Pragmatic Accounts of Justification, Epistemic Analyticity, and Other Routes to Easy Knowledge of Abstracta

Brett Topey
✉ brett.topey@plus.ac.at

Abstract

One common attitude toward abstract objects is a kind of platonism: a view on which those objects are mind-independent and causally inert. But there's an epistemological problem here: given any naturalistically respectable understanding of how our minds work, we can't be in any sort of contact with mind-independent, causally inert objects. So platonists, in order to avoid skepticism, tend to endorse epistemological theories on which knowledge is easy, in the sense that it requires no such contact—appeals to Boghossian's notion of epistemic analyticity are particularly common here, as are appeals to some broadly pragmatic account of the good standing of basic beliefs. I argue, though, that these appeals are hopeless: an argument adapted from the Benacerraf–Field challenge shows that, even if some such theory can deliver the verdict that our beliefs about abstract objects have some prima facie good standing, this good standing will inevitably be defeated.

1   Introduction

My concern here is with theories of abstract objects that are platonistic in the following minimal sense: they entail that there are such objects and that those objects are both mind-independent and causally inert. Theories meeting this description are widely endorsed, but they're also epistemologically puzzling. It's clear enough that our beliefs can't, on any naturalistically respectable story of how our cognitive

faculties work, be a result of any sort of contact with a realm of mind-independent, causally inert objects. So how can we be in good epistemic shape with respect to such objects?

Platonists, if they're to have any hope of avoiding skepticism, must provide a satisfying answer to this question. This is why platonism is so often paired with an epistemological theory on which it's *easy* for us to be in good epistemic shape with respect to such objects, in the sense that we need not be in cognitive contact with those objects in order to know things about them (or to have justified beliefs about them, or to be entitled to our beliefs about them—different epistemological theories will attribute different sorts of good standing to the beliefs in question, but these differences won't matter much for our purposes).

It bears noting, though, that epistemological theories on which good standing doesn't require cognitive contact—we can call them *easy knowledge theories*—are themselves puzzling.[1] The reason there's an epistemological puzzle here in the first place is that, if we're not in any sort of cognitive contact with (say) a realm of abstract mathematical objects, there doesn't seem to be any mechanism by which we can determine what those objects are like, or even whether they exist. So when we ask how we can be in good epistemic shape with respect to those objects, the underlying worry is that it's utterly mysterious how it is that we're supposed to have arrived at true beliefs about them. And easy knowledge theories can be of no help whatsoever in resolving that mystery—to endorse such a theory is to attempt to finesse our epistemological notions so that we can deliver nonskeptical verdicts *without* resolving the mystery, so that we may simply marvel at how unaccountably lucky it is that, despite our lack of cognitive contact with the objects in question, we've managed to arrive at the truth about them. Or, to put the point another way: to respond to our worry by invoking an easy knowledge theory is not to answer the worry but to dismiss it as epistemologically irrelevant. Insofar as we're inclined to take the worry at all seriously, then, we might suspect that the alleged epistemological benefits of easy knowledge theories are, in the end, just the familiar old advantages of theft over honest toil.

Here I argue that this suspicion is correct. In particular, I show that, if our goal is to make a genuine case for the claim that our beliefs about abstract objects are in good epistemic standing, appeals to easy knowledge theories are hopeless. For even if we can make plausible, by appeal to some such theory, that those beliefs do indeed have some prima facie good standing, it turns out, given the mystery

---

[1] Note that the sense in which knowledge (or justification, or entitlement) is "easy", on the theories I have in mind, is not the same as the sense that's in play in discussions of what Cohen (2002) has called the "problem of easy knowledge". What's relevant here is whether a theory implies that good epistemic standing doesn't require cognitive contact, not whether a theory implies that an agent can (say) bootstrap her way to knowledge that her color vision is reliable.

described above, that this good standing will inevitably be defeated. And the argument I give for this claim—an argument adapted from the Benacerraf–Field challenge—doesn't depend on any particular epistemological theory. Indeed, its only substantive epistemological premise is a claim that *all* reasonable epistemological theories, including easy knowledge theories, must accommodate: that one can't be in good standing in believing something while at the same time accepting that, insofar as that belief is true, that it's true is merely a lucky coincidence.

What I show, in short, is that, no matter what *epistemological* theory we endorse, it's just inevitable that platonism will lead to skepticism. And if that's right, finessing our epistemological theory isn't a viable strategy here; if we want to deliver the verdict that we're in good epistemic shape with respect to abstract objects, our only option is to insist on a nonplatonistic *metaphysical* picture—i.e., a picture on which abstract objects aren't both mind-independent and causally inert.

2   Easy knowledge theories

We can begin with the question of how, exactly, easy knowledge theories are supposed to work. So: suppose we endorse a platonistic theory of (say) arithmetic. How might we nevertheless deliver the verdict that our arithmetical beliefs are in good epistemic standing?

One standard approach is to endorse some form of *reliabilism*, some epistemological view on which the main condition a belief must satisfy in order to be in good standing is for it to have been reliably formed—i.e., formed in such a way that, at some particular class of (metaphysically) possible worlds (where this may be the class containing just the actual world, the class containing all worlds sufficiently similar to the actual world, or the class containing all worlds that are "normal" in Goldman's (1986) sense), a sufficient proportion of beliefs formed in that way are in fact true.[2] It's clear enough that our arithmetical beliefs meet this reliability condition: we've formed those beliefs by standard mathematical methods such as calculation and proof, and these ways of forming beliefs are *extremely* reliable despite the fact that they don't work by putting us into cognitive contact with natural

---

[2] The lack of specificity in this reliability condition is intentional—different ways of filling in the details will be associated with different reliabilist theories. The condition is consistent, for instance, with (standard and modalized versions of) both process reliabilism (see, e.g., Goldman 1979, 1986) and virtue reliabilism (see, e.g., Sosa 1991, Greco 1999). It's also consistent with safety theories (see, e.g., Sosa 1999, Williamson 2000, Pritchard 2005)—what's needed in this case is a precisification on which the relevant class of worlds is the class containing those sufficiently similar to the actual world and on which a *way of forming beliefs* is narrow enough that it produces only beliefs with the same content as the belief in question (or perhaps only beliefs with sufficiently similar contents—see, e.g., Williamson 2009). The differences between these theories won't be relevant for our purposes here.

numbers.[3] (And this is so, notice, irrespective of what class of worlds we take to be relevant—the arithmetical facts don't vary across worlds.) So reliabilism does seem to make available a story on which our arithmetical beliefs have some prima facie good standing.

Another approach is to insist that true arithmetical sentences are analytic in the epistemic sense introduced by Boghossian (1996)—i.e., are such that grasp of their contents (where this consists in being disposed to infer in accordance with the right axioms and rules) suffices for the prima facie good standing of beliefs in those contents. The thought here is that sentences can be epistemically analytic, can have the epistemological features traditionally associated with analyticity, without being analytic in the traditional metaphysical sense of being true purely in virtue of their (conventional) linguistic meanings.[4] If true arithmetical sentences are indeed analytic in this merely epistemic sense, then the prima facie good standing of our arithmetical beliefs requires only grasp of those sentences' conceptual contents and so doesn't require cognitive contact with the natural numbers: grasp of those contents is, once again, just a matter of being disposed to infer in accordance with the right axioms and rules.

Of course, an explanation is still needed of *how* a sentence can have the epistemological features associated with analyticity without being true purely in virtue of its meaning—in the absence of such an explanation, to appeal to the epistemic analyticity of true arithmetical sentences is not to give a full account of the good standing of our arithmetical beliefs but is merely to issue a promissory note. And though Boghossian has attempted various sorts of explanation here (see his 1996, 1997, 2000), the only one that turns out to have any hope of succeeding (see his 2003a, 2003b) is one that goes roughly as follows: certain of our concepts are such that, for one purpose or another, we *need* them, and so there are sentences—namely, those that we must accept in order to grasp the concepts in question—that we aren't blameworthy for accepting, even in the absence of any positive reason for taking those sentences to be true. As Schechter and Enoch (2006) have pointed out, though, the notion of concept possession appears to be an unnecessary detour in this explanation; what's essential is just the thought that there are purposes for which we *must* accept certain sentences, in which case we aren't blameworthy for

---

[3]Of course, in insisting that standard mathematical methods are reliable, we're taking for granted that our platonistic theory of arithmetic is *correct*—i.e., that mind-independent natural numbers exist and have the basic features we take them to have. But that's not a problem, by reliabilist lights. For discussion, see, e.g., Kornblith 2004.

