

Conceptual Limitations, Puzzlement, and Epistemic Dilemmas

Michael Deigan
Rutgers University

Conceptual limitations restrict our epistemic options. One cannot believe, disbelieve, or doubt what one cannot grasp. I show how, even granting an epistemic ought-implies-can principle, such restrictions might lead to epistemic dilemmas: situations where each of one's options violates some epistemic requirement. But we might instead react to such cases by taking epistemic norms to be sensitive to one's options in ways that ensure dilemmas never arise. I propose, on behalf of the dilemmist, that we treat puzzlement as a kind of epistemic residue, roughly analogous with guilt, appropriate only when one has violated an epistemic requirement. Sometimes, in bumping up against the limits of one's concepts, it is appropriate to be puzzled no matter what one believes. This lends some support to epistemic dilemmism.

If you are unable to do something, morality cannot require that you do it. So says the familiar moral ought-implies-can principle, when contraposed.¹ Some philosophers endorse ought-implies-can but still take there to be genuine moral dilemmas: situations in which every option one has involves violating a moral requirement. Sophie can choose to save either of her two children, but can only choose one. She is not required to save both, since she cannot. But she can save one and she can save the other. So it may be wrong to let one die and also wrong to let the other die, in which case Sophie does something wrong no matter what she does.² On this view, although individual

¹This principle, in various formulations, has been the topic of much controversy. It has been claimed to be subject to counterexamples (Stocker 1971; King 2014; Graham 2011), undermined by folk judgements Turri and Buckwalter (2015), and have only self-undermining theoretical motivation (Talbot 2016). For defenses of it, see Howard-Snyder (2006), Streumer (2007), Vranas (2007, 2018), Littlejohn (2012), Wedgwood (2013), and Leben (2018). I accept a moral ought-implies-can principle, but the main claims in this paper are not about moral requirements, and so will not depend on its truth.

²Putting things this way involves denying that moral requirements agglomerate, as Williams (1965) and others have observed. Some dilemmists, like Lemmon (1962,

moral requirements are restricted to what we are able to do, when taken together they may outstrip our abilities to fulfill them. Other philosophers reject this. They say that even when faced with the most horrifying choices there must always be at least some option that is permissible. Sophie does nothing wrong, according to them, if she picks a least bad option.

Arguments between the dilemmists and anti-dilemmists have not been conclusive, but they have produced a clearer picture of what packages of views are available and what the costs and benefits of these packages are. Considerations of simplicity and deontic logic have been one important source of considerations and constraints. Dilemmists must either reject agglomeration or that 'ought' implies 'can', for instance, whereas anti-dilemmists must find ways to circumscribe or rank what may have looked to be simple, general, and overriding moral prohibitions.³ Another important such source, which will be of more interest to us, is the appropriateness of moral emotions or sentiments like guilt, regret, remorse, and blame.

Dilemmists have observed that in some cases, including Sophie's, it seems appropriate for an agent to feel guilt about whatever they do, even a least bad option. Plausibly, guilt is a kind of moral 'residue', appropriate only if one has violated a moral requirement. This would indicate that in the relevant cases the agent does indeed do something wrong no matter what they do.⁴

This is an advance, but not a decisive one, for the anti-dilemmists can and have disputed both steps. Perhaps it is not appropriate for Sophie to feel guilt, given that she had no better option.⁵ Or perhaps guilt has different appropriateness conditions which do not imply that one has done wrong, but only that one has, say, causally contributed to a bad outcome.⁶ Considering moral sentiments has not settled the

150, n. 8), Sinnott-Armstrong (1988), and Tessman (2015), have instead preferred to reject ought-implies-can.

³On more recent developments of the deontic logic involved in allowing and in disallowing dilemmas, see Horty (2003) and Goble (2013).

⁴See Williams (1965), van Fraassen (1973), Marcus (1980), Sinnott-Armstrong (1988), Greenspan (1995), and Tessman (2015).

⁵This seems to be the line that Foot (1983, pp. 387–390) takes, for instance. McConnell (1978, p. 278) and Sayre-McCord (ms) suggest it is regret rather than guilt or remorse that is appropriate.

⁶See Zhao (2020) for a defense of such a view of guilt without making the connection to dilemmas. Conee (1982) in rebutting the residue argument for dilemmas takes guilt to be appropriate when one has violated a *prima facie* duty.

debate, but it has helped in mapping the philosophical terrain.

Epistemologists have recently noticed that a similar question about dilemmas arises in their neck of the woods. Can our epistemic requirements, when taken together, outstrip our abilities to fulfill them? Are there genuine epistemic dilemmas?⁷ Daniel Greco ([forthcoming](#)) has suggested that we search for epistemic parallels of the moral emotions that have been useful for making progress in the debate about moral dilemmas. He concludes pessimistically. Greco takes the best candidate for epistemic residue to be withdrawal of trust, on the model of blame. But this does not, he argues, behave in the right way in order to support the contention that there are epistemic dilemmas. And without some such residue, he thinks, it's difficult even to see what is at issue between the epistemic dilemmists and anti-dilemmists, beyond the terminology they use in expressing their epistemic evaluations.

I agree with Greco on the usefulness of searching for epistemic residue, but disagree with his pessimism. In this paper I identify an alternative candidate for an emotion that is an epistemic parallel of guilt: puzzlement. I will claim that in some cases it seems appropriate to be puzzled no matter what one believes. And plausibly, puzzlement is appropriate only if one has violated an epistemic requirement. This would indicate in the relevant cases that one does something epistemically wrong no matter what one believes. Puzzlement, then, can play the same role in supporting epistemic dilemmism that guilt plays in supporting moral dilemmism.

The most plausible cases of unavoidable appropriate puzzlement are ones in which an agent lacks the conceptual resources needed to make adequate sense of their evidence. We will thus begin by discussing conceptual limitations, the restrictions they place on agents' epistemic options, and how such restrictions might lead to an epistemic dilemma. Then we will bring in considerations about puzzlement as a way to support the view that these apparent dilemmas are genuine.

My own reaction to these considerations is to lean somewhat in favor of accepting that there are epistemic dilemmas, but I do not think there is a decisive case here, any more than there is in the moral analogue. My aim is primarily to help map the terrain, not settle the

⁷For the current state of the debate, see especially the essays in Hughes ([forthcoming\[b\]](#)) and McCain, Stapleford, and Steup ([forthcoming](#)). But see also Odegard (1993), Priest (2002), Conee (2006), Christensen (2007, 2010, 2016), Ross (2010), Turri (2012), Lasonen-Aarnio (2014, 2020), Pryor (2018), Hughes (2019), and Leonard (2020).

debate.

Additionally, the discussion will bring to light and start to chip away at two topics that I think deserve more attention from epistemologists: (i) the nature of puzzlement and its role in inquiry and (ii) how, in trying to understand the world, one should go about straining against the limits of one's conceptual repertoire. To this end, I conclude the paper by outlining a view from the dilemmist's perspective about why it makes sense for inquirers to sometimes seek out situations where they face unavoidable puzzlement.

1. Conceptual limitations restrict epistemic options

We can't do just anything. Rather, a given agent in a given circumstance has a limited range of acts which they can perform. They have a limited range of *practical options*. A traditional ought-implies-can principle means that any action that one is morally required to do must be among these options. What must I do right now? If I could end poverty by snapping my fingers, then perhaps I would be morally required to do that. But I cannot do that—it is not among my practical options—so I am not so required.

This limitation is not specific to me. Nobody has an option to end poverty by snapping their fingers, even if some suitably idealized agent could do so. Some limitations on options are even more pervasive: no possible agent has options to do what is logically impossible. Other limitations are more local. I cannot jump over 7 feet into the air, whereas a few people can. I cannot, without consulting various resources, translate the previous sentence into Italian, but many people can. How exactly to delineate an agent's options in a given circumstance is a delicate question, one on which a variety of important issues may well turn. The basic claims about options made so far should suffice for our purposes, however.

I am making these observations so that I can draw attention to a connection between options and the possibility of dilemmas. Moral dilemmas, if there are any, depend on there being restrictions on practical options. Sophie must choose which of her children will be killed and which will be saved. If she does face a genuine dilemma, it is in part because saving both is not an option for her. This is, of course, only a contingent limitation, one that some possible agents would not

face in similar circumstances, but it is a restriction nonetheless. And even if we take there to be moral dilemmas that are unavoidable by any possible agent, these too would result in part from the fact that such agents cannot do the impossible and avoid the unavoidable.

