erkenntnis* vol. 77, no. 1: 65-94.
- Mele, Alfred. 2002. 'Agents' Abilities.' *Noûs* vol. 37, no. 3: 447-470.
- Moss, Sarah. 2015a. 'Time-Slice Epistemology and Action Under Indeterminacy.' In *Oxford Studies in Epistemology* vol. 5, eds., John Hawthorne and Tamar Szabó Gendler. Oxford University Press, 172-194.
- . 2015b. 'Credal Dilemmas.' *Noûs* vol. 49, no. 4: 665-683.
- Mizrahi, Moti. 2012. 'Does 'Ought' Imply 'Can' From an Epistemic Point of View?' *Philosophia* vol. 40, no. 4: 829-840.
- Mueller, Andy and Jacob Ross. 2017. 'Knowledge Dethroned.' *Analytic Philosophy* vol. 58, no. 4: 283-296.
- Neta, Ram. 2014. 'The Epistemic "Ought."' In *Naturalizing Epistemic Virtue*, eds., Abrol Fairweather and Owen Flanagan. Cambridge University Press, 36-52.
- Podorski, Abelard. 2016. 'A Reply to the Synchronist.' *Mind* vol. 125, no. 499: 859-871.
- . 2017. 'Review of *Reasons Without Persons* by Brian Hedden.' *Utilitas* vol. 29, no. 2: 253-256.
- Pryor, Jim. 2018. 'The Merits of Incoherence.' *Analytic Philosophy* vol. 59, no. 1: 112-141.
- Rinard, Susanna. 2018. 'Reasoning One's Way Out of Skepticism.' In *The Mystery of Skepticism*, eds., Kevin McCain and Ted Poston. Brill, 240-264.
- Ross, Jacob. 2010. 'Sleeping Beauty, Countable Additivity, and Rational Dilemmas.' *The Philosophical Review* vol. 119, no. 4: 411-447.
- Ryan, Sharon. 2003. 'Doxastic Compatibilism and the Ethics of Belief.' *Philosophical Studies* vol. 114, no. 1: 47-79.

- . 2013. 'In Defense of Moral Evidentialism.' *Logos & Episteme* vol. 6, no. 4: 405-427.
- Schechter, Joshua. 2017. 'No Need for Excuses.' In *Knowledge First: Approaches in Epistemology and Mind*, eds., J. Adam Carter, Emma C. Gordon, and Benjamin Jarvis. Oxford University Press, 132-160.
- Shah, Nishi. 2003. 'How Truth Governs Belief.' *The Philosophical Review* vol. 112, no. 4: 447-482.
- Simion, Mona. 2024. 'Resistance to Evidence and the Duty to Believe.' *Philosophy and Phenomenological Research* vol. 108, no. 1: 203-216.
- Sinott-Armstrong, Walter. 1984. "'Ought" Conversationally Implies "Can".' *The Philosophical Review* vol. 93, no. 2: 249-261.
- Smithies, Declan. 2015. 'Ideal Rationality and Logical Omniscience.' *Synthese* vol. 192, no. 9: 2769-2793.
- Snedgar, Justin. 2017. 'Time-Slice Rationality and Filling in Plans.' *Analysis* vol. 77, no. 3: 595-607.
- Southwood, Nicholas. 2016. 'Does 'Ought' Imply 'Feasible'?' *Philosophy & Public Affairs* vol. 44, no. 1: 7-45.
- Srinivasan, Amia. 2015. 'Normativity Without Cartesian Privilege.' *Philosophical Issues* vol. 25, no. 1: 273-299.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford University Press.
- Velleman, J. David and Nishi Shah. 2005. 'Doxastic Deliberation.' *The Philosophical Review* vol. 114, no. 4: 497-534.
- Wedgwood, Ralph. 2002. 'The Aim of Belief.' *Philosophical Perspectives* vol. 16: 267-297.
- . 2013. 'Rational "Ought" Implies "Can."' *Philosophical Issues* vol. 23 no. 1: 70-92.
- . 2017. *The Value of Rationality*. Oxford University Press.
- Williams, Bernard. 1973. 'A Critique of Utilitarianism.' In *Utilitarianism: For and Against*, eds., J. J. C. Smart and Bernard Williams. Cambridge University Press, 77-150.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.

———. Forth. 'Justification, Excuses, and Sceptical Scenarios.' In *The New Evil Demon*, eds., Julien Dutant and Fabian Dorsch. Oxford University Press.