

LUCK AND CREDIT IN THE SPACE OF REASONS

J. Adam Carter
Stirling University

ABSTRACT: This essay will advance the view that the McDowellian theory of knowledge fails to satisfy the requirements of an adequate anti-luck epistemology. Section 1 presents two twin anti-luck desiderata that, I shall argue, an account must accommodate: (i) If S knows p , then S could not have easily been wrong that p ; (ii) If S knows p , then S is credit-worthy for her true belief that p . In Section 2, I outline the salient differences between McDowell's iconoclastic anti-luck strategy and traditional strategies. Section 3 offers reasons for thinking that McDowell fails to satisfy what I have presented as the first anti-luck desideratum; section 4 offers reasons for thinking that McDowell fails the second desideratum.

1. *Epistemic Luck: Two Twin Desiderata*

I shall begin by clarifying what has, particularly in the second half of the past century, been taken as a platitude in epistemology: *Knowledge can be undermined by luck.*¹ This is not to be confused with the stronger claim that luck *always* undermines knowledge; rather, the idea is that for some agent S and proposition p , if S believes p , and it is by luck that p is true, then S does not know p . It is in this specific respect that knowledge *cannot* depend on luck, and hence, that knowledge can be undermined by luck.²

To take this platitude seriously, we must make precise the conditions under which some event is lucky. In the past, attempts to define a 'lucky' event have been sparse and unhelpful. Luck has either been taken as a basic concept requiring no recursive definition, or defined in such a way that intuitive features of luck—such as subjective significance—have been ignored. Consequently, luck, when defined, has been conflated in the literature³ with logically distinct concepts such as 'accident' and 'chance.'

The most comprehensive recent account of luck is offered by Duncan Pritchard (2003), who cashes out lucky events modally. On Pritchard's view:

Lucky Event (L1) If an event is lucky, then it is an event that occurs in the actual world but which does not occur in a wide class of the nearest possible worlds where the relevant initial conditions for that event are the same as in the actual world.⁴

This modal definition is akin to a less technical idea, namely: if S knows p , then S couldn't easily have been wrong that p . This idea captures an important part of what it is about Gettier cases that leads us to reject them as cases of genuine knowing. For example, consider Gettier's famous disjunctive 'Brown is in Barcelona' example.⁵ When asked whether Smith 'knows' that Brown is in Barcelona, a salient feature of our denying that he does will be that Smith—even though he is right—could have easily been wrong.

The story, I wish to press, does not end here. The explanation for why Smith doesn't know should not end with the mere pronouncement that he could have easily been wrong. To come to grips with a different explanation, let us consider the following non-modal epistemic platitude:

Credit-worthiness (C1): Knowledge is an achievement, creditable to the agent.⁶

Along with failing the luck platitude (LI), our Gettier protagonist also fails (C1). For example, we could just as easily have said that Smith doesn't know that Brown is in Barcelona because he doesn't deserve *credit* for his belief's being true. *Either* being such that one could have easily been wrong, or being such that one doesn't deserve credit for her belief's being true is sufficient for undermining one's knowledge.

A useful way to assess whether knowledge is undermined via recourse to (C1) is to understand 'credit-worthiness' in terms of what is explanatorily salient in an individual's belief formation. This idea has been advanced in recent work by John Greco (2006a)⁷, who assesses agents' credit-worthiness by considering whether the agent's intellectually virtuous character is an explanatorily salient feature of the causal explanation for why the agent has formed the belief. Whether the agent's character is a sufficiently salient part of the causal explanation will itself depend, on Greco's view, on the particular context of inquiry.

A helpful way to apply Greco's view is to consider Russell's famous 'grandfather clock' Gettier case. As Russell's story goes, a man walks downstairs and forms the belief 'It is 11:00' after consulting a grandfather clock that has always been reliable in the past. Unbeknownst to the man, the clock was broken and the hands were stuck on 11:00. Fortunately, the clock broke exactly 12 hours before he consulted it, and so his belief that it is 11:00 is true.

While it is true to say both that our protagonists' intellectual character (i.e. leading to his having formed his belief on reliable evidence, etc.) *and* a fortunate favor from the world (i.e. that the clock happened to be stuck on the correct time) are part of the causal explanation for his having formed a true belief that it is 11:00, the favor from the world—and not the agent's virtuous weighing of evidence—will, intuitively, be the more salient feature of our explanation. Our 'Gettiered' protagonist does not know that it is 11:00, and a reason for this is that the lack of explanatory salience of his agency in his reaching a true belief precludes him from being credit-worthy.