That agents have various practical limitations is, of course, a matter on which moral dilemmists and anti-dilemmists agree. They disagree only about whether Sophie's limited options must include at least one that is permissible. It is nevertheless worth emphasizing the role that restriction of options might play in generating moral dilemmas, since the fact that one's epistemic options are limited has implications which epistemologists have yet to fully appreciate, including for the possibility of epistemic dilemmas.

We can't believe just anything. And I don't mean we can't believe just anything at will. What I have in mind is orthogonal to the issue of doxastic voluntarism.⁸ Rather, the point is that a given agent in a given circumstance has a limited range of range of doxastic attitudes which they are capable of having. Agents have, in this sense, limited *epistemic options*. Providing a detailed theory of epistemic options is an important task, but not one can be undertaken here. It will suffice to present one especially clear kind of limitation on epistemic options: conceptual limitations.

There are many contents which I can have various attitudes towards. I can believe or disbelieve *that there's an even number of chairs in this room*, have this or that degree of credence *that it's later than 5:00*, be certain or doubtful *that there are hydrogen atoms in my glass of water*, and suspend judgment or settle my mind about *whether it will rain tomorrow*. To be able to have these attitudes, I must have certain basic cognitive capacities; I have to be the sort of thing that can have beliefs, credences, and so on. Pebbles and ball-point pens do not have any

⁸The most prominent discussion of epistemic ought-implies-can principles stems from Alston (1988), who argued that since we cannot believe at will, our beliefs cannot be subject to an epistemic 'ought'. Many have denied this kind of ought-implies-can principle in favor of some weaker variant (see, e.g., Hieronymi (2008) and McHugh (2012)). For present purposes we do not need any principle that would require our belief be under our voluntary control, or indeed under any kind of control.

Nor do we need a principle that rules out epistemic 'ought's in cases of compulsive belief of the kind taken by Mizrahi (2012), Côté-Bouchard (2019), and Buckwalter and Turri (2020) to be problematic for epistemic ought-implies-can principles.

For uses of various thinner epistemic ought-implies-can principles more in the spirit of the one I will be appealing to, see, for instance, Bykvist and Hattiangadi (2007), Chuard and Southwood (2009), Greco (2012), Helton (2020), and Forrai (2021).

epistemic options, since they aren't the kinds of things that can have any attitudes at all. Some relatively simple creatures may be able to have some kinds of attitudes, like belief, without being able to have others, such as suspension of judgement, in which case having these others will not be among their epistemic options.

Besides the cognitive capacities for having the relevant kinds of attitudes, though, I must also have certain concepts in order to have these attitudes towards those particular contents. I need to be able to think of chairs and even numbers, of later-than-ness, of rain, and of hydrogen atoms. There are some people who do not have the concept of an atom, so they cannot be certain that there are hydrogen atoms in my glass. This is not because they are compelled to doubt or reject the proposition. They can't do those things either. Nor does this have to do with whether their beliefs are under their control. Rather, they cannot be certain that there are hydrogen atoms in the glass because they don't have the concepts required to have any attitudes towards this proposition. When one doesn't have the concepts required for having attitudes towards some proposition, one is *conceptually limited*.

This idea of a conceptual limitation immediately runs us into two difficult questions: what is it to have a given concept? And what determines what concepts are required for having an attitude towards a proposition? One might hope for a proper theory of concepts and their possession which, together with a theory of propositions and a theory of attitudes, would show how failure to grasp this or that concept (or any of some group of concept packages) results in an agent's not being able to have doxastic attitudes towards the proposition. Alas, I have no such theories, at least not beyond the various controversial options available in the literature.⁹ But no matter: the arguments of this paper will rely on only minimal assumptions about these issues. I assume that (i) some doxastic attitudes cannot be had without the agent having sufficient grasp of certain concepts and (ii) sufficient grasp of some of these concepts requires cognitive capacities that some agents do not have.¹⁰ Assumption (i), of course, is required for there to be conceptual limitations as I understand them. Assumption (ii) will guarantee that

⁹For an entry point into that literature, see Margolis and Laurence (2014).

¹⁰I find it natural and helpful to put this in terms of concepts, but note that we can cut out the middleman: we can just say that certain doxastic attitudes cannot be had without certain cognitive capacities. So the arguments could be restated in terms acceptable to those who, like Machery (2009), prefer to the term 'concept'.

there are conceptual limitations of a certain deep sort which clearly restrict one's epistemic options.

It will be useful for us to distinguish different ways one might be conceptually limited.

Fleeting: the agent (in some sense) has the concepts, but in a moment of inattention or failure of recall, cannot access them in the way needed to have the relevant attitude.

Eliminable: the agent does not have the concepts needed, but has the capacities needed to acquire them.

Deep: the agent has neither the concepts needed nor the capacities needed to acquire them.

Riddles often turn on fleeting conceptual limitations. "I can't operate on him; he's my son," says the surgeon truly of a boy whose father has just died in a car wreck. How is this possible? Quite obvious once one accesses one's concept of a female surgeon (or, for that matter, a non-biological concept of parenthood). If it's a conceptual limitation that keeps people from realizing the answer, it's at most a fleeting one. They need only bring to bear a way of thinking they already possess, rather than acquire a new concept.

The people mentioned above who can't be certain that there are hydrogen atoms in my glass are likely to have only an eliminable conceptual limitation. Normal humans, given the right instruction, are able to acquire the required concepts without too much difficulty, at least to the degree necessary for having this relatively rudimentary thought. But they will need to learn to think something new, rather than just access what they already have.¹¹

My dog, however, is conceptually limited with respect to this proposition in a deeper way. It's not just that he will lack the required concepts until he does a bit of reading. Rather, he just does not have the cognitive capacities needed to acquire these concepts. His inability to grasp propositions about atoms is like my inability to fly.

The distinctions between these types of conceptual limitations aren't sharp, but we have no need to sharpen them. Nor do we need to take

¹¹Putting aside the views of certain Platonists and Fodorians, anyways.

them to be exhaustive.¹² What we need is the rough idea of a deep conceptual limitation, as distinguished from the shallower limitations. And now we can see what we can get from assumption (ii): that there are some deep conceptual limitations.

Now we put this idea to use. Agents' epistemic options are restricted by their deep conceptual limitations. More specifically, if an agent is deeply conceptually limited with respect to p —lacks the concepts required for having doxastic attitudes towards p and cannot acquire those concepts—then doxastic attitudes towards p are not among her epistemic options. What are the epistemic options available to my dog? It's hard to say exactly, but none of them involve believing anything about hydrogen atoms. This is because he has a deep conceptual limitation with respect to propositions about hydrogen atoms. What are the epistemic options available to me? Even harder to say, especially for me, but none of them involve having beliefs towards propositions which I cannot acquire the concepts required for grasping.

This is not to say that *only* deep conceptual limitations restrict epistemic options. Allowing that other sorts of conceptual limitations restrict epistemic options would only broaden the implications of the arguments below. Epistemic options may also be unavailable for reasons that are independent of conceptual limitations of any kind. But whatever other restrictions we take there to be, we should at least take an agent's epistemic options to be restricted by their deep conceptual limitations. This is enough to see how some conceptual limitations might give rise to epistemic dilemmas.

¹²One interesting intermediary between eliminable and deep conceptual limitations: grasping one concept requires, given the capacities of the agent, the giving up of another concept. Then we could get cases where it's true of each doxastic attitude that the agent could (at least under possible conceptual revision) have that attitude, but certain combinations of these attitudes would be impossible for this agent. Depending on how we individuate and combine concepts, we could think of this as a deep limitation on a compound concept.

Also interesting are concepts that can normally only be acquired with heavy reliance on culture that is not present to the agent. Would the average human living in 500 B.C. lack the concept of an internet search engine in an eliminable or deep way? That depends on what we wish to include in their capacities. Similarly for concepts that may require some other external cognitive aid to grasp.