[4]The claim that arithmetical sentences are *metaphysically* analytic isn't consistent with platonism: if arithmetical sentences were true purely in virtue of their conventional linguistic meanings, the objects those sentences are about couldn't be mind-independent. This, perhaps, explains why philosophers appealing to analyticity tend to insist that they're committed only to the epistemic variety (see, e.g., Chalmers 2012, Thomasson 2015, Nyseth 2021).

accepting them. So it's not clear that the approach here is genuinely analyticity-based after all.

Instead, what we appear to have is an instance of a third sort of approach, one on which the good standing of our arithmetical beliefs is ultimately to be accounted for in *pragmatic* terms. The thought is just that there are certain cognitive projects, certain forms of inquiry, that we're prima facie rational in engaging in, and the fact that we can't be faulted for engaging in these projects is supposed to be what explains why, when a claim is a presupposition of one of these projects—i.e., is such that we must accept it in order to engage in the project successfully—we can't be faulted for taking that claim for granted.[5] And notice: whether a claim is a presupposition of a given cognitive project has nothing whatsoever to do with whether we're in cognitive contact with the part of the world that claim is about. So, insofar as a case can be made that foundational arithmetical claims are presuppositions of some cognitive project of the right sort, this pragmatic approach makes available another kind of story on which our arithmetical beliefs have some prima facie good standing despite our lack of contact with the natural numbers.[6]

There's certainly much more to be said about the epistemological theories underlying each of these approaches. But most of the fine details aren't going to matter for our purposes here—we aren't going to be criticizing the theories themselves. Indeed, we can grant for the sake of argument that one of these theories is correct, so that our arithmetical beliefs *do* have some sort of prima facie good standing. My claim, once again, is that even if this is so, it turns out—given platonism—that those beliefs will inevitably be subject to defeat. So what's crucial to notice here is just that, according to (plausible versions of) all of the approaches I've discussed, the sort of good standing enjoyed by our arithmetical beliefs is *defeasible*.

To see why, consider that, for any of the above proposed conditions for prima facie good standing, there are cases in which a belief meets that condition while at the same time the agent has very good reason to think that it's nevertheless not likely to be true; in such cases, whatever good standing the belief has is surely going to be defeated. Suppose, for instance, that I've formed some belief about my environment by relying on my vision. Presumably, this way of forming beliefs is reliable, in which

---

[5] Versions of this sort of pragmatic account have been developed by Enoch and Schechter (2008) and by Wright (2004a, 2004b, 2014, 2016). Related (though in some respects more permissive) accounts have been developed by Putnam (see, e.g., his 1976/1983), Field (see, e.g., his 2000), Harman (see, e.g., his 2003), and Ebbs (see, e.g., his 2022).

[6] It's worth emphasizing that, on this story, what is accounted for by the relevant pragmatic considerations is the good *epistemic* standing of our beliefs. One might instead set epistemology aside altogether and insist only on a story on which we're justified in some *pragmatic* sense in holding on to the beliefs in question, but this isn't the sort of story that's under discussion here. (It also isn't the sort of story that Boghossian, Enoch and Schechter, and Wright are trying to tell.) The claims I'm targeting in this paper are claims about our beliefs' *epistemic* status.

case, according to reliabilist views, my belief has some prima facie good standing. But if I then gain strong evidence that I've been hallucinating, it's clear enough that I shouldn't hold on to my belief: it's no longer in good standing.[7] And there are going to be analogous cases of defeat for beliefs that plausibly are epistemically analytic and for beliefs that plausibly are presuppositions of some cognitive project that we're prima facie rational in engaging in.[8]

The point: platonists can do no better, by appeal to an easy knowledge theory, than to deliver the verdict that their beliefs about abstract objects are *initially* and *defeasibly* in good standing. Now to explain why it's inevitable that this good standing will in fact be defeated.

3   Epistemic luck and the Benacerraf–Field challenge

Suppose I've formed some empirical belief—say, a belief that Shirley won yesterday's lottery—on the basis of a report in what I, with very good reason, take to be today's newspaper. Presumably, this belief is (defeasibly) in good standing, on any reasonable account of what it takes for a belief to be in good standing. But suppose I then gain strong evidence that the newspaper in question isn't today's newspaper at all but is instead a novelty newspaper that was printed several years ago. It's clear enough that, on gaining this new evidence, I should give up my belief that Shirley won the lottery: that belief's good standing has been defeated. What I want to discuss is how, exactly, this defeat works.

Here, I take it, is one good way of describing what's going on in this case. There are, from my perspective, many ways the world might be with respect to who won the lottery. After I've read the newspaper report, though, I should be confident that one of those possibilities in particular—the one in which Shirley won—is the actual one. This is because I ought to take the newspaper report to be a good guide to the lottery results: given what I know about how newspapers work, I should take the newspaper's reporting to be (causally) related to what the world is really like in such a way that, if I form beliefs on the basis of that report, those beliefs will coincide with the facts. But then, when I gain the new evidence against the newspaper's authenticity, I should become confident that there's no such relationship after all.

[7]Quite a few different strategies have been tried for incorporating the phenomenon of defeat into a reliabilist epistemology (see, e.g., Goldman 1979, Alston 1988, Bergmann 2006, Grundmann 2009, Constantin 2020, Beddor 2021), but it's widely accepted that a reliabilist theory, in order to be adequate, must have *some* way of allowing for the possibility of defeat.

[8]Note also that those who defend the epistemological views that are relevant here are themselves committed to the possibility of defeat. Boghossian (2003b: 27), for his part, says that "no one should expect more than a defeasible entitlement, even from concept possession". Enoch and Schechter (2008) are explicit about their commitment to defeasibility as well, as is Wright (2016).

From my own perspective, then, my situation is as follows: of the many ways the world might be with respect to who won the lottery, I've formed the belief that one possibility in particular is the actual one, but given the way I've done so, there just isn't any connection between my having this belief and the lottery results being what they are. So what I should think is that, if the belief coincides with the fact of the matter about who won the lottery—i.e., if the belief is true—this is just a matter of luck. This, I take it, explains why the good standing of the belief has been defeated: it would be perverse of me to continue to insist that Shirley won the lottery even after I've become convinced that, insofar as my belief that she did is true, its truth is merely accidental, nothing more than a lucky coincidence. And notice that, in order for my belief to be defeated in this way, it doesn't need to be *true* that the belief is merely accidentally true if true at all; what's doing the work is just that I'm required to come to *believe* that it's accidentally true if true at all—even if my evidence for this is in fact misleading. The problem, then, is that there's a kind of incoherence involved in believing *both* that Shirley won the lottery *and* that insofar as that belief is true, its truth is just a lucky coincidence.

To make this vivid, consider that, if I hold on to both of these beliefs, I thereby commit myself to concluding, without corroboration, that I have indeed gotten unimaginably lucky—i.e., that I've managed to hit on the truth about who won the lottery despite the fact that my belief about this matter is, by my own lights, no better than a wild guess. I hope the reader will agree that accepting this conclusion would be an absurd response to my situation.[9]

This suggests a general epistemological principle, something like a coherence constraint, along the following lines:

Lucky Coincidence Thesis. An agent can't be in good epistemic standing in holding on to a belief while at the same time accepting that, insofar as that belief is true, its truth is merely a lucky coincidence.

And this principle is indeed plausible. Cases of defeat that have the general structure of our lottery case—i.e., cases in which the good standing of an agent's belief that *p* plausibly is defeated by evidence indicating that, given the way the belief was formed, there's no connection between the agent's having that belief and the part of the world relevant to whether it's true that *p* being the way it is, so that if the belief is true, its truth is a lucky coincidence—can be multiplied without end. And

---

[9]Note also that if I do have corroboration, I *still* can't reasonably conclude that the truth of my belief is unimaginably lucky: if I have corroboration that Shirley won the lottery, I shouldn't accept that insofar as my belief that she did is true, its truth is just a lucky coincidence. (This isn't to say that I can't reasonably conclude, on the basis of some newly acquired corroboration, that the truth of my *past* belief was unimaginably lucky. But given that I now have corroboration, I shouldn't regard the truth of my *present* belief as lucky.)

part of the explanation for this sort of defeat does seem to be that there would be something absurd about insisting both that *p* and that one's belief that *p*, if true, is merely accidentally true. Furthermore, it bears emphasizing that there isn't anything particularly subtle going on here, nor are there any sectarian commitments in play; one doesn't need to be in the grip of any particular epistemological theory, or even to have thought very much about epistemology, in order to recognize that in our lottery case, it would indeed be absurd for me to conclude, without corroboration, that, by relying on a way of forming beliefs that's no better from my perspective than wildly guessing, I've managed to arrive at the truth about who won the lottery.[10]

What I want to suggest, then, is that—even if we grant for the sake of argument that some easy knowledge theory is correct—we're in a position, now that the Lucky Coincidence Thesis is on the table, to see that platonists' beliefs about abstract objects just *cannot* be in good standing. Here's a statement of the argument—call it the *simple argument*:

(1) Insofar as we accept a platonistic theory of abstract objects of a particular sort, we're committed to the claim that our beliefs about those objects are accidentally true if true at all.