Of course, we could also point out that in the Russell case, the clock consulter could have easily been wrong. Given that this is so, and given that a failure of the credit-worthiness condition doesn't necessarily imply a failure of the modal luck condition, and vice versa, we reach the following conclusion: (i) meeting each condition is necessary for knowing, and thus, failing either is sufficient for undermining knowledge. In the Russell case, neither necessary condition is met. In other cases in which knowledge is undermined, one condition might be met while the other is not met.⁸ Whether the anti-luck desiderata are such that their joint satisfaction constitutes a mutually sufficient condition for knowing is a question I shall not explore here.⁹

I offer that we take from our discussion to this point the following:

1. An adequate theory of knowledge must preserve the platitude that luck can undermine knowledge.
2. When luck undermines knowledge, at least one of the following two twin desiderata will not be met: (i) the agent could not have easily been wrong; (ii) the agent must be credit worthy for holding a true belief.

For the remainder of this paper, I will refer to (i) and (ii) as the ‘twin desiderata’ that must be satisfied by a knower. Any account of knowledge that rules as knowledge a belief that fails either anti-luck desiderata will fail to do justice to what is intuitive about the relationship between knowledge and luck that motivates the platitude in (1).

2. *Factive Reasons: McDowell’s Departure from the ‘Hybrid Conception’*

Theoretical anti-luck machinery crafted in the face of these desiderata are myriad; one particular strategy is to address our first desideratum through ‘low-standards’ fallibilism. The following could serve as an anti-luck-motivated template argument for fallibilism.

Luck-‘assuaging’-argument for fallibilism

1. If knowledge requires that S’s evidence for p entails p , then we know (basic everyday) propositions only if our evidence for these propositions entails their truth.
2. For any proposition p , evidence for p will not entail p ; also necessary will be a favor from the world—namely, that the world is, in fact, the way that S takes it to be.
3. We do know a great deal of basic, everyday propositions.
4. Therefore, knowledge does not require that S’s evidence for p entails p .

This sort of argument is an attempted reconciliation between two apparently conflicting ideas: first, that subjective evidence that the world is thus and so is insufficient for entailing that the world is thus and so¹⁰; secondly is the neo-Moorean idea that, in light of this first idea, it is nonetheless counterintuitive to deny that we know a great deal of what we take ourselves to know. These ideas are captured in premises 2 and 3.

Stewart Cohen has defended a brand of fallibilism along these lines¹¹ (albeit, a fallibilism that is applicable only in certain low-stakes, everyday contexts)¹². At the crux of Cohen’s contextualist-based fallibilism, and other versions, is P2, the motivation for which is on brute display in lottery cases. To borrow an example from John Hawthorne (2004), suppose that you are wagering on a roulette wheel that contains (to exploit the idea) one trillion red spaces and one white space, each of uniform size. You form the belief that the ball will land on red, and are correct. Your evidence for this belief is clearly as good, if not much better, than your evidence that, for example, your car is currently parked where you last left it. However, it strikes us as highly counterintuitive to say that you ‘knew’ the ball would land on red; as John McDowell points out, “for all you knew” the ball might have landed on white.¹³

Given that our evidence in the lottery case is (or can be, given that we can increase the wheel to be near-infinitesimally salted with red spaces) better than it will be for our everyday propositions—

such as that our car is where we have parked it—it seems that we must settle for an inevitably severed connection between justifying evidence and the world. That is to say, so long as there is a ‘window’ open for fortune to render the world as other than how our evidence suggests it is, then that window will be invariably large enough to slip in a Gettier case.

John McDowell (1998b) limns an account of knowledge constructed to, as it were, ‘seal this window’ before it has a chance to have opened; his account, among other things, attempts to do justice to both our anti-luck desiderata while rejecting the temptation to accept the sort of fallibilism that arises from, for example, the lottery-type cases that motivate P2.