2. Restricted options might lead to epistemic dilemmas

A traditional ought-implies-can principle extended to epistemology would mean that any doxastic attitude that one is epistemically required to have must be among one's epistemic options. Given what we have said above, this would mean one cannot be epistemically required to believe propositions that one is prevented by deep conceptual limitations from believing. This seems right. Perhaps my dog is missing out on some epistemic value by lacking beliefs about hydrogen atoms, but I don't think we should take him to be epistemically required to have such beliefs, even if he were to have evidence that would generate such a requirement for an agent not limited in this way. Similarly, I may be missing out on some epistemic value by not believing those propositions which are far beyond what I could come to grasp, whatever those might be, but I am not epistemically required to believe them.

Epistemic requirements seem to be sensitive to one's epistemic options, much in the way that moral requirements seem to be sensitive to one's practical options. But this does not imply that there are no epistemic dilemmas. Even if restrictions on one's epistemic options do not keep one from satisfying any single epistemic requirement, they might keep one from satisfying all of one's epistemic requirements taken together.

This point is best illustrated with examples. To make the illustrations sharp, I will use a simple model of mental states and epistemic in terms of possible worlds, along the lines of Stalnaker (1984), Lewis (1996), and others. An agent's belief state, in this model, will include certain worlds—the ways the world might actually be if the agent's beliefs are accurate—and exclude others. A belief that *P* in such a model amounts to *P* being true in all of the worlds in one's belief state.¹³

Epistemic norms will impose various requirements on these belief states. A plausible evidentialist norm, for example, says if one has conclusive evidence against *P*, then one is required to not have any *P*-worlds in one's belief state. One's belief state should exclude those worlds that one's evidence excludes. A plausible consistency norm will require that one's belief state be non-empty. If one has

¹³To be able to explicitly model what we have said about conceptual limitations above, this specification of belief will need revision, since some truth I am unable to grasp may be true at all the worlds in my belief state. For one natural way to revise the model in this direction, see Yalcin (2018). For another, see economists' models of unawareness (Schipper 2015; Steele and Steffánsson 2021).

already eliminated all of the P -worlds, then one would violate this requirement if one also eliminates all of the $\neg P$ -worlds as well. If an update would lead one to this trivial belief state with no worlds, something has gone wrong, and one is required to either not make the update or else re-include some worlds that will not be eliminated by it.¹⁴

Example 1: false belief

Developmental psychology provides a useful stock of examples of conceptual limitations. Children at certain stages of development seem to be conceptually limited in various interesting ways, many of them extensively studied.¹⁵ The concept of a false belief is an important concept whose development has received plenty of attention; let's take a look.

On the traditional understanding of the classic experiments using the false-belief task, typical children don't acquire this concept until as late as age four.¹⁶ Here is one well-known variant of the task.

SALLY-ANNE TASK

There were two doll protagonists, Sally and Anne. . . . Sally first placed a marble into her basket. Then she left the scene, and the marble was transferred by Anne and hidden in her box. Then, when Sally returned, the experimenter asked the critical Belief Question: "Where will Sally look for her marble?" If the children point to the previous location of the marble, then they pass the Belief Question by appreciating the doll's now false belief. If however, they point to the marble's current location, then they fail the question by not taking into account the doll's belief.

Baron-Cohen, Leslie, and Frith (1985)

Suppose, on a given occasion, Sally returns and looks in the basket

¹⁴This dynamic way of describing the consistency norm in terms of updates corresponds to non-prioritized models of rational belief update (Hansson 1999). It is a natural fit with the synchronic non-emptiness norm (which is equivalent to the consistency condition of Stalnaker (1984), at least assuming worlds themselves are consistent), but it is not required.

¹⁵See, for instance, Carey (2009) and Barner and Baron (2016).

¹⁶For discussion, see Barlassina and Gordon (2017, §6.3), Ravenscroft (2016, §2.2), and sources cited therein. We need not be concerned with whether the traditional understanding of these experiments is the right one; I am skeptical that it is. All we need is a *possible* case of a dilemma, so this should suffice. And once we have presented a possible case, it shouldn't be too difficult to carry over the basic structure to genuine cases of conceptual limitations, whatever those turn out to be.

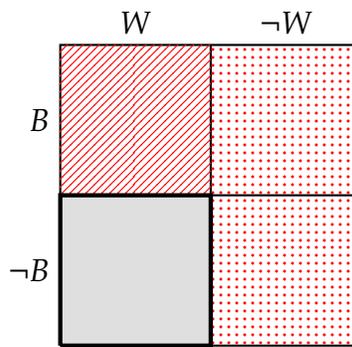
and seems surprised not to find it there. She continues looking for her marble and when she finally finds it in the basket she is pleased and ends her search. What is one to believe about this?

Let B be the proposition that Sally, at the time of her return, believed that the ball is in the basket. And let W be the proposition that Sally wanted the ball. We can divide up the space of worlds according to whether these propositions are true:

	W	$\neg W$
B		
$\neg B$		

One salient piece of evidence is that Sally does not go straight for the basket to get the ball. Given that there are no other interfering factors, this is, let us suppose, conclusive evidence against the proposition $W \wedge B$. The evidentialist norm mentioned above, then, will require of anyone who has this evidence that all $W \wedge B$ -worlds be eliminated. I will call this requirement R1. Another salient piece of evidence is that Sally is at first looking around for something, then is happy and concludes her search once she has found the ball. In this scenario, let us suppose, this is conclusive evidence against the proposition $\neg W$. The evidentialist norm will thus impose another requirement on those who possess this evidence: that one not have any $\neg W$ -worlds in one's belief state. I will call this requirement R2. Finally, the consistency norm requires of all agents to have non-empty belief states. I will call this requirement R3.

For an adult, it is easy enough to comply with these three requirements. They should believe that Sally wants to get the marble but does not believe that the marble is in the box. Limiting our diagram to the relevant worlds, they should be in the following belief state:

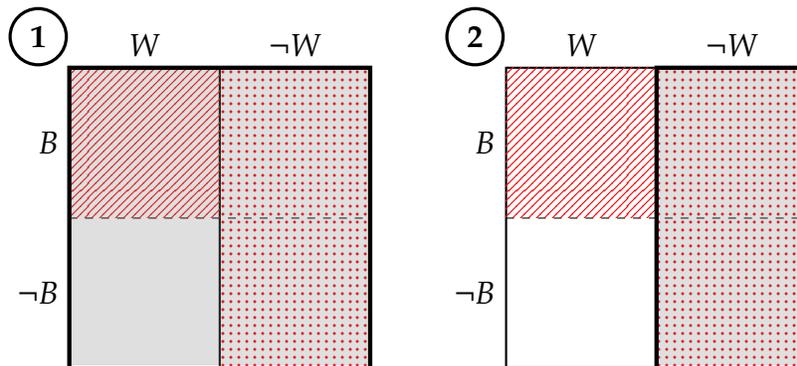


belief state: R1: R2:

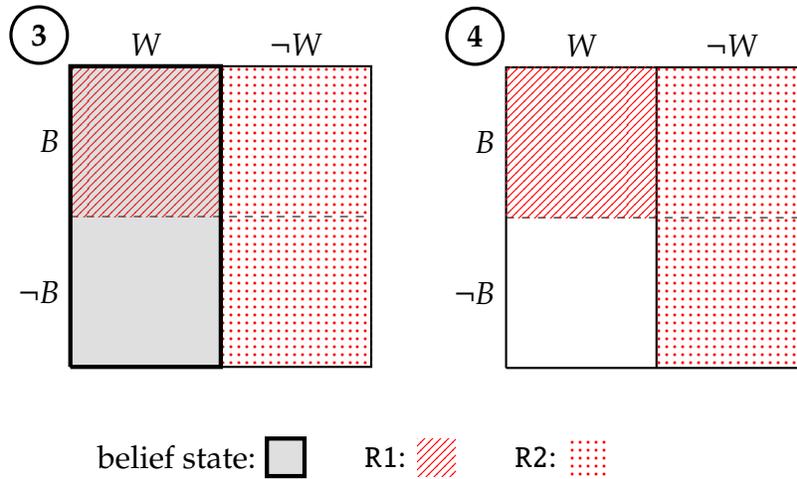
There are no $W \wedge B$ -worlds, so R1 is satisfied. There are no $\neg W$ -worlds, so R2 is satisfied. And the belief state is non-empty, so R3 is satisfied.