(2) We can't be in good standing in holding on to a belief while at the same time accepting a theory that commits us to the claim that this belief is accidentally true if true at all. (From the Lucky Coincidence Thesis.)

(C) We can't be in good standing in being nonskeptical platonists about abstract objects of any sort. (From (1) and (2).)

And from (C) it of course follows that, insofar as we can be in good standing in holding on to a platonistic theory of (say) arithmetic at all, we can do so *only* by giving up our arithmetical beliefs. So even if platonists can, by appeal to some easy knowledge theory, deliver the verdict that their arithmetical beliefs have some prima facie good standing, (C) entails that this good standing will inevitably be defeated.

---

[10]Notice: the thesis that self-ascriptions of epistemic luck are never rational—a thesis Skipper calls "No Luck" and suggests is so intuitively plausible that denying it without explaining away this plausibility would involve "an egregious amount of bullet-biting" (2023: 390)—is closely related to the Lucky Coincidence Thesis. But this latter thesis is weaker than No Luck in that it disallows only self-ascriptions of *veritic* luck (i.e., judgments to the effect that, given the way one's belief was formed, it's merely a lucky coincidence that the belief is true); it says nothing about self-ascriptions of other sorts of epistemic luck, such as judgments to the effect that it's a matter of luck that one is disposed to form beliefs in the way one does rather than in some less truth-conducive way. As a result, the sorts of putative counterexamples to No Luck discussed by Skipper—cases in which an agent learns that it's a matter of luck that she has the particular fundamental standards of reasoning she does but in which she's nevertheless purportedly rational in holding on to beliefs formed in accordance with those standards (see Elga 2008, White 2010, Schoenfield 2014)—are, regardless of whether they succeed in showing that No Luck is false, no threat to the Lucky Coincidence Thesis.

Now: the simple argument is valid, and (2) is an obvious consequence of the Lucky Coincidence Thesis. So, insofar as the Lucky Coincidence Thesis is correct, as I've suggested it is, all that's left for us to do, in order to establish our conclusion, is to show that (1) is true as well—to show, in particular, that platonists, as a result of their commitment to a view according to which they're not in any sort of cognitive contact with abstract objects, are committed also to taking any beliefs they've formed about those objects to be accidentally true if true at all.

Before we turn to that task, though, some clarificatory remarks about how the simple argument works are in order. Note, first of all, that it's *not* a radical skeptical argument. It doesn't present a challenge that platonists must answer without relying on their own beliefs—the question for platonists is instead whether, *according to their own theory of the world*, it turns out to be no more than a lucky coincidence that their beliefs about abstract objects are true. Nor does the argument depend on any controversial premises about the epistemic status of any particular belief or about what sorts of grounds a belief must have in order to be in good standing—the only substantive epistemological principle that's in play here is the Lucky Coincidence Thesis, which, as I've already suggested, is a fairly minimal coherence constraint. So we aren't going to be able to use an argument with the structure of the simple argument to call into question our beliefs in areas the epistemology of which isn't mysterious in the way the epistemology of abstracta is.

In our lottery case, for example, the theory of the world to which I initially subscribe—a theory that includes beliefs about what sorts of causal relationships there are between the newspaper's reporting and what the world is really like—is *not* a theory according to which my straightforward empirical belief that Shirley won the lottery is merely accidentally true, and so the Lucky Coincidence Thesis doesn't suggest that this belief fails to be in good standing. It's only when I gain the evidence against the newspaper's authenticity that the belief is called into question. There just doesn't seem to be any real danger that the argumentative strategy here will lead us into anything like global skepticism.

In this respect, the simple argument bears a striking resemblance to the Benacerraf-inspired argument against mathematical platonism that appears in Field's *Realism, Mathematics and Modality* (1989), an argument known variously as the Benacerraf–Field challenge, the Benacerraf problem, the access problem, or the reliability challenge. And the resemblance is no accident—what I'm offering here is one way of construing (and generalizing) that challenge.[11] The main epistemologi-

---

[11] This generalization isn't novel: it's well known already that, though Field's discussion is about mathematical platonism in particular, the challenge applies in just the same way to platonism about abstract objects of other sorts (for discussion, see, e.g., Liggins 2010). The challenge plausibly generalizes even further, to certain sorts of realism about logic (see, e.g., Schechter 2010, 2013, 2018) and morality (see, e.g., Enoch 2010), but our focus here is on theories of abstract objects in particular.

cal premise of Field's argument is that, insofar as it "*appears in principle impossible to explain*" how our beliefs about mathematical objects "can so well reflect the facts about them", this "tends to *undermine* the belief in mathematical entities, *despite* whatever reason we might have for believing in them" (1989: 26). And this claim, I take it, is plausible. But I also suspect that what *makes* it plausible is the thought that, if it is indeed impossible to explain the coincidence between our mathematical beliefs and the facts about mathematical objects—i.e., if there's no explanation available here, even in principle—then this coincidence is nothing more than an accident, just a matter of luck. If I'm right about this, Field's argument is best understood as working by (implicit) appeal to something like the Lucky Coincidence Thesis. And in that case it's more or less equivalent to the simple argument.[12]

It bears emphasizing, though, that in the literature on the Benacerraf–Field challenge, this construal isn't often endorsed (or even discussed).[13] In particular, this construal is *not* equivalent (though it is superficially similar) to a sort of interpretation that's quite common: one on which the primary epistemological principle that's in play is some general principle of coincidence avoidance. The basic thought behind this latter interpretation is that, all else being equal, a theory of the world according to which fewer inexplicable coincidences have occurred is better than one according to which more of them have occurred, in which case nonskeptical platonism has theoretical costs that its alternatives don't have.[14] And maybe this thought is correct. But it isn't what's in play in the simple argument. What's relevant there is not how many inexplicable coincidences an agent's theory commits her to positing but only whether it commits her to self-ascribing a particular sort of epistemic luck. This is significant because it means that many of the responses to the Benacerraf–Field challenge that have appeared in the literature—in particular, those responses that target the coincidence avoidance interpretation, aiming to show either that platonism's benefits outweigh its costs or simply that these costs are more minimal than we might have thought (see, e.g., Burgess and Rosen 1997, Enoch 2010, Setiya 2012)—are just irrelevant to the simple argument.[15]

---

[12]This isn't an exegetical claim—I'm suggesting that an argument with the structure of the one given by Field *ought* to be understood as deriving its force from the plausibility of the Lucky Coincidence Thesis, not necessarily that Field himself understands his argument this way.

[13]Though I've endorsed it elsewhere (see my 2021, 2022).

[14]Those who construe the Benacerraf–Field challenge in roughly this way include, among others, Enoch (2010), Liggins (2010), Schechter (2010, 2013), Setiya (2012), Baras (2017), Berry (2020), and Bhogal (2022).

[15]Another standard response to challenges like Field's is that they go by way of principles that, while suitable to the epistemology of the empirical, aren't appropriately applied in the context of the epistemology of abstracta (see, e.g., Hale 1994, Horwich 2006, Linnebo 2006). If the principle in question is some general principle of coincidence avoidance, this claim may be correct: perhaps the fact that inexplicable coincidences are improbable is something we've discovered empirically, not something we can assume to hold in abstract domains (though see Schechter 2018 for some reasons

With these clarifications out of the way, we can now move on to the question of whether the simple argument's first premise is true—i.e., the question of whether platonists are indeed committed to taking their beliefs about abstract objects to be accidentally true if true at all.