Part of what makes McDowell’s account so interesting, and iconoclastic, is that he distances himself from an idea that has been central to our discussion of luck thus far, which is that knowledge of some proposition p seems to require, over and above our evidence that justifies us in believing p , some favor from the world. A feature of McDowell’s view is that, if S knows that p , then S couldn’t be mistaken that p . (If knowledge that p requires a favor from the world—no matter how small a favor—then it follows that independent of how strong your evidence for p is, you could be mistaken that p just in case the ‘favor’ is not granted.) This mistake-precluding infallibilist feature has proved dauntingly difficult to defend in the literature for the reason that it opens itself up to the charge that it will foment a Cartesian-like skepticism. For example, consider the following line of argument:

1. If knowledge requires that the knower couldn’t be mistaken, then all known (contingent) propositions are such that those who know them could not be wrong about them.
2. If it is possible that we are BIVs, then it is possible that we are wrong about all (contingent) propositions we take ourselves to know.
3. It is possible that we are BIVs.
4. Therefore, if knowledge requires that the knower couldn’t be mistaken, then there are no known (contingent) propositions (From 1, 2 and 3)

Such an argument need not scare McDowell off, however, regardless of what aversions he might have toward the skeptical conclusion that meretriciously flows from his endorsement of P1. McDowell is able to maintain P1 whilst not courting the obvious skeptical worries associated with its endorsement because he rejects the very framework within which these skeptical worries become a problem for views averring P1. That framework is, as he describes it, the *hybrid* conception. The hybrid conception (HC), to which many epistemologists have wedded themselves dogmatically, amounts to the following idea:

Hybrid conception (HC): For agents S and T, S and T can be alike “in the respect of their satisfactoriness of their [respective standings] in the space of reasons, although only one of them is a knower.¹⁴”

An agent’s ‘standing in the space of reasons’ is a locution McDowell borrows from Wilfred Sellars. An agent’s standing in the space of reasons will be a function of what reasons the agent takes for her belief. The Sellarsian idea, which McDowell defends, is that knowing amounts to a certain satisfactory standing in the space of reasons. Crucial to McDowell’s view is that *facts* can be taken as reasons: for

example, if I know that there is a chaffinch in front of me, my satisfactory standing in the space of reasons could be comprised by my seeing *that* there is a chaffinch before me—my reason for believing it to be so. Facts, for McDowell, shape the space of reasons ‘as I find it’, and so what reasons will be available to me will depend in part on how the world happens to be. It follows that I might well have facts available to me to use as reasons that someone having my identical experiences does not have.

To clarify this idea, suppose that my recently envatted duplicate has the experience of seeing what appears to be a chaffinch before him; suppose further that my envatted duplicate’s visual experience is qualitatively identical to mine. On McDowell’s view, I can be a knower whilst my envatted duplicate is not; this is because there are facts available for me as reasons which are not available to my envatted duplicate—namely, the fact that I see *that* a chaffinch is before me. Because an agent’s standing in the space of reasons is determined by what she takes as reasons¹⁵, my duplicate and I do not occupy identical standings in the space of reasons; thus, I can maintain a satisfactory standing whilst my duplicate will have a defective (even if blameless) standing.

McDowell’s take on the chaffinch scenario allows us to see wherein lies his departure from the view characterized as HC. HC will rule that my envatted duplicate and I are alike in our standings in the space of reasons, whilst only I am a knower. HC is motivated, in part, by the thought that our standing in the space of reasons must be alike because, for both of us, what is reflectively accessible to us is identical: namely, that there appears to be a chaffinch in front of us. HC endorser will allow that my counterpart and I both take as reasons merely that there appears to be a chaffinch in front of us. Fortunately for me, I have received the needed favor from the world—namely the favor that I am not being fed hallucinatory experiences—and *this* is what distinguishes me from my duplicate as a knower.

McDowell boldly rejects the thought that it boils down to this ‘favor’, or bit of lucky world cooperation, that is the demarcating feature between a knower and a non-knower. He is frankly unsympathetic with the thought that such a demarcating feature should be one that is out of the control of the agent or, as he says, out of the agent’s rational powers. Here’s McDowell:

And if [the demarcating feature] is outside the reach of her rational powers, how can its being so be the crucial element in an intelligible conception of knowing that it is so—what makes the relevant difference between her and the other subject [in the hallucinatory case]? Its being so is conceived as external to the only thing that is supposed to be epistemologically significant about the knower herself, her satisfactory standing in the space of reasons.¹⁶

The HC view, which commits this sin, is taken to be a *hybrid* conception by McDowell because it understands knowing as requiring two parts: one internal component, namely, the agent’s standing in the space of reasons (comprised of what evidence is available to the agent), and secondly, an external component, namely that the world is thus and so when we take it to be thus and so.