But now consider 3-year-old Francine, who fails the Sally-Anne false belief task because she, as a normal 3-year-old (as the traditional account would have it), cannot make a distinction between agents with true beliefs and agents with false beliefs. We should take this to be a deep conceptual limitation, given Francine’s current cognitive capacities, and so conclude that it’s not among her epistemic options to believe that Sally falsely believes the marble is in the basket. Given this, Francine will not be able to adequately make sense of all her evidence about Sally’s behavior. Indeed, given our assumptions, Francine will face a dilemma (or rather trilemma), not being able to satisfy all of R1, R2, and R3.

For Francine the only options are, at most:¹⁷



¹⁷I say ‘at most’ because it is not clear to me that we should take option 4 to be a genuine option. Removing it would not weaken the example; it would just mean we would not need to appeal to the consistency norm requirement R3.



Observe, first, that each of the three requirements are ones that Francine is able to satisfy. Options 2 and 4 satisfy R1; options 3 and 4 satisfy R2; and options 1, 2, and 3 satisfy R3. So the epistemic ought-implies-can principle is compatible with taking each of these to be requirements for Francine, even given her conceptual limitations. Nevertheless, each of Francine’s options violates one of the proposed requirements: option 1 violates both R1 and R2; option 2 violates R2; option 3 violates R1; and option 4 violates R3. This appears, then, to be an epistemic dilemma. Francine is able to satisfy of a trio of plausible epistemic requirements, but given her epistemic options, which are restricted by a deep conceptual limitation keeping her from having attitudes towards propositions about false beliefs, she cannot satisfy all of them. No matter which of her options she takes, she will violate an epistemic requirement.

Example 2: uncountably infinite sets.

We can also look for examples in theoretical developments over the course of human history. The concepts in currency today far exceed those which have been used for most of humanity’s existence. Many have been developed in order to make sense of evidence previously impossible to make sense of. Take your pick of a conceptual revolution. Were an agent deeply limited to the pre-revolution concepts, they might face a dilemma in making sense of what the post-revolutionary concepts help us make sense of.¹⁸

¹⁸For a range of approaches to thinking about conceptual development in science and mathematics, see Hempel (1952), Kuhn (1996), Lakatos (1976), Kitcher (1978), Gillies (1992), Thagard (1992), Andersen, Barker, and Chen (2006), Wilson (2006, 2020), Nersessian (2008), Strevens (2012), and Haueis (forthcoming).

Here's an example: cardinalities of sets larger than the infinite cardinality of the natural numbers.¹⁹ Cantor proves with a diagonalization argument that there is no surjective function from the natural numbers to the reals. The set of real numbers is of a cardinality greater than that of the natural numbers. The reals are uncountably infinitely many, whereas the natural numbers are merely countably infinitely many.

Now suppose someone, call him George, has a concept of infinite sets, but is deeply conceptually limited with respect to the concept of infinite sets of cardinality different than that of the natural numbers. He just can't grasp the concept of infinite cardinalities greater than countably infinitely many. Though perhaps no actual humans are like him, I don't think George is so hard to imagine. One might even sympathize by recalling one's own or one's students' first encounters with the diagonalization argument.

George could follow each stage of the proof—he could have workable concepts of the natural and real numbers, as well as sets, functions, and cardinality. In particular, he could know all the principles that are used in Cantor's argument, and so he would correctly find each step in the argument convincing. But the conclusion that these two infinite sets have different cardinalities is something that George just won't be able to grasp, let alone believe. George might be able to believe "The conclusion of the argument is true, whatever it means", but this is different from believing the conclusion. Compare: you might believe that that whatever I have written in my secret diary is true, but that doesn't mean you believe any of the propositions expressed there. So even if George can believe that the conclusion, whatever it is, is true, believing the conclusion is not among his epistemic options.²⁰

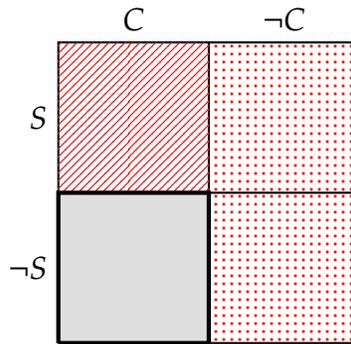
A dilemma is near at hand. Let S be the proposition that the naturals and reals have the same cardinality and C be the proposition that there

¹⁹See Hodges (1998) for a statement of Cantor's proof and an entertaining review of some common errors made in attempts to refute it. For ease of exposition, I won't focus on the issues that motivated Cantor, nor what surprised him about his results (for that, see Dauben (1992)), but instead something simpler.

²⁰Given that we are talking about mathematical truths which are presumably true in all metaphysically possible worlds, we must here be appealing to a different kind of possibility. See Chalmers (2011) and Berto and Jago (2019) for a couple ways of doing this. This is a controversial issue, and there are other ways of handling this coming from a worlds-based model, most notably by appeal to fragmentation (Stalnaker 1984; Pérez Carballo 2016; Yalcin 2018; Elga and Rayo 2020). I don't see reason for doubting that the example would carry over when translated to other frameworks, but I will not work through how to do so here.

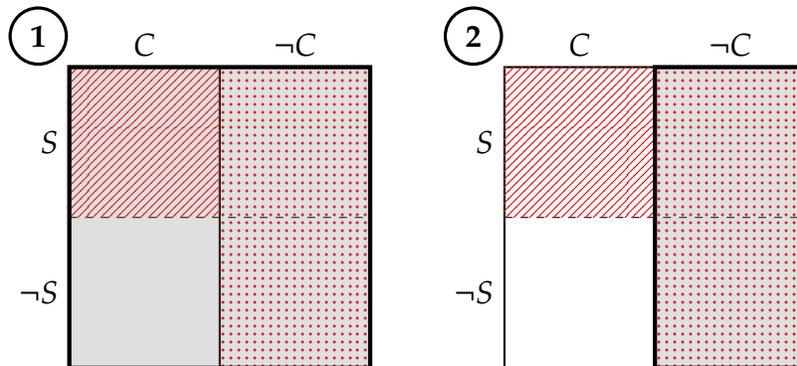
is no one-to-one function between the naturals and the reals. Let us suppose that there is conclusive evidence that S and C cannot both be true, from the definition of cardinality, and so there is an epistemic requirement $R1^*$ to not have any $C \wedge S$ -worlds in one's belief state. And let us suppose that there is conclusive evidence from Cantor's proof that C , and so there is an epistemic requirement $R2^*$ to not have any $\neg C$ -worlds in one's epistemic state. And again we suppose there is a consistency requirement $R3$ to have a non-empty belief state.

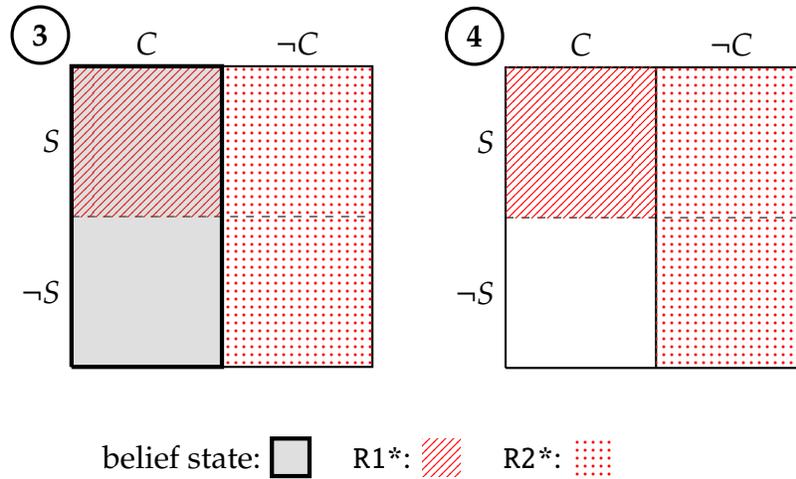
For those of us who have the concepts of countably and uncountably infinite cardinalities, there is a clearly permissible option available when we are faced with this evidence: we simply believe that the naturals and the reals have different cardinalities, despite both being infinite.



belief state: $R1^*$: $R2^*$:

But the permissible option is just not available to George, who cannot distinguish $\neg S$ -worlds from S -worlds. His options are, at most:





Like Francine, George can satisfy each of his epistemic requirements, but he cannot satisfy all of them. It seems he faces an epistemic dilemma: no matter which option he takes, he will violate an epistemic requirement.

Requirement Reformulation.