4   Lucky coincidence, responsiveness, and sensitivity

We can begin here by returning once again to our lottery case. In virtue of what am I committed to the claim that insofar as my belief that Shirley won the lottery is true, its truth is merely a lucky coincidence?

As suggested above, the explanation is roughly as follows. My perspective is one on which there are many ways the world might be with respect to the lottery results, one of which—the one in which Shirley won—I believe to be actual. But after I've gained the new evidence against the newspaper's authenticity, my perspective is *also* one on which, since there's no connection between my having that belief and the lottery results being what they in fact are, the belief isn't the product of any genuine ability on my part to discriminate among the various possible results: there's no sense in which the belief is *responsive* to what the results in fact are. These two features of my perspective, it seems, suffice to commit me to taking the belief to be accidentally true if true at all.

The general thought behind this explanation is as follows: if an agent is such that

(i) her perspective is one on which there are multiple ways a certain part of the world might be, one of which she believes to be actual, and

(ii) her perspective is also one on which the procedure by which she's chosen that possibility is in no sense responsive to the way that part of the world really is,

then the agent should take it to be no more than a lucky coincidence if the possibility she's chosen happens to be the actual one. So, if the explanation is indeed a good one, what it suggests is that this general thought (a thought that, I take it, is independently plausible) is correct—i.e., that (i) and (ii) amount to a sufficient

to think otherwise). But the claim is not at all plausible if the relevant principle is instead the Lucky Coincidence Thesis. It's clear enough that, irrespective of whether the content of a given belief is that Shirley won the lottery or that the continuum hypothesis is true, there's a kind of incoherence involved in holding on to that belief while at the same time accepting that insofar as the belief is true, its truth is just a lucky coincidence. (Of course, there remains the question of what the conditions are in which we should accept, in the case of a belief about something nonempirical, that insofar as the belief is true, its truth is just a lucky coincidence. I return to this question in the next section.)

condition for an agent to be committed to taking her belief to be accidentally true if true at all.[16]

   This sufficient condition, notice, is incompatible with at least one influential account of epistemic luck: the account developed and defended by Pritchard (see, e.g., his 2005, 2007), on which the truth of a belief is nonaccidental just in case the belief is *safe*—i.e., just in case, roughly, there are very few nearby possible worlds in which the agent falsely maintains that belief. To see why, suppose I both believe that *p* and take the nearest possible worlds where it's not the case that *p* to be far from the actual world. Then my belief that *p* is, by my lights, safe, since, if there are no nearby worlds where it's not the case that *p*, there certainly are no nearby worlds where I falsely believe that *p*.[17] But clauses (i) and (ii) may nevertheless be satisfied: whether I take the nearest possible worlds where it's not the case that *p* to be far from the actual world tells us nothing at all about whether I take my belief to be a result of choosing one possibility among many in a way that fails to be responsive to what the relevant part of the world is really like.

   Why, then, should we prefer our sufficient condition to a safety condition in the spirit of Pritchard's? The answer is that cases in which safety and responsiveness come apart in the way just described are precisely cases in which the verdicts delivered by safety-based accounts seem obviously incorrect. Consider, for example,

---

[16]There's a potential worry here. Suppose I (having seen no newspaper report) form the belief that Shirley *didn't* win the lottery, just on the basis of my knowledge that her ticket is one of an extremely large number only one of which is the winning ticket. (This is a version of a case discussed in Donaldson 2014.) Presumably, I'm not committed to taking this belief to be accidentally true if true at all, but the belief isn't in any straightforward sense responsive to what the lottery results are. So we might think this case is a counterexample to our sufficient condition. There isn't space to do full justice to this worry, but note that the case here does *not* seem to be one in which, in forming the belief that Shirley didn't win the lottery, I've chosen one possibility among many: though my perspective is indeed one on which there are multiple ways the relevant part of the world might be, the *vast majority* of those possibilities are ones in which Shirley didn't win. (Indeed, this is what led me to form the belief that she didn't win in the first place.) So clause (i) isn't satisfied, which means we can, consistently with our sufficient condition, insist that I'm not committed to taking my belief to be accidentally true if true at all. (I suspect, incidentally, that beliefs formed by enumerative induction, insofar as they fail to be responsive to their corresponding facts, are going to be amenable to a broadly similar treatment—see fn. 28 below.)

[17]Note that Pritchard's account is an account of the conditions under which a belief is *in fact* veritically lucky, while we're interested instead in whether my *perspective* is one on which my belief that *p* is veritically lucky. In order to apply Pritchard's theory in the present context, then, what we must determine is whether my belief is safe *by my lights*—i.e., whether, according to my own theory of the world, Pritchard's safety condition is met. Note also that what we're interested in is *not* whether my belief suffers from what Pritchard calls "reflective" luck—the conditions a belief must meet in order to avoid this sort of luck (conditions that, incidentally, are far more demanding than is justified by the motivations Pritchard offers for introducing the notion of reflective luck—see Leite 2006) have to do not with what follows from the agent's own theory of the world but with what follows from "what the agent is able to know by reflection alone" (Pritchard 2005: 175).

a variant of our lottery case in which, even before reading the newspaper report indicating that Shirley was the winner, I was aware that there was an extremely well-coordinated conspiracy to rig yesterday's lottery, so that, whatever the lottery results in fact are, there aren't any nearby worlds in which they're different. (We can fill in the details so as to make the results as modally robust as we like.) It's quite clear that in this case, gaining the new evidence against the newspaper's authenticity should still lead me to think that, insofar as my belief that Shirley won is true, its truth is merely a lucky coincidence—the fact that I take the lottery results to be modally robust is just irrelevant here. But according to a safety-based account of epistemic luck, this is *not* what I should think, since my perspective is one on which, insofar as my belief is true, there aren't any nearby worlds where it's false.[18] So this is a case in which safety-based accounts go wrong, and do so precisely by neglecting the close connection between responsiveness and nonaccidental truth that our sufficient condition is designed to capture.[19]

Now: with our sufficient condition in hand, we can see that platonists are indeed committed to the claim that insofar as their beliefs about abstract objects are true, this is just a lucky coincidence. Clause (ii) seems clearly to be satisfied in the case of these beliefs; platonists are, again, committed to a view on which they're not in any sort of cognitive contact with abstract objects, which means their view is *not* one on which there's any mechanism by which their beliefs about those objects can possibly be responsive to what those objects are like. (Indeed, this is what makes platonistic theories epistemologically puzzling in the first place.) And as for clause (i): easy knowledge theories, recall, can do no more than to make available the verdict that platonists' beliefs about abstract objects initially enjoy some sort of *defeasible* good standing, and accepting that one's own belief is merely defeasibly in good standing involves acknowledging that the belief might turn out not to be true—i.e., that there are other ways the relevant part of the world might turn out

---

[18]Lackey (2006) takes cases with roughly this structure to show that Pritchard's account of epistemic luck is fundamentally on the wrong track.

[19]Nor would it help to move to a stricter safety condition according to which, for a belief that p to be safe, there must be very few nearby worlds where the agent is mistaken about whether p, so that the belief, in addition to failing to be safe if there are nearby worlds where the agent falsely believes that p, also fails to be safe if there are nearby worlds where the agent falsely *disbelieves* that p (see, e.g., Miščević 2007, Pritchard 2009). For even if some principled motivation can be offered for adopting this stricter safety-based account of epistemic luck (which is doubtful; see, e.g., the discussion in Luper 2012 of the "false negatives avoidance" condition), we can again adjust the details of our lottery case, this time in such a way that my perspective is one on which there are no nearby worlds in which I fail to believe that Shirley won the lottery—say, because my new evidence, in addition to indicating that the report is from a novelty newspaper, also indicates that I was somehow *bound* to form my belief on the basis of that newspaper—and it remains entirely clear that, on gaining this new evidence, I should think that, insofar as my belief that she won is true, its truth is just a lucky coincidence.

to be.[20] So it's clear enough that (i) is satisfied as well.

I take it that these considerations are conclusive. Since (i) and (ii) are both satisfied, the first premise of the simple argument is true: platonists are committed to taking their beliefs about abstract objects to be accidentally true if true at all. And as we've already noted, the second premise is true as well, since that premise is just a consequence of the Lucky Coincidence Thesis. So the conclusion of that argument is true: it's simply not possible to be in good standing in being a nonskeptical platonist about abstract objects.

But this, we might think, is suspiciously easy. If things are as straightforward as all that, why is there so much controversy about whether the Benacerraf–Field challenge is a cogent argument against nonskeptical platonism?