By allowing reasons to be factive, McDowell smuggles what is taken by HC endorsers as external to the space of reasons *into* the space of reasons. Making this move allows him to preserve the Sellarsian idea that knowledge is certain standing in the space of reasons whilst not committing the foul of *interiorizing* the space of reasons—the foul committed by any view that severs the connection between the agent’s standing in the space of reasons, and the world.

This severed connection, allowed by the HC view, will leave it inevitable that what one takes as a reason for a belief p will not entail p . Consequently, the HC view—via requiring the external component over and above the agent’s standing—for knowledge, will allow that we (for example) know that what sits in front of us is a chaffinch only if something beyond our rational powers, namely that the chaffinch is in fact there, happens to coalesce with the evidence we have that this is so. Because the view McDowell pronounces doesn’t require this extra external component (as it would already be built into the agent’s factive standing) it appears on the surface that McDowell has developed a view that avoids the painstaking difficulties previously discussed and which beguile views attempting to preserve the twin epistemic luck desiderata. *Prima facie*, (or, as McDowell would have it) the satisfaction of these desiderata engenders quandaries such as the one Zagzebski noted, as a result of prior commitments to HC. And so, as the story would go, rejecting the HC framework offers an anti-luck alternative to HC that preserves our desiderata without needing to laboriously craft (for example) Gettier-assuaging conditions and playing, as Guy Axtell has called it, the Gettier game.

Such confidence, though, is premature. Although rejecting the HC framework does, in fact, eliminate one feature of an account that facilitates the infiltration of luck, we should not be so quick to conclude that the rejection of HC serves to eliminate all methods by which an account can fail one of our two anti-luck desiderata. Put another way, adopting HC is the cause of ‘some’ anti-luck troubles, but not all. Getting rid of it does not vouchsafe a theory from the ills of luck altogether.

In the next section, I advance an argument for the conclusion that McDowell’s view cannot preserve the first anti-luck intuition: namely, *that If S knows p, then S could not have easily been wrong that p.*

3. *McDowell and Safety: The First Desideratum*

I shall begin this section with a story, which I intend to serve two functions. First, it will provide a scenario within which one could argue that, within the McDowellian framework, if S knows p , then p must be ‘safe.’ I will formalize the argument to this conclusion, and then demonstrate that it has a false premise. I shall then make clear just why the falsity of the premise amounts to what I think is an intractable failure, for McDowell’s view, toward the prospective end of satisfying the first anti-luck desideratum.

To borrow as a starting point for our story a modified example from Alan Millar¹⁷, suppose you are a seasoned horticulturalist and find yourself standing before a large hill, in Augusta, Georgia, covered with azaleas. Your friend, Mr. Vat, an equally competent horticulturalist, is at that moment looking at a strikingly similar hill of azaleas, which sits miles away, in Atlanta.

Unbeknownst to either of you, a lunatic had been on the loose in the area; the previous night, the lunatic planted one ‘silk’ azalea on the Atlanta hill, and one ‘silk’ azalea on the Augusta hill. The silk azaleas are visually indistinguishable from actual azaleas, even to the most trained eye.

Coincidentally, at about the same time, both you and Mr. Vat pick one ‘azalea’ on your respective hills. Mr. Vat, unluckily, swipes up his hill’s lone silk-pedaled fraud. You, on the other hand, get lucky. You reach for the lone fraud on your hill, but at the last second, grab a genuine azalea. You each, at the same time, form the belief “What I hold is an azalea.”

Mr. Vat’s does not know he has an azalea in his hand because he, in fact, does not have one. McDowell’s view allows that you, on the other hand, know you have an azalea.

Your reason for believing that you hold an azalea is that you see *that* you are holding one. This reason is not available to Mr. Vat, even though the hill he is viewing is relevantly similar to yours, and even though he is equally as epistemically blameless a belief former as you, given how things, for both of you, appear.

I think the appropriate response to have at this point is that luck distinguishes, on McDowell’s view, the knower from the non-knower in a manner relevantly analogous to the way luck distinguishes a knower and a non-knower on the HC account McDowell criticizes.