My aim with these two examples is to illustrate how an epistemic dilemma might arise given a restriction on options imposed by a conceptual limitation, not to establish that there must be epistemic dilemmas. There are plausible ways for the anti-dilemmist to accommodate them. The most plausible way to do this, in my view, is find a way to reformulate the epistemic requirements on Francine and George.

One way to do this is to leave the evidentialist norm as is, but take evidential relations to be relativized to agents based at least in part on their conceptual limitations. We could say, for instance, that the fact that Sally doesn't go directly to the basket to get the ball is evidence against $W \wedge B$ for an adult, but not for a child who lacks the concept of belief. Then even with an evidentialist norm there is no requirement R1 for Francine, and so she does not face a dilemma, since her option 2 is permissible.

An alternative way to reformulate the requirements is to reformulate the norms that generate them. We can hedge them in various ways, making them at most prima facie requirements which can be overridden by other considerations. And we can say that conflicting with the demands of other norms, given an agent's conceptual limitations, is just the kind of thing that can keep a prima facie requirement from becoming and ultima facie one. We might, for instance, allow the consistency

norm to be violated by agents in positions like Francine's and George's, rejecting R3 in these cases. Or we might allow an evidential requirement to be overridden by a conflicting evidential requirement that is at least as strong, rejecting at least one of R1(*) or R2(*) in each case.²¹

Allowing that this reformulation strategy is tenable, it is not clear why we should prefer the dilemmic picture over the non-dilemmic picture. Indeed, it is unclear what the substantive difference between these positions is. One camp will say agents in certain unfortunate positions have no very good options and the best they can do is go with a least bad option, and do not violate any epistemic requirements if they do this. The other will say these agents have no very good options and the best they can do is go with a least bad option, but do violate some epistemic requirement even if they do this. Both approaches seem to give the same advice, so what is the difference?

3. Puzzlement as epistemic residue

The requirement reformulation strategy is not only available in the epistemic case. As noted at the outset, we can replace dilemma-generating moral requirements with restricted, dilemma-avoiding counterparts. And for the same reasons it might lead us to doubt the substantiveness of the issue.

On the moral side, an influential line of argument for maintaining that we should stick to the absolute requirements and take there to be dilemmas—such as in Williams (1965), van Fraassen (1973), and Marcus (1980)—involves an appeal to the moral sentimental 'residue' which is taken to indicate wrongness and which remains after the apparently dilemmic choice has been made, no matter which option one takes. It seems that in some of the relevant situations it is appropriate or fitting to feel guilt no matter what one has chosen. If we make the natural posit that guilt is only fitting when one has done something wrong, this implies that in these situations one does something wrong no matter what one does. This argument is not irresistible, since one can deny that guilt is fitting in these situations or deny that fitting guilt implies wrongness. But even so, considering the moral emotions adds some

²¹See Lasonen-Aarnio (2014) and Worsnip (2018) for proposals along these lines for dealing with puzzles of higher-order evidence.

more substance to the debate about dilemmas: the two sides at least have different implications for when guilt is appropriate or what its appropriateness conditions are.

Greco ([forthcoming](#)) extends this idea to the epistemic case, suggesting that if we could find an epistemic counterpart of moral residue, we could give more substance to the debate about epistemic dilemmas, and that if we could find situations where this epistemic residue persists no matter what one believes, this would be evidence in favor of the pro-epistemic-dilemmas position. Greco takes appropriate withdrawal of trust to be the most plausible candidate for epistemic residue, as a kind of epistemic blame.²² He argues, however, that appropriate withdrawal of trust does not behave in the way epistemic residue would need to behave to support the view that there are epistemic dilemmas. I do not aim to dispute Greco's argument about withdrawal of trust.²³ Instead, I will propose another candidate for epistemic residue and claim that it does seem to behave in a way that makes sense from a dilemmic perspective. If this is right, it seems that the case for epistemic dilemmas is about as strong as the one for moral dilemmas, and the debate about as substantive.

I propose that we take appropriate puzzlement to be epistemic residue, an analogue of guilt. The idea is that it's appropriate to be puzzled only if one has failed to meet some epistemic requirement. If this is right, then we can make progress on evaluating our putative cases of epistemic dilemmas by seeing whether every option leaves this residue of appropriate puzzlement.

Roughly, puzzlement is an emotional response that involves a kind of uneasy feeling that something doesn't make sense and a pressure to make sense of that thing. I do not know how to pin down puzzlement more precisely with a less rough description. But I don't know how to do so for guilt, anger, or joy, either. I will proceed on the assumption that we all have experienced and recognize the state of being puzzled, as we have experienced and recognized the state of feeling guilt.

Like guilt, puzzlement can be felt inappropriately or appropriately. When I make some error in calculating and end up with a wildly

²²For different approaches to epistemic blame, see Brown (2020) and Boulton (2020, 2021). An interesting project would be to work out whether all accounts of epistemic blame lead to similar results to Greco's withdrawal of trust, or whether there are some understandings of it more useful to the dilemmist, but as far as I know this has yet to be pursued.

²³For that, see Hughes ([forthcoming\[a\]](#)).

implausible result, I may appropriately feel puzzled until I correct my mistake. And when I fail to think of any place I could have left my keys—something I feel I should be able to do—I may appropriately feel puzzled until I think of somewhere I haven't yet checked. But we can also imagine cases where I feel puzzled even though nothing has gone wrong. I might have the same feeling even if I have made the calculation correctly and the result doesn't conflict with anything else I believe, or if I know of several places where my keys might be. This would be odd, but doesn't seem impossible. In such a case puzzlement seems inappropriate; what is there to be puzzled about? This would be like feeling guilt even when one has done nothing morally criticizable. It seems that the clearest cases of appropriate puzzlement all involve intellectual error or failing on the part of the agent. Something has to be going epistemically wrong for puzzlement to be an appropriate reaction.

Guilt is at least often directed at one's past actions. It usually can only motivate us to atone and avoid violating similar moral requirements in the future. Guilt seems to say to us, "You've done something (morally) wrong; make up for it!" Puzzlement, by contrast, seems to be directed towards the present rather than the past. It motivates us to think our way out of what's puzzling us and arrive at an acceptable way of thinking. Puzzlement seems to say to us, "You're doing something (epistemically) wrong; fix it!"

As we can plausibly take guilt to be a kind of moral residue, one only fitting when one has failed to meet a moral requirement, so too we can plausibly take puzzlement to be a kind of epistemic residue, one only fitting when one is failing to meet an epistemic requirement.²⁴

Now let us return to our putative epistemic dilemmas. Is it appropriate for Francine and George to feel puzzled no matter which of their epistemic options they take? It seems that it is. If Francine believes that Sally doesn't want the marble after all, it would be appropriate to be puzzled about why Sally said she wants it, why she is satisfied when she eventually finds it, and so on. But if she doesn't believe this this, she could be appropriately puzzled about why Sally isn't looking

²⁴Note that I'm not saying that puzzlement is appropriate for all cases of failure to meet an epistemic requirement. The proposal is that failure to meet an epistemic requirement is necessary for appropriate puzzlement, not vice versa. Perhaps violations of only some kinds of requirements are appropriately puzzlement inducing. And perhaps some violations are equivalent to blameless violations of moral requirements.

for it in the box. This contradicts, after all, what a child like Francine typically expects to happen. Similarly, it's appropriate for Cantor's proof to puzzle George no matter what he ends up thinking about it.

What about the option of resigning oneself to incomprehension and withholding judgment about some of propositions that led to the problem?²⁵ This option may be the best response, given that understanding really is beyond the agents' capacities. Faced with something paradoxical, it may indeed be reasonable to admit one doesn't know what to believe, and perhaps won't ever know what to believe.

This may also help quell the puzzlement, in the sense that it's natural to proceed without giving the issue much more thought, at least assuming it's not of great practical importance, and so not having the associated emotional response. But this reaction doesn't seem to make any remaining puzzlement inappropriate. George, after some initial struggling with Cantor's proof, may give up on making sense of what's going on with the cardinalities of infinite sets. But it would not be unfitting for him to feel puzzled again if he once again considers the issue. For various reasons, he should not let it take over his life, but puzzlement here would not be inappropriate in the same way that it is inappropriate in a situation where everything is going epistemically just fine. The kind of intermittent puzzlement George might experience is familiar to philosophers. I may give up for a while thinking about the liar paradox, or the hard problem of consciousness, or whatever. But unless I've made some breakthrough, it's not inappropriate for me to feel puzzled about such things when I consider them again.