I suggest that there are two main reasons for this. The first of these is that there's some temptation to insist that, because abstract objects are the way they are as a matter of metaphysical necessity, the kind of responsiveness that's in question here is irrelevant to whether we should take our beliefs about those objects to be, if true, merely accidentally true.[21] The world, the thought goes, can't *help* but be such that the objects have the features they do, and so there's no need to be responsive to what the world is in fact like in order to (nonaccidentally) form true beliefs about those objects.[22]

The second reason stems from the fact that, although platonistic theories' implications for the responsiveness of our beliefs are at the heart of what's epistemologically puzzling about platonism, the Benacerraf–Field challenge is almost never formulated in an explicitly responsiveness-based way. What tends to happen instead

---

[20]Ebbs appears to think otherwise—he suggests that, if an agent finds herself unable to "specify a way in which a belief…may actually be false", it follows that it's not epistemically possible for her that the belief is false, even if she herself acknowledges that the belief isn't immune to revision (2022: 137), and he takes this sort of inability to specify a way in which a belief may be false to have a role to play in explaining why we're reasonable in holding on to (e.g.) certain of our mathematical beliefs. I doubt that this sort of inability to specify is really what's epistemologically relevant here—I take it that what's important, as I discuss below, is whether the agent takes there to be a *risk* that she's gone wrong, and whether this is so isn't determined by whether she can specify exactly *how* she might have gone wrong. But even if I'm mistaken about this, Ebbs's strategy is, it seems to me, going to be of no help to platonists. After all, it's straightforward to specify a way in which their beliefs about abstract objects may be false: the world may simply fail to contain any objects that are both mind-independent and causally inert.

[21]It's sometimes suggested (e.g., by Field himself) that truths about abstract objects aren't necessary in any genuine sense, in which case this sort of response to the Benacerraf–Field challenge is just based on a mistake. But this move isn't available in the present context. The simple argument, remember, isn't a radical skeptical argument: our question is whether *platonists' own theory of the world*—where this includes their theory of what the metaphysically necessary truths are—is one on which their beliefs about abstract objects are merely accidentally true.

[22]Lewis (1986: sect. 2.4), Hale (1994), Pust (2004), and Baras (2017) all express something in the neighborhood of this thought. See also fn. 24 below.

is that the challenge is formulated in terms of a modal property that's supposed to serve as a *proxy* for responsiveness—namely, sensitivity, where an agent's belief that *p* counts as sensitive just in case, roughly, in the nearest worlds where it's not the case that *p*, the agent refrains from believing that *p* (or, equivalently: just in case, if it weren't the case that *p*, the agent wouldn't believe that *p*).[23] Perhaps the thought is that we have a better grip on the modal notions relevant to sensitivity than we do on the notion of responsiveness itself. The problem, though, is that the usual formulations of the challenge in terms of sensitivity turn out *not* to be cogent: even setting aside cases of metaphysical necessity, there are lots of beliefs that clearly aren't undermined despite the fact that they very obviously fail to be sensitive. The result is that there's some temptation to retreat from the project of advancing a critique of platonism in the spirit of the simple argument, where this retreat involves either adopting some entirely different way of understanding the Benacerraf–Field challenge (such as the coincidence avoidance interpretation mentioned above) or concluding that the challenge isn't cogent after all.

In fact, though, neither of these reasons turns out to reveal any genuine problem for the simple argument; further reflection makes clear that both of the temptations just mentioned are to be resisted. The first is just based on a misunderstanding of what the problem for platonism is supposed to be. And as for the second, the correct diagnosis here is not that our responsiveness-based critique of platonism fails to be cogent but simply that sensitivity turns out to be an imperfect proxy for responsiveness. We can see why these things are so by examining the temptations in turn.

4.1 Responsiveness and metaphysical necessity

Once again, the first of these temptations is to claim that, if our perspective is one on which abstract objects have the features they do as a matter of metaphysical necessity, responsiveness isn't relevant to whether we should take it to be a matter of luck whether our beliefs about those objects are true: since there's only one possible way for the world to be with respect to what those objects are like, we need not be able to discriminate among possibilities in order to get at the truth about those objects. Sometimes—in particular, when what's under discussion is a sensitivity-based formulation of the Benacerraf–Field challenge—this takes the form of the claim that when it's necessarily true that *p*, there aren't any possible worlds where

---

[23]Roush (2005) is explicit about the fact that sensitivity is intended to capture a kind of responsiveness to what the world is like, though her concern is with accounts of knowledge in the spirit of Nozick's (see his 1981: chap. 3) rather than with the Benacerraf–Field challenge. Sensitivity-based formulations of the latter challenge have been discussed by Field himself as well as by Linnebo (2006), Liggins (2010), Donaldson (2014), Clarke-Doane (2016a, 2016b, 2020a, 2020b, 2022), and Berry (2020), among others.

it isn't the case that *p*, in which case, if we believe that *p*, this belief is trivially sensitive (since, on the standard semantics for counterfactuals, a counterfactual whose antecedent isn't true in any possible worlds—i.e., a counterpossible—is vacuously true), or perhaps isn't evaluable for sensitivity at all (since, if there aren't any possible worlds where it's not the case that *p*, there's no answer to the question of what agent believes in the nearest possible worlds where it's not the case that *p*).[24] The underlying thought, though, remains the same: according to a view on which the realm of abstracta is the way it is as a matter of metaphysical necessity, there aren't multiple possibilities among which to discriminate, which means asking whether our beliefs about abstract objects result from an ability to discriminate among possibilities isn't going to be of any help in determining whether those beliefs are, if true, merely accidentally true.

What this amounts to, in the context of our sufficient condition, is the thought that if an agent's perspective is one on which her beliefs about abstract objects are true as a matter of metaphysical necessity, then clause (i) fails to be satisfied: the agent's perspective isn't one on which there are multiple possibilities one of which she takes to be actual. It of course follows that our sufficient condition isn't met, irrespective of whether (ii) is satisfied (i.e., irrespective of whether the agent's perspective is one on which she's formed her beliefs about abstract objects in a way that fails to be responsive to what those objects are like), in which case we can't show, by appeal to that sufficient condition, that the agent is committed to taking those beliefs to be accidentally true if true at all.

But when we put things this way, it's immediately evident that the thought here is based on a misunderstanding of what our sufficient condition is supposed to be doing. The reason failures of responsiveness are supposed to be worrying in the first place is that, from the agent's own perspective, if her beliefs fail to be responsive to the part of the world they're about, this indicates that the beliefs don't result from any genuine ability to pick out, from all the candidates for being her actual situation, the situation she's in fact in. Clause (i), then, is just intended to leave room for us to allow that there are cases in which getting at the truth doesn't *require* any such ability (e.g., cases in which most or all of the candidates are situations where the beliefs in question are true). So to insist that this clause fails to be satisfied in cases in which the agent takes her beliefs to be true as a matter of metaphysical necessity is to insist that, from the agent's perspective, the candidates for being her actual situation are worlds that, according to her theory, are metaphysically possible. But

---

[24]Linnebo (2006), Clarke-Doane (2016a, 2016b), Schechter (2018), and Berry (2020) all express some sympathy with this sort of response to sensitivity-based formulations of the Benacerraf–Field challenge (though Clarke-Doane has more recently suggested (see his 2020a: sect. 5.7) that this isn't a good response, on the grounds that truths about abstracta aren't genuinely necessary in any "maximal" sense).

this is a mistake. The candidates, as previously mentioned, are simply ways things might, from the agent's perspective, turn out to be—i.e., situations the actuality of which the agent hasn't definitively ruled out from the start—and so, unless she takes herself to be infallible about what the metaphysical possibilities are, whether she takes her beliefs to be true as a matter of metaphysical necessity is irrelevant to whether clause (i), properly understood, is satisfied.