On the HC view, what was ‘lucky’ was the fact that the world happens to be the same way that our evidence leads us to take it to be. McDowell’s view does vitiate this particular insidious feature of luck. But he does so only to open the door to another. Consider that, on McDowell’s view, what distinguishes you from Mr. Vat is that you are fortunate enough to be situated in the world in which you are able to take the fact that you see *that* you are holding an azalea as a reason for believing that this is so. Mr. Vat doesn’t have this luxury. We can suppose further, and I think importantly, that if Mr. Vat had been standing on your hill of azaleas in Augusta and you had been standing on his hill of azaleas in Atlanta, then you would pick an fraud and hold a false belief and he would pick an azalea and hold a true belief.¹⁸

Against this background, it is not clear how McDowell is going to preserve the first anti-luck desideratum: namely, that knowers couldn’t easily have been wrong. For, surely, you could just as easily have been wrong as Mr. Vat! McDowell appears to allow that what is epistemically significant between you and Mr. Vat is a lucky and an unlucky swipe, respectively, a turn of fortune/misfortune that could just as easily have been reversed. I’ll turn now to what I take to be the best case the McDowellian has available toward preserving our first anti-luck desideratum.

McDowell’s First Line of Defence

P1: For any agent S and belief ϕ , if S 's reasons for believing ϕ entail ϕ , then S couldn't have been mistaken that ϕ .

P2: On the McDowellian account, all known propositions are such that agent's reasons for believing them entail that they are true. (Factive reason thesis).

(3): Therefore, on the McDowellian account, all known propositions are such that those who know them couldn't have been mistaken about them (from 1 and 2).

I shall now engage with what I think is most crucially flawed about McDowell's attempt to satiate the first anti-luck desideratum through this sort of argument. Before proceeding, I shall point out that, if the above argument is sound, then McDowell *would* satisfy the first desideratum. The satisfaction of that desideratum would be entailed by the conclusion (3).

P1 of the argument, however, is false. Consider that P1 claims:

[P1] For any agent S and belief ϕ , if S 's reasons for believing ϕ entail ϕ , then S couldn't have been mistaken that ϕ .¹⁹

So long as reasons are factive, then P1 will always be true in the following sense: it is impossible for S to have a reason r for a belief ϕ , such that r entails ϕ , and S could have been mistaken that ϕ . For, if r entails ϕ , then it is logically impossible for:

[A]: $[\Box(r \rightarrow \phi) \wedge \Diamond(r \wedge \neg\phi)]$

However, in this respect, the fact that the agent couldn't be mistaken turns upon the idea that we are granting that the agent is already in a factive state. And so, *once you have adopted a factive reason*, then you couldn't be mistaken. P1 is trivially true on this reading of [P1]; this is because the criterion that "the agent with the factive reason couldn't be wrong about the subsequent belief" is taken to be satisfied on this reading by the analytic truth that [A] is false.

There is a reading of [P1], however, under which [P1] comes out false rather than trivially true. The motivation for the second, and I think more appropriate, reading of P1 comes once we call to scrutiny the following conditional (C) McDowell seems to ignore in his recent discussions of luck²⁰:

[C] If S is in a position such that her reason r for believing ϕ entails ϕ , then it couldn't have easily been that the reason she had for believing ϕ *does not* entail ϕ .

McDowell's 'First Line of Defense' seems to come apart at the seams once we recognize that [C] is false, and that any non-trivially true interpretation of P1 in the argument would depend on [C] being true. For, if [C] is false, then the conditional claim in P1 has a true antecedent and a false consequent; that is, a counterexample can be constructed against P1 such that an agent can have a factive reason r that entails ϕ , *and* be such that she could have easily been mistaken; that is, it could have easily been that case that the reason she had did not entail ϕ .

Our azalea case illuminates the falsity of [C]. 'Your' reason for believing that that you are holding an azalea is the fact that you see that you are holding one. The antecedent of [C] is thus satisfied. However, it *could* have easily been the case that your reason for believing that you are

holding an azalea did not entail this. For example, suppose that, rather than had you at the last second grabbed a genuine azalea, you grabbed the silk-pedaled fake. If this had happened (and even suppose that in most nearby worlds, it *would* have happened, because in most nearby worlds, you do not avert your hand at the last second), then you would not longer have a factive reason available to you with respect to your belief ϕ (that you are holding an azalea). [C], then, is open to counterexample, and *ipso facto*, P1 of the ‘First Line of Defense’ argument fails, and even more importantly, the most obvious way McDowell has available for him to satisfy the first anti-luck desideratum loses all the promise it, against the HC view, initially appeared to have.