The only way for puzzlement to be inappropriate in the cases of Francine and George, it seems to me, is for them to have beliefs towards contents which they cannot grasp. There's nothing puzzling about Sally's behavior once one can recognize she has a false belief. And there's nothing puzzling about Cantor's proof once one believes infinities can have different cardinalities. It would be inappropriate to remain puzzled once one has taken these options. But given their deep conceptual limitations, the only options Francine and George are ones

²⁵The option of withholding or suspending judgement has been important to the discussion of epistemic dilemmas, since it seems to offer a way out to agents facing putative epistemic dilemmas for which there is no clear analogue for agents facing putative moral dilemmas. For anti-dilemmic takes on this, see Flowerree ([forthcoming](#)), Lord and Sylvan ([forthcoming](#)), and Simion ([forthcoming](#)); for pro-dilemmic takes, see Odegard (1993), Turri (2012), and Hughes (2019).

that make puzzlement appropriate.²⁶

Certain conceptual limitations, then, seem to lead to situations where it is appropriate to be puzzled no matter which of one's epistemic options one takes. And since puzzlement seems like a good candidate for epistemic residue, this counts in favor of taking there to be epistemic dilemmas that arise from conceptual limitations.

It does not count conclusively (or even strongly) in favor of dilemmas, though, at least not without further supporting argument. In the case of moral dilemmas, we've noted, one can reject the judgements about guilt being appropriate or else reject the proposal that moral wrongdoing is necessary for appropriate guilt. In the epistemic case one can reject the judgements about puzzlement being appropriate no matter what one thinks, or else reject the proposal that epistemic wrongdoing is necessary for puzzlement.

The latter option seems to me more promising, as well as a more likely source of interesting further debate. I don't aim to explore the matter deeply here, nor to settle it one way or the other, but I do want to indicate one direction this debate could go.

A natural first step for the anti-dilemmist would be to try to find a case of appropriate puzzlement without violation of any epistemic requirement. Here's one suggestion:²⁷ suppose Henrietta throws a cupful of sand into the air, expecting it to land on a sheet of paper in a more-or-less normally distributed random pattern on a large sheet of paper. Instead, it falls into a strikingly clear smiley face pattern. Henrietta, of course, is very puzzled by this. She tries to come up with various explanations, but there are none to be found. She remains puzzled, and reasonably so.

In fact, let us suppose, this was just a freak coincidence. There is no explanation of why the sand fell into this smiley face pattern, at least not beyond the low level explanation that due to the physical laws and initial conditions, sand grain 1 fell like this, sand grain 2 fell

²⁶In cases like these, it also seems that besides puzzlement being appropriate, *lacking* puzzlement is inappropriate. If one were simply to take one of the bad options available without any puzzlement, one would apparently be missing something. And if it's brought to one's attention but one *still* feels nothing like puzzlement, then something is more seriously off. This is similar to guilt. A moral wrongdoer can lack guilt through inattention or self-deception, in which case they're missing something, but the wrongdoer can also recognize wrongdoing but feel no guilt, which seems a more serious problem.

²⁷Thanks to [redacted] for this case.

like that, and so on. That the sand landed in a smiley face pattern was not due to some miracle or trick, it just happened. In this case, it is hard to see how Henrietta could have failed to meet any epistemic requirements. She started with a reasonable expectation, then she updated correctly when provided with new perceptual evidence, and while she didn't find the explanation she sought for, this was not due to any criticizable lack of creativity or mistaken reasoning. But it also seems that Henrietta's puzzlement was appropriate. Indeed, it would be odd if she weren't puzzled. This looks like a problem for the proposal to treat appropriate puzzlement as as a residue of epistemic requirement violation. Henrietta is appropriately puzzled, but hasn't failed to meet any epistemic requirements.

The dilemmist might respond by finding some plausible epistemic requirement that Henrietta does fail to meet. Perhaps they will say that she must have a false (even if justified) belief that there is some explanation she's missing, or that she is wondering a question with a false presupposition, and in doing so violates an epistemic requirement. I suspect, though, that responses along these lines will return the debate to a relatively unproductive stalemate, with dilemmists insisting that we count something as a requirement that anti-dilemmists insist on giving some weaker status.

A more effective response is to draw attention to the distinction between some attitude's being appropriate and its being justified. Given one's evidence, it might make sense to be angry even there is nothing to be angry about and one's justified anger turns out not to be appropriate. Suppose you've been told by a very reliable source that an acquaintance has insulted you. You would be justifiably angry at this acquaintance, but there is a sense in which this anger is not in fact appropriate. It would have been appropriate only if they had done what you reasonably thought they did.

Given Henrietta's evidence—which makes it extremely likely that she is missing some explanation for why the sand landed in a smiley pattern—her puzzlement is justified. But in fact there is nothing to be puzzled about here, since there is no explanation that she is missing. Henrietta's puzzlement is justified, but not appropriate. It would have been appropriate only if what Henrietta reasonably thought about the situation had been true.

A dilemmist can say, then, that Henrietta's puzzlement was justified but not appropriate, and that what entails an epistemic norm violation

is appropriate puzzlement, not justified puzzlement. They can thus grant that Henrietta has violated no epistemic requirement without giving up on treating puzzlement as a kind of epistemic residue.²⁸

This may answer the objection from the case of Henrietta, but the debate should not end here. More cases should be brought forth and more thoroughgoing accounts of puzzlement's appropriateness conditions should be defended in more detail. What the participants in the debate about epistemic dilemmas should be aiming for is to develop the best package of views. Part of this task will involve giving an account of what kinds of epistemic requirements there are and where these requirements come from. But another part, currently much less developed, will be to give an account of puzzlement's appropriateness conditions, ideally incorporated into a theory of emotions' appropriateness conditions more generally.

4. Seeking Puzzlement

Bertrand Russell says "the point of philosophy is to start with something so simple as not to seem worth stating, and to end with something so paradoxical that no one will believe it" (Russell 1918, p. 20). I am not sure whether he meant this in earnest, since it is surely false. Even if there is some point we can ascribe to philosophy in general, it is not this. Nevertheless, I do think there is an important insight behind Russell's remark. Inquiring well, particularly in philosophy, often involves seeking out puzzlement. Finding a new paradox that leaves everyone deeply puzzled counts as a great success, as Russell knew well. But why is that? Why try to make ourselves more puzzled than we already are? One might have thought that inquiry should aim to reduce the puzzlement we start out with, not add to it.

This may seem to be a problem for the dilemmist-friendly proposal I

²⁸I have used the terms 'justified' and 'appropriate', but what matters here isn't the terminology—we could mark the distinction with 'objective' as opposed to 'subjective appropriateness', 'fitting' as opposed to 'warranted', or 'correct' as opposed to 'rational'. The point is just that there is some status that puzzlement doesn't have in Henrietta's case that it does have in the cases we've considered involving failures to meet some epistemic requirement. The dilemmist should say that it is *that* status that indicates failure to meet an epistemic requirement, and that is the status that Francine's and George's puzzlement would have, no matter which of their options they take.

have made about puzzlement. If it is appropriate puzzlement inquirers are after, they will have to be in violation of an epistemic requirement in order to get what they seek, which hardly sounds like good inquisitive practice. But seeking inappropriate puzzlement sounds no better.

In fact, however, I think the view of puzzlement as an epistemic residue is well positioned to make sense of this phenomenon. What we are aiming for is not to produce new violations of epistemic requirements and feel puzzlement in response to that. Rather, it is uncover ways in which we are already failing to meet epistemic requirements. This is indeed an admirable thing for an inquirer to do, at least within limits, just as it is admirable to try to uncover ways in which one is failing to meet moral requirements, finding matters about which one should but does not yet feel guilt.

This also helps us see the use in finding epistemic dilemmas, where due to some conceptual limitation none of our current epistemic options are acceptable. In the literature on moral dilemmas, it is sometimes claimed that it is important to think of certain cases as moral dilemmas because it encourages us to arrange our lives and societies so as to reduce their occurrence. Nussbaum (2000) thinks we should ask not just the obvious question of what we ought to do in a given situation, but also what she calls the “tragic question” of whether any available option is morally acceptable. When an agent sees that there are cases where even his best is unacceptable because of some conflict between moral requirements, he will be prompted, as Marcus (1980, p. 133) puts it, “to arrange his own life and encourage social arrangements that would prevent, to the extent that it is possible, future conflicts from arising”. I think the analogous question is worth raising in the epistemic case. We should ask not just “What is the best to believe about this?” but also “Are any of my options epistemically acceptable?”.