We can make this vivid by considering variants of our lottery case in which the content of the newspaper report is not a claim about lottery results but is instead some claim that, according to my own theory of the world, is necessarily true if true at all—e.g., that Shirley's biological parents (insofar as Shirley exists at all) are Geraldine and Leslie, or that every glucose molecule contains twelve hydrogen atoms, or that the continuum hypothesis is true. In these variants, it's quite clear that, despite my commitment to the metaphysical necessity of the facts in the relevant domains, gaining the new evidence against the newspaper's authenticity should still lead me to give up my belief in the claim in question, and for the same reason as in the original case: my perspective is one on which I've chosen among the various ways the relevant part of the world might be using a procedure that's no better than wildly guessing. And the reason this is worrying is precisely that, from my perspective, there's a risk that it will turn out that I've guessed incorrectly, that (say) Shirley's parents will turn out *not* to be Geraldine and Leslie; unless I take there to be some risk that I've gone wrong, even coming to accept that I've formed my belief by wildly guessing won't lead me to judge that, insofar as the belief is true, its truth is a lucky coincidence. But whether I take there to be any such risk is just determined by whether I take there to be *candidates for being my actual situation* where Shirley has other parents, and so, since I don't take myself to be infallible about what the metaphysical possibilities are, whether my theory of the world entails that it's metaphysically possible that she has other parents is irrelevant.

If what's under discussion is a sensitivity-based formulation of the Benacerraf–Field challenge in particular, our conclusion here can be formulated as the claim that when an agent is evaluating her own beliefs in the context of this challenge, the worlds she should take to be relevant to whether those beliefs are sensitive are not the worlds she takes to be metaphysically possible but the worlds that are *epistemically* possible for her, where what it takes for there to exist an epistemically possible world where *p*, for an agent, is just for the agent not to have definitively ruled out from the start that her actual situation is a situation where *p*.[25] But the larger point

[25] For a somewhat different argument for this same conclusion, see my "Realism, Reliability, and Epistemic Possibility" (2021), and for an explanation of how to construct a semantics that includes worlds that are epistemically possible in the relevant sense, see Collin 2018, which refines some ideas presented in Chalmers 2011. Incidentally, some of our epistemically possible worlds are presumably going to turn out to be metaphysically impossible, and it's sometimes suggested (see, e.g., Berry 2020) that appealing to metaphysically impossible worlds in this context is problematic,

remains the same: whether platonists take the realm of abstracta to be the way it is as a matter of metaphysical necessity has no bearing whatsoever on whether there are multiple possibilities here, in the sense relevant to whether platonists are committed to taking their beliefs about abstract objects to be accidentally true if true at all. What guarantees that there are indeed multiple such possibilities—i.e., that clause (i) of our sufficient condition is satisfied—is, as suggested above, simply that platonists are committed to accepting that the beliefs in question are merely defeasibly in good standing.

4.2   Insensitive but nonaccidentally true beliefs

As for the second temptation, it arises, again, from the fact that the usual sensitivity-based formulations of the Benacerraf–Field challenge fail to be cogent. The problem with these formulations is that one of the principles on which they depend—namely, that an agent's coming to accept that one of her beliefs is insensitive suffices to undermine the good standing of that belief—is false: there are cases in which an agent's belief clearly remains in good standing despite clearly failing to be sensitive (even when it's uncontroversial that there are multiple possibilities here—i.e., that clause (i) is satisfied). Typically these are cases in which some insensitive belief is such that its truth is obviously entailed by the truth of a belief that (by the agent's lights) *is* sensitive. Consider, for example, White's case of the jury members who, on gaining a great deal of evidence of the defendant's guilt, believe both that the defendant is guilty and that it's not the case that "the defendant was framed and all the evidence was planted" (2010: 581). Insofar as the former belief is in good standing—which it is—the latter one is in good standing as well, since the jury members can be sure that any world where the former one is true is thereby a world where the latter one is true.[26] This despite the jury members' awareness that the latter belief is insensitive, since, had the defendant been framed, they'd have exactly the same evidence they now have. But sensitivity, recall, is intended to be a modal proxy for responsiveness. So the temptation, in the face of cases like this, is to abandon our responsiveness-based critique of platonism—i.e., to allow that a

---

on the grounds that we haven't fully worked out how to think about similarity in a semantics that includes such worlds. But the utility of this semantics doesn't require our having fully specified a similarity metric, any more than it does in the case of a more traditional semantics that includes only metaphysically possible worlds (a semantics for which, notoriously, we *also* haven't fully specified a similarity metric). What's important for our purposes is just that platonism, since it's a view on which there's no contact between the causal realm and the realm of abstracta, is thereby a view on which, on *any* reasonable way of specifying our similarity metric, the nearest worlds where the realm of abstracta is different in some respect aren't going to be worlds where the causal realm is also different.

[26]I assume, with White, that good standing is closed under this sort of obvious entailment, though this has sometimes been denied—see, e.g., Nozick 1981: chap. 3.

belief can remain in good standing even in a case in which clauses (i) and (ii) of our sufficient condition are both satisfied.

But to abandon our critique in this way would be to take entirely the wrong lesson from cases like White's. To see why this is, note that White's jury members do have one belief that, according to their own theory of the world, *is* responsive to what the relevant part of the world is like: the belief that the defendant is guilty. What the truth of this belief requires is just that a particular part of the world be a particular way—i.e., that it be such that the defendant committed the crime—and the jury members' theory is one on which, had this part of the world not been that way, they wouldn't have gained the evidence they now have and so wouldn't believe that the defendant is guilty. But note also that, insofar as this part of the world is indeed such that the defendant committed the crime, it's thereby such that the defendant wasn't framed in the relevant way, so that the truth of the jury members' belief that the defendant is guilty is sufficient for the truth of their belief that it's not the case that the defendant was framed and all the evidence was planted. Furthermore, the jury members *accept* that the truth of the former belief is sufficient for the truth of the latter one. Indeed, this is part of what explains why they have the latter belief in the first place: by their own lights, the grounds on which they believe that the defendant is guilty are certain also to be sufficient to ground the belief that the defendant wasn't framed. So the latter belief, despite its obvious insensitivity, does turn out to be responsive (albeit in an inherited way) to what a particular part of the world is like such that this part of the world being the way it is suffices for the truth of that belief.[27]

In short, this just isn't a case in which the agents should conclude that the belief in question fails to be responsive to the part of the world relevant to its truth. So clause (ii) of our sufficient condition isn't satisfied. And if that's right, we can, consistently with that condition, insist that the agents aren't committed to taking the belief to be accidentally true if true at all, which means we don't have a route, via the Lucky Coincidence Thesis, to the claim that the good standing of the belief is undermined. What cases with this structure indicate, then, is not that there's something wrong with our responsiveness-based critique of platonism.[28] It's simply that

[27] For a more thorough discussion of this point, see my "Saving Sensivity" (2022).

[28] It's sometimes suggested (see, e.g., Korman 2014) that there are also cases of a different sort in which beliefs remain in good standing despite failing to be sensitive: cases of beliefs formed by enumerative induction (e.g., my belief that the sun will rise tomorrow). Now: it's plausible that, in cases in which such beliefs are true, this isn't an accident. And insofar as we should think that such beliefs genuinely fail to be responsive to the parts of the world relevant to their truth—as it's standardly assumed that we should (though see, e.g., Cross 2010, Wallbridge 2018)—these are cases in which clause (ii) is satisfied. So we might think these cases are counterexamples to our sufficient condition. But notice that in these cases, there's a belief in the background that pretty clearly *is* responsive to the part of the world relevant to its truth: the belief that (roughly) there's some lawlike

sensitivity turns out not to be necessary for responsiveness.

We can, if we like, generate a better modal proxy for responsiveness by modifying our sensitivity condition so as to handle cases with this structure. (This is precisely what I do in my "Saving Sensitivity" (2022): I devise a recursive condition for responsiveness such that a belief satisfies the base clause just in case it's sensitive and such that the inductive clause captures exactly the sort of inherited responsiveness exhibited by the jury members' belief that the defendant wasn't framed, and he argues that this condition (with one other minor modification in place) turns out to be entirely counterexample-free.[29]) But it's not clear that we need to do so. We already have enough of a grip on the sort of responsiveness that's relevant here that we can conclude, with confidence, that agents in cases with this structure should indeed take their beliefs to be responsive to the parts of the world relevant to their truth, which means our sufficient condition isn't met. And this is all that's required in order for it to be clear that these cases aren't counterexamples to our sufficient condition and so have no tendency whatsoever to suggest that our critique of platonism fails to be cogent.[30]

---

explanation of why the agent might have observed the pattern in question in the first place. (If there's no such belief in the background, there's no reason to expect the next instance to fit the pattern. And the belief is responsive to the relevant part of the world simply because, if there hadn't been anything lawlike going on, the agent wouldn't have observed any pattern.) And though there's no inherited responsiveness here, there's inheritance of another kind: if we narrow the class of possibilities to those where there the background belief is true—i.e., where there is in fact a lawlike explanation of the relevant kind—then the *vast majority* of those possibilities turn out to be ones where the next instance does indeed fit the pattern. (The vast majority, but not all: the fact that there's a lawlike explanation of why the sun has risen every day up to this point doesn't guarantee that the sun will rise tomorrow.) So it seems that, just as in the case discussed in fn. 16 above, clause (i) isn't satisfied, which means we don't have a counterexample to our sufficient condition.