Before moving on to my second thesis, which is that McDowell fails the second anti-luck desideratum, I want to bring to surface a potential objection to what I’ve argued for in this section thus far.

McDowell, given what has been presented, only runs in to problems if it is stipulated that his view counts *as knowledge* cases in which an agent could have easily been mistaken. One ‘McDowellian’ line of thought, which might be taken here by Alan Millar, would be to reject that ‘you’ know you are holding an azalea for the following reason: to have perceptual knowledge that you are holding an azalea, it must be the case that you have exercised the appropriate perceptual-recognitional abilities in judging this to be so. Millar writes:

“This ability is not *to judge correctly enough of the time*. It is an ability to tell by looking, and thus come to know, whether or not something is an azalea. It is exercised only if one ends up knowing²¹.”

Millar’s ‘only if’ here is important in that it shows that he requires the exercising of appropriate perceptual-recognitional abilities for knowledge, but does not make the stronger claim that, in all cases in which one exercises these abilities, one ends up with knowledge. Whether one knows will also depend, on Millar’s view, on features of the environment within which inquiry is taking place. For example, if there are a multitude of fake azaleas in the area, then, even if you hold an actual azalea, you would not know that you do; in such an environment, you would not have the ability to tell by looking whether or not something is an azalea. And so, on this line, perceptual-recognitional abilities are indexed to particular environments.²²

It is unclear to me whether a lone fraud azalea on a hill of genuine azaleas would, on Millar’s account, render the environment a ‘bad’ case. Regardless, the possibilities seem to be as follows: If the protagonist in the azalea case ‘knows’ that what is held is an azalea, then, as has been discussed, the first anti-luck desideratum is not met because he could easily have been mistaken. If the protagonist in the azalea case *does not* know that what is held is an azalea, then two problems arise:

- (i) The reason for denying her knowledge could not be applied, without contradiction, to what McDowell has to say elsewhere about testimonial knowledge.
- (ii) The idea that, on McDowell’s framework, one cannot ever ‘know that one knows’ is exploited.

Regarding (i): McDowell, in “Knowledge by Hearsay”, delineates conditions under which testimonial knowledge is possible within his framework. An example he uses to illustrate his view regards a tourist who, in a foreign city, asks a passerby where a historic cathedral is located. The passerby gives directions, which are correct. I will address more carefully in the next section McDowell’s reasoning for claiming that the tourist can, from the passerby’s testimony, know where the cathedral is. It suffices, to make the point here, that McDowell *does* think that knowledge is attributable to the tourist, in light of the fact that the environment of inquiry is such that at least some passersby will give incorrect information. I wish to draw this analogy: some ‘false-speaking’ passersby in the city in which the tourist is inquiring amounts to a relevantly similar environment as the horticulturalist encounters when approaching a hill of a lone fake azalea amidst a crop of genuine ones. And so, it seems that McDowell would be hardpressed to deny knowledge to the protagonist in the azalea case whilst maintaining (without inconsistency) that the tourist can ‘know’ where the cathedral is.

Regarding (ii): The second problem that arises if McDowell denies knowledge to the protagonist in the azalea case is that a direct consequence will be a concession that one can never be in a position to know *that* she knows some proposition that she knows. Consider the following argument:

1. On McDowell’s view, whether one knows some contingent proposition p will depend, in part, on facts about the environment within which inquiry is taking place.
2. To know *that* one knows some contingent proposition p , one must have knowledge of all relevant antecedent facts about the environment in which she is inquiring. (i.e. Are there any ‘hologram’ x’s in the area, etc.)
3. Knowers of contingent propositions rarely (if ever) are in command of all relevant antecedent facts about their environments of inquiry.
4. Therefore, on McDowell’s view, if some agent knows a proposition p , she will rarely (if ever) know *that* she knows this proposition.

The crucial premise in this argument is (3). If we rule that the protagonist in the azalea case doesn’t know, then the reason would be that he doesn’t know a relevant fact about the environment, namely, that there is a lone fraud azalea—one that he couldn’t distinguish by looking from a real azalea. Given that he *can’t* make this distinction, he can’t know whether his environment is such that he can appropriately exercise his perceptual-recognitional abilities. *Ipso facto*, he cannot know if he has satisfied a necessary condition for perceptual knowledge. The extent to which this consequence is problematic to McDowell’s view turns on the extent to which an account ought to