When we face an epistemic dilemma due to a deep conceptual limitation, there’s not much we can do to avoid future dilemmas arising from the same source. But we can still mark the area as one where apparently good reasoning might lead us astray and perhaps to try to lay the groundwork for future inquirers to surpass us.

We might also take dilemmas to arise from eliminable conceptual limitations. Finding such dilemmas would have a more immediate payoff: we can aim to avoid them by eliminating the conceptual limitation. Sometimes this will require learning concepts that others have already mastered, but sometimes this will involve devising new

concepts. Asking the tragic question can be a way to find where conceptual development is most needed. This fits well with a traditional picture of philosophy aiming at conceptual development through consideration of thought experiments and resolution of puzzles.²⁹ We search for cases for which our current concepts are inadequate—where no matter what we believe, feel appropriately puzzled, and so must be doing something wrong. Then we attempt to eliminate the inadequacy by expanding our epistemic options to include some new ones that are epistemically permissible, not merely the best we can do.

Another picture takes philosophy to aim at reminding us of the concepts we already have, as a kind of therapy to quell our philosophical puzzlement. Some advocates of this view seem to hold that this can be the *only* aim of philosophy. They think we cannot sensibly broaden our abilities to think, but must leave everything as it is, merely drawing our attention to what we already know.³⁰ I am either less optimistic about the conceptual repertoires we find ourselves with, or more optimistic about our ability to improve them, or both. But I do think this kind of therapy is sometimes just what one needs. If we take a view on which epistemic options are restricted not just by our deep and eliminable but also by our fleeting conceptual limitations, successful therapeutical reminders can also help us avoid dilemmas by expanding our options. Raising the tragic question—if it is still fitting to call it that—is useful here as well.

It is indeed a kind of success to find that one should be puzzled no matter which of one's options one takes. The purpose of this, I suggest, is the recognition, management, and, ideally, future dissolution of epistemic dilemmas.

²⁹For a recent articulation of a view along these lines, see Scharp (2020).

³⁰See Wittgenstein (1953/2009, §§121–129).

References

- Alston, William P. (1988). "The Deontological Conception of Epistemic Justification". In: *Philosophical Perspectives* 2, pp. 257–299.
- Andersen, Hanne, Peter Barker, and Xiang Chen (2006). *The Cognitive Structure of Scientific Revolutions*. Cambridge University Press.
- Barlassina, Luca and Robert M. Gordon (2017). "Folk Psychology as Mental Simulation". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2017. Metaphysics Research Lab, Stanford University.
- Barner, David and Andrew Scott Baron, eds. (2016). *Core Knowledge and Conceptual Change*. Oxford University Press.
- Baron-Cohen, Simon, Alan M. Leslie, and Uta Frith (1985). "Does the autistic child have a "theory of mind"?" In: *Cognition* 21, pp. 37–46.
- Berto, Franz and Mark Jago (2019). *Impossible Worlds*. Oxford University Press.
- Boult, Cameron (2020). "There is a distinctively epistemic kind of blame". In: *Philosophy and Phenomenological Research*.
- (2021). "Epistemic blame". In: *Philosophy Compass*.
- Brown, Jessica (2020). "What is epistemic blame?" In: *Noûs* 54.2, pp. 389–407.
- Buckwalter, Wesley and John Turri (2020). "Inability and Obligation in Intellectual Evaluation". In: *Episteme* 17.4, pp. 475–497.
- Bykvist, Krister and Anandi Hattiangadi (2007). "Does thought imply ought?" In: *Analysis* 67.4, pp. 277–285.
- Carey, Susan (2009). *The Origin of Concepts*. Oxford University Press.
- Chalmers, David J. (2011). "The Nature of Epistemic Space". In: *Epistemic Modality*. Ed. by Andy Egan and Brian Weatherson. Oxford University Press, pp. 60–107.
- Christensen, David (2007). "Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals". In: *Oxford Studies in Epistemology*. Ed. by Tamar Szabó Gendler and John Hawthorne. Vol. 2, pp. 3–31.
- (2010). "Higher-Order Evidence". In: *Philosophy and Phenomenological Research* 81.1, pp. 185–215.
- (2016). "Conciliation, Uniqueness and Rational Toxicity". In: *Noûs* 50.3, pp. 584–603.
- Chuard, Philippe and Nicholas Southwood (2009). "Epistemic Norms without Voluntary Control". In: *Noûs* 43.4, pp. 599–632.

- Conee, Earl (1982). "Against Moral Dilemmas". In: *Philosophical Review* 91, pp. 87–97.
- (2006). "Against an epistemic dilemma". In: *Australasian Journal of Philosophy* 72.4, pp. 475–481.
- Côté-Bouchard, Charles (2019). "'Ought' implies 'can' against epistemic deontology: beyond doxastic involuntarism". In: *Synthese* 196, pp. 1641–1656.
- Dauben, Joseph (1992). "Conceptual revolutions and the history of mathematics: two studies in the growth of knowledge". In: *Revolutions in Mathematics*. Ed. by Donald Gillies. Oxford University Press.
- Elga, Adam and Agustín Rayo (2020). "Fragmentation and logical omniscience". In: *Noûs*.
- Flowerree, A. K. (forthcoming). "When in Doubt, Withhold: A Defense of Two Grounds of Rational Withholding". In: *Epistemic Dilemmas: New Arguments, New Angles*. Ed. by Kevin McCain, Scott Stapleford, and Matthia Steup. Routledge.
- Foot, Philippa (1983). "Moral Realism and Moral Dilemma". In: *The Journal of Philosophy* 80.7, pp. 379–398.
- Forrai, Gábor (2021). "Doxastic Deontology and Cognitive Competence". In: *Erkenntnis* 86, pp. 687–714.
- Gillies, Donald, ed. (1992). *Revolutions in Mathematics*. Oxford University Press.
- Goble, Lou (2013). "Prima Facie Norms, Normative Conflicts, and Dilemmas". In: *Handbook of Deontic Logic and Normative Systems*. Ed. by Dov Gabbay et al. College Publications, pp. 241–352.
- Graham, Peter A. (2011). "'Ought' and Ability". In: *Philosophical Review* 120.3, pp. 337–382.
- Greco, Daniel (2012). "The Impossibility of Skepticism". In: *Philosophical Review* 121.3.
- (forthcoming). "On the Very Idea of an Epistemic Dilemma". In: *Essays on Epistemic Dilemmas*. Ed. by Nick Hughes. Oxford University Press.
- Greenspan, P. S. (1995). *Practical Guilt: Moral Dilemmas, Emotions, and Social Norms*. Oxford University Press.
- Hansson, Sven Ove (1999). "A survey of non-prioritized belief revision". In: *Erkenntnis* 50.2, pp. 413–427.