[29]The point of the second modification is to ensure that a belief meets the condition only if it's responsive to *all* the features of the world that are required for its truth. This is needed in order to handle what Kripke (2011) calls the "absorption phenomenon": the fact that, if an agent has a conjunctive belief that (say) her left shoelace is untied and tomorrow she'll inherit a fortune from some as-yet-unknown relative, then this belief, if true, is sensitive, just because the nearest worlds where her shoelace is tied are nearer than the nearest worlds where she doesn't inherit the fortune.

[30]Clarke-Doane suggests another route to the conclusion that some beliefs (even in cases in which clause (i) is satisfied) aren't undermined despite clearly failing to be responsive. He begins by claiming that (insofar as there can be nonvacuous counterfactuals whose antecedents are false in all metaphysically possible worlds) our metaphysical beliefs about the composition conditions of macroscopic objects don't seem to be sensitive: "had particles arranged 'paper-wise' failed to compose a piece of paper, it seems that we still would have believed that they did" (2020a: 145). Nor would modifying our sensitivity condition be helpful here; the beliefs in question genuinely don't seem to be responsive to what the relevant bridge laws in fact are. At first sight, this seems unobjectionable—the epistemology of the metaphysics of composition seems to be mysterious in broadly the same way that the epistemology of abstracta is (see, e.g., Korman 2014). But Clarke-Doane argues that our beliefs about composition conditions *must* be in good standing because, if they weren't, an unacceptable skepticism about everyday perceptual claims would result: "If I

5   Concluding remarks

In the end, then, things *are* as straightforward as I've suggested they are. By our sufficient condition—a well-motivated condition that stands up to scrutiny—platonists are indeed committed to taking their beliefs about abstract objects to be accidentally true if true at all: their commitment to a view on which those beliefs are merely defeasibly in good standing ensures that clause (i) of the condition is satisfied, and their commitment to a view on which they're not in cognitive contact with abstract objects ensures that clause (ii) is satisfied. That is, the first premise of the simple argument is indeed true. So, given the second premise—which is, again, just a consequence of the Lucky Coincidence Thesis—platonists, if they want to remain in good standing in holding on to their beliefs about abstract objects, must give up their platonism. The simple argument is sound after all.

What this means is that there's no hope that an easy knowledge theory can save platonism from skepticism; no amount of epistemological maneuvering will change the fact that the simple argument is sound. If we wish to avoid skepticism, we must adopt a nonplatonistic *metaphysical* picture, one on which cognitive contact with abstract objects is possible.[31] And this is just to say that we must adopt

believe that I am looking at a piece of paper, but my belief in the metaphysically necessary bridge law that particles arranged 'paper-wise' compose a piece of paper is undermined, then it is hard to see how my belief that I am looking at a piece of paper could fail to be" (2020a: 145; largely the same argument also appears in his 2016a, 2016b, 2020b, 2022).

It seems to me, though, that this line of reasoning is just confused. When an agent, in her everyday life, takes herself to be looking at a piece of paper, she either is or isn't thereby taking a stand on some substantive metaphysical thesis about composition. If she isn't—if her belief, in order to be correct, requires nothing more than that she be looking at some particles arranged paperwise—then there's no failure of responsiveness here, in which case there's no danger of skepticism. (My own opinion, fairly widely shared among metaphysicians of various stripes (see, e.g., van Inwagen 1990, Dorr 2005, Thomasson 2007, Chalmers 2009, Cameron 2010, Sider 2013), is that it's overwhelmingly plausible that this is in fact the situation.) But if she is—if her belief, in order to be correct, requires, for example, that the substantive metaphysical thesis of mereological nihilism be false—then the resulting skepticism is not unacceptable: the epistemology of the metaphysics of composition is indeed mysterious (unless, *contra* Clarke-Doane, our beliefs here *are* responsive to the relevant metaphysical facts, as Korman (2015: chap. 7) insists they are), and so, insofar as our everyday perceptual beliefs in fact commit us to substantive metaphysical theses about composition, their epistemology is mysterious as well. In short, Clarke-Doane is presupposing *both* that our questions about the composition conditions of macroscopic objects are substantive metaphysical questions *and* that it's a trivial matter to answer those questions by appeal to our everyday perceptual beliefs. But that's absurd. I can't refute the metaphysical thesis of mereological nihilism simply by insisting that I'm looking at a piece of paper, though it would certainly be convenient if I could.

[31]Notice that merely adopting some plenitudinous form of platonism (see, e.g., Balaguer 1998) isn't going to be of any help here. We can grant, at least for the sake of argument, that on this picture, our belief that (say) the continuum hypothesis is true is responsive to its corresponding fact, since the nearest worlds where that hypothesis isn't true are worlds where we've chosen to talk about a

a picture on which such objects aren't both mind-independent and causally inert. Nothing less radical than this will suffice.[32]

References

Alston, W. P. 1988. An internalist externalism. *Synthese*, *74*, 265–283.

Balaguer, M. 1998. *Platonism and anti-platonism in mathematics*. Oxford University Press.

Baras, D. 2017. Our reliability is in principle explainable. *Episteme*, *14*, 197–211.

Beddor, B. 2021. Reasons for reliabilism. In J. Brown and M. Simion (Eds.), *Reasons, justification, and defeat* (pp. 146–176). Oxford University Press.

Bergmann, M. 2006. *Justification without awareness: A defense of epistemic externalism*. Oxford University Press.

Berry, S. E. 2020. Coincidence avoidance and formulating the access problem. *Canadian Journal of Philosophy*, *50*, 687–701.

Bhogal, H. 2022. What's the coincidence in debunking? *Philosophy and Phenomenological Research*. Advance online publication. doi: 10.1111/phpr.12907

Boghossian, P. A. 1996. Analyticity reconsidered. *Noûs*, *30*, 360–391.

Boghossian, P. A. 1997. Analyticity. In B. Hale and C. Wright (Eds.), *A companion to the philosophy of language* (pp. 331–368). Blackwell.

Boghossian, P. A. 2000. Knowledge of logic. In P. Boghossian and C. Peacocke (Eds.), *New essays on the a priori* (pp. 229–254). Oxford University Press.

Boghossian, P. A. 2003a. Blind reasoning. *Proceedings of the Aristotelian Society, Supplementary Volume*, *77*, 225–248.

Boghossian, P. A. 2003b. Epistemic analyticity: A defense. *Grazer Philosophische Studien*, *66*, 15–35.

Burgess, J. P., and Rosen, G. 1997. *A subject with no object: Strategies for nominalistic interpretation of mathematics*. Oxford University Press.

Cameron, R. P. 2010. How to have a radically minimal ontology. *Philosophical Studies*, *151*, 249–264.

Chalmers, D. J. 2009. Ontological anti-realism. In D. Chalmers, D. Manley, and R. Wasserman (Eds.), *Metametaphysics: New essays on the foundations of ontology* (pp. 77–129). Oxford University Press.

Chalmers, D. J. 2011. The nature of epistemic space. In A. Egan and B. Weatherson (Eds.), *Epistemic modality* (pp. 60–107). Oxford University Press.

mathematical universe in which it isn't true. But there remains a problem: adopting plenitudinous platonism also involves having beliefs about the multiverse itself—e.g., the belief that it contains a universe corresponding to any consistent mathematical theory—and these beliefs aren't in any sense responsive to the relevant facts.

Chalmers, D. J. 2012. *Constructing the world.* Oxford University Press.

Clarke-Doane, J. 2016a. Debunking and dispensability. In U. D. Leibowitz and N. Sinclair (Eds.), *Explanation in ethics and mathematics: Debunking and dispensability* (pp. 23–36). Oxford University Press.

Clarke-Doane, J. 2016b. What is the Benacerraf Problem? In F. Pataut (Ed.), *Truth, objects, infinity: New perspectives on the philosophy of Paul Benacerraf* (pp. 17–43). Springer.

Clarke-Doane, J. 2020a. *Morality and mathematics.* Oxford University Press.

Clarke-Doane, J. 2020b. Set-theoretic pluralism and the Benacerraf problem. *Philosophical Studies, 177,* 2013–2030.

Clarke-Doane, J. 2022. *Mathematics and metaphilosophy.* Cambridge University Press.

Cohen, S. 2002. Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research, 65,* 309–329.

Collin, J. H. 2018. Towards an account of epistemic luck for necessary truths. *Acta Analytica, 33,* 483–504.

Constantin, J. 2020. Replacement and reasoning: A reliabilist account of epistemic defeat. *Synthese, 197,* 3437–3457.

Cross, T. 2010. Skeptical success. In T. S. Gendler and J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 3, pp. 35–62). Oxford University Press.

Donaldson, T. M. E. 2014. If there were no numbers, what would you think? *Thought, 3,* 283–287.

Dorr, C. 2005. What we disagree about when we disagree about ontology. In M. E. Kalderon (Ed.), *Fictionalism in metaphysics* (pp. 234–286). Oxford University Press.

Ebbs, G. 2022. Putnam on trans-theoretical terms and contextual apriority. In J. Conant and S. Chakraborty (Eds.), *Engaging putnam* (pp. 131–156). De Gruyter.

Elga, A. 2008. *Lucky to be rational.* Paper presented at Bellingham Summer Philosophy Conference.

Enoch, D. 2010. The epistemological challenge to metanormative realism: How best to understand it, and how to cope with it. *Philosophical Studies, 148,* 413–438.

Enoch, D., and Schechter, J. 2008. How are basic belief-forming methods justified? *Philosophy and Phenomenological Research, 76,* 547–579.

Field, H. H. 1989. *Realism, mathematics and modality.* Basil Blackwell.

Field, H. H. 2000. Apriority as an evaluative notion. In P. Boghossian and C. Peacocke (Eds.), *New essays on the a priori* (pp. 117–149). Oxford University Press.

Goldman, A. I. 1979. What is justified belief? In G. S. Pappas (Ed.), *Justification and knowledge: New studies in epistemology* (pp. 1–23). D. Reidel.

Goldman, A. I. 1986. *Epistemology and cognition.* Harvard University Press.

Greco, J. 1999. Agent reliabilism. *Philosophical Perspectives, 13,* 273–296.

Grundmann, T. 2009. Reliabilism and the problem of defeaters. *Grazer Philosophische Studien, 79,* 65–76.

Hale, B. 1994. Is platonism epistemologically bankrupt? *Philosophical Review, 103,* 299–325.

Harman, G. 2003. Skepticism and foundations. In S. Luper (Ed.), *The skeptics: Contemporary essays* (pp. 1–11). Ashgate Publishing.

Horwich, P. 2006. A world without isms: Life after realism, fictionalism, non-cognitivism, relativism, reductionism, revisionism, and so on. In P. Greenough and M. P. Lynch (Eds.), *Truth and realism* (pp. 188–202). Oxford University Press.

Korman, D. Z. 2014. Debunking perceptual beliefs about ordinary objects. *Philosophers' Imprint*, *14*, 1–21.

Korman, D. Z. 2015. *Objects: Nothing out of the ordinary*. Oxford University Press.

Kornblith, H. 2004. Does reliabilism make knowledge merely conditional? *Philosophical Issues*, *14*, 185–200.

Kripke, S. A. 2011. Nozick on knowledge. In *Collected papers: Vol. 1. Philosophical troubles* (pp. 162–224). Oxford University Press.

Lackey, J. 2006. Review of Duncan Pritchard, *Epistemic luck*. *Philosophical Quarterly*, *56*, 284–289.

Leite, A. 2006. Review of Duncan Pritchard, *Epistemic luck*. *Notre Dame Philosophical Reviews*, *2006.04.20*.

Lewis, D. 1986. *On the plurality of worlds*. Basil Blackwell.

Liggins, D. 2010. Epistemological objections to platonism. *Philosophy Compass*, *5*, 67–77.

Linnebo, Ø. 2006. Epistemological challenges to mathematical platonism. *Philosophical Studies*, *129*, 545–574.

Luper, S. 2012. False negatives. In K. Becker and T. Black (Eds.), *The sensitivity principle in epistemology* (pp. 207–226). Cambridge University Press.

Miščević, N. 2007. Armchair luck: Apriority, intellection and epistemic luck. *Acta Analytica*, *38*, 48–73.

Nozick, R. 1981. *Philosophical explanations*. Harvard University Press.

Nyseth, F. 2021. Linguistic conventionalism and the Truth-Contrast Thesis. *Philosophical Quarterly*, *71*, 264–285.

Pritchard, D. 2005. *Epistemic luck*. Oxford University Press.

Pritchard, D. 2007. Anti-luck epistemology. *Synthese*, *158*, 277–297.

Pritchard, D. 2009. Safety-based epistemology: Whither now? *Journal of Philosophical Research*, *34*, 33–45.

Pust, J. 2004. On explaining knowledge of necessity. *Dialectica*, *58*, 71–87.

Putnam, H. 1983. "two dogmas" revisited. In *Realism and reason: Philosophical papers, volume 3* (pp. 87–97). Cambridge University Press. (Reprinted from *Contemporary aspects of philosophy*, pp. 202–213, by G. Ryle, Ed., 1976, Oriel Press)

Roush, S. 2005. *Tracking truth: Knowledge, evidence, and science*. Oxford University Press.

Schechter, J. 2010. The reliability challenge and the epistemology of logic. *Philosophical Perspectives*, *24*, 437–464.

Schechter, J. 2013. Could evolution explain our reliability about logic? In T. S. Gendler and J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 4, pp. 214–239). Oxford University Press.

Schechter, J. 2018. Is there a reliability challenge for logic? *Philosophical Issues*, *28*, 325–347.

Schechter, J., and Enoch, D. 2006. Meaning and justification: The case of Modus Ponens. *Noûs*, *40*, 687–715.

Schoenfield, M. 2014. Permission to believe: Why permissivism is true and what it tells us about irrelevant influences on belief. *Noûs*, *48*, 193–218.

Setiya, K. 2012. *Knowing right from wrong*. Oxford University Press.

Sider, T. 2013. Against parthood. In K. Bennett and D. W. Zimmerman (Eds.), *Oxford studies in metaphysics* (Vol. 8, pp. 236–293). Oxford University Press.

Skipper, M. 2023. Can arbitrary beliefs be rational? *Episteme*, *20*, 377–392.

Sosa, E. 1991. *Knowledge in perspective: Selected essays in epistemology*. Cambridge University Press.

Sosa, E. 1999. How must knowledge be modally related to what is known? *Philosophical Topics*, *26*, 373–384.

Thomasson, A. L. 2007. *Ordinary objects*. Oxford University Press.

Thomasson, A. L. 2015. *Ontology made easy*. Oxford University Press.

Topey, B. 2021. Realism, reliability, and epistemic possibility: On modally interpreting the Benacerraf-Field challenge. *Synthese*, *199*, 4415–4436.

Topey, B. 2022. Saving sensitivity. *Philosophical Quarterly*, *72*, 177–196.

van Inwagen, P. 1990. *Material beings*. Cornell University Press.

Wallbridge, K. 2018. Sensitivity, induction, and miracles. *Australasian Journal of Philosophy*, *96*, 118–126.

White, R. 2010. You just believe that because…. *Philosophical Perspectives*, *24*, 573–615.

Williamson, T. 2000. *Knowledge and its limits*. Oxford University Press.

Williamson, T. 2009. Replies to critics. In P. Greenough and D. Pritchard (Eds.), *Williamson on knowledge* (pp. 279–384). Oxford University Press.

Wright, C. 2004a. Intuition, entitlement and the epistemology of logical laws. *Dialectica*, *58*, 155–175.

Wright, C. 2004b. Warrant for nothing (and foundations for free)? *Proceedings of the Aristotelian Society, Supplementary Volume*, *78*, 167–212.

Wright, C. 2014. On epistemic entitlement (II): Welfare state epistemology. In D. Dodd and E. Zardini (Eds.), *Skepticism and perceptual justification* (pp. 213–247). Oxford University Press.

Wright, C. 2016. Abstraction and epistemic entitlement: On the epistemological status of Hume's Principle. In P. A. Ebert and M. Rossberg (Eds.), *Abstractionism: Essays in philosophy of mathematics* (pp. 161–185). Oxford University Press.