- Haueis, Philipp (forthcoming). "A Generalized Patchwork Approach to Scientific Concepts". In: *The British Journal for the Philosophy of Science*. doi: <https://doi.org/10.1086/716179>.
- Helton, Grace (2020). "If You Can't Change What You Believe, You Don't Believe It". In: *Noûs* 54.3, pp. 501–526.
- Hempel, Carl G. (1952). *Fundamentals of Concept Formation in Empirical Science*. The University of Chicago Press.
- Hieronymi, Pamela (2008). "Responsibility for believing". In: *Synthese* 161, pp. 357–373.
- Hodges, Wilfrid (1998). "An Editor Recalls Some Hopeless Papers". In: *The Bulletin of Symbolic Logic* 4.1, pp. 1–16.
- Horty, John F. (2003). "Reasoning with Moral Conflicts". In: *Noûs* 37.4, pp. 557–605.
- Howard-Snyder, Frances (2006). ""Cannot" Implies "Not Ought"". In: *Philosophical Studies* 130.2, pp. 233–246.
- Hughes, Nick (2019). "Dilemmic Epistemology". In: *Synthese* 169, pp. 4059–4090.
- (forthcoming[a]). "Epistemic Dilemmas Defended". In: *Essays on Epistemic Dilemmas*. Ed. by Nick Hughes. Oxford University Press.
- ed. (forthcoming[b]). *Essays on Epistemic Dilemmas*. Oxford University Press.
- King, Alex (2014). "Actions That We Ought, But Can't". In: *Ratio* 27.3, pp. 316–327.
- Kitcher, Philip (1978). "Theories, Theorists and Theoretical Change". In: *The Philosophical Review* 87.4, pp. 519–547.
- Kuhn, Thomas S. (1996). *The Structure of Scientific Revolutions*. 3rd ed. Chicago University Press.
- Lakatos, Imre (1976). *Proofs and Refutations: The Logic of Mathematical Discovery*. Ed. by John Worrall and Elie Zahar. Oxford University Press.
- Lasonen-Aarnio, Maria (2014). "Higher-Order Evidence and the Limits of Defeat". In: *Philosophy and Phenomenological Research* 88.2, pp. 214–345.
- (2020). "Enkrasia or evidentialism? Learning to love mismatch". In: *Philosophical Studies* 177, pp. 597–632.
- Leben, Derek (2018). "In Defense of 'Ought Implies Can'". In: *Oxford Studies in Experimental Philosophy*. Ed. by Tania Lombrozo, Joshua Knobe, and Shaun Nichols. Oxford University Press, pp. 151–166.

- Lemmon, E. J. (1962). "Moral dilemmas". In: *The Philosophical Review* 71, pp. 139–158.
- Leonard, Nick (2020). "Epistemic dilemmas and rational indeterminacy". In: *Philosophical Studies* 177.3, pp. 573–596.
- Lewis, David (1996). "Elusive knowledge". In: *Australasian journal of Philosophy* 74.4, pp. 549–567.
- Littlejohn, Clayton (2012). "Does 'Ought' Still Imply 'Can'?" In: *Philosophia* 40, pp. 821–828.
- Lord, Errol and Kurt Sylvan (forthcoming). "Beginning in Wonder: Suspensive Attitudes and Epistemic Dilemmas". In: *Essays on Epistemic Dilemmas*. Ed. by Nick Hughes. Oxford University Press.
- Machery, Edouard (2009). *Doing Without Concepts*. Oxford University Press.
- Marcus, Ruth Barcan (1980). "Moral Dilemmas and Consistency". In: *The Journal of Philosophy* 77, pp. 121–136.
- Margolis, Eric and Stephen Laurence (2014). "Concepts". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Spring 2014. Metaphysics Research Lab, Stanford University.
- McCain, Kevin, Scott Stapleford, and Matthias Steup, eds. (forthcoming). *Epistemic Dilemmas: New Arguments, New Angles*. Routledge.
- McConnell, Terrance C. (1978). "Moral Dilemmas and Consistency in Ethics". In: *Canadian Journal of Philosophy* 8.2, pp. 269–287.
- McHugh, Conor (2012). "Epistemic Deontology and Voluntariness". In: *Erkenntnis* 77, pp. 65–94.
- Mizrahi, Moti (2012). "Does 'Ought' Imply 'Can' from an Epistemic Point of View?" In: *Philosophia* 40, pp. 829–840.
- Nersessian, Nancy J. (2008). *Creating Scientific Concepts*. The MIT Press.
- Nussbaum, Martha C. (2000). "The Costs of Tragedy: Some Moral Limits of Cost-Benefit Analysis". In: *The Journal of Legal Studies* 29.2, pp. 1005–1036.
- Odegard, Douglas (1993). "Resolving epistemic dilemmas". In: *Australasian Journal of Philosophy* 71.2, pp. 161–168.
- Pérez Carballo, Alejandro (2016). "Structuring Logical Space". In: *Philosophy and Phenomenological Research* 42.2, pp. 460–491.
- Priest, Graham (2002). "Rational Dilemmas". In: *Analysis* 62.1, pp. 11–16.
- Pryor, James (2018). "The Merits of Incoherence". In: 59.1, pp. 112–141.
- Ravenscroft, Ian (2016). "Folk Psychology as a Theory". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2016. Metaphysics Research Lab, Stanford University.

- Ross, Jacob (2010). "Sleeping Beauty, Countable Additivity, and Rational Dilemmas". In: *Philosophical Review* 119.4, pp. 411–447.
- Russell, Bertrand (1918). *The Philosophy of Logical Atomism*. First published in *The Monist*. Routledge 2010 edition.
- Sayre-McCord, Geoffrey (ms). "A Moral Argument Against Moral Dilemmas". <http://philosophy.unc.edu/files/2013/10/A-Moral-Argument-Against-Moral-Dilemmas.pdf>. Draft of May 24, 2013.
- Scharp, Kevin (2020). "Philosophy as the Study of Defective Concepts". In: *Conceptual Engineering and Conceptual Ethics*. Ed. by Herman Cappelen Alexis Burgess and David Plunkett. Oxford University Press, pp. 396–416.
- Schipper, Burkhard C. (2015). "Awareness". In: *Handbook of Epistemic Logic*. Ed. by Hans van Ditmarsch et al. College Publications, pp. 77–146.
- Simion, Mona (forthcoming). "Skepticism about Epistemic Dilemmas". In: *Epistemic Dilemmas: New Arguments, New Angles*. Ed. by Kevin McCain, Scott Stapleford, and Matthia Steup. Routledge.
- Sinnott-Armstrong, Walter (1988). *Moral Dilemmas*. Basil Blackwell.
- Stalnaker, Robert (1984). *Inquiry*. MIT Press.
- Steele, Katie and H. Orri Steffánsson (2021). *Beyond Uncertainty: Reasoning with Unknown Possibilities*. Cambridge University Press.
- Stocker, Michael (1971). "'Ought' and 'Can'". In: *Australasian Journal of Philosophy* 49, pp. 303–316.
- Streumer, Bart (2007). "Reasons and Impossibility". In: *Philosophical Studies* 136, pp. 351–384.
- Strevens, Michael (2012). "Theoretical terms without analytic truths". In: *Philosophical Studies* 160.1, pp. 167–190.
- Talbot, Brian (2016). "The Best Argument for 'Ought Implies Can' Is a Better Argument Against 'Ought Implies Can'". In: *Ergo* 3.14.
- Tessman, Lisa (2015). *Moral Failure: The Impossible Demands of Morality*. Oxford University Press.
- Thagard, Paul (1992). *Conceptual Revolutions*. Princeton University Press.
- Turri, John (2012). "A Puzzle About Withholding". In: *The Philosophical Quarterly* 62.247, pp. 355–364.
- Turri, John and Wesley Buckwalter (2015). "Inability and Obligation in Moral Judgment". In: *PLoS ONE* 10.8.
- Van Fraassen, Bas (1973). "Values and the Heart's Command". In: *The Journal of Philosophy* 70, pp. 5–19.

- Vranas, Peter B. M. (2007). "I ought, therefore I can". In: *Philosophical Studies* 136, pp. 167–216.
- (2018). "'Ought' Implies 'Can' but Does Not Imply 'Must': An Asymmetry between Becoming Infeasible and Becoming Overridden". In: *Philosophical Review* 127.4, pp. 487–514.
- Wedgwood, Ralph (2013). "Rational 'Ought' Implies 'Can'". In: *Philosophical Issues* 23, pp. 70–92.
- Williams, Bernard (1965). "Ethical Consistency". In: *Proceedings of the Aristotelian Society, Supplementary Volumes* 39, pp. 103–138.
- Wilson, Mark (2006). *Wandering Significance: An Essay on Conceptual Behavior*. Oxford University Press.
- (2020). *Innovation and Certainty*. Cambridge Elements in the Philosophy of Mathematics. Cambridge University Press.
- Wittgenstein, Ludwig (1953/2009). *Philosophical Investigations*. Trans. by G. E. M. Anscombe, P. M. S. Hacker, and Joachim Schulte. Wiley-Blackwell.
- Worsnip, Alex (2018). "The Conflict of Evidence and Coherence". In: *Philosophy and Phenomenological Research* 96, pp. 3–44.
- Yalcin, Seth (2018). "Belief as Question-Sensitive". In: *Philosophy and Phenomenological Research* 97.1, pp. 23–47.
- Zhao, Michael (2020). "Guilt Without Perceived Wrongdoing". In: *Philosophy & Public Affairs* 48.3, pp. 285–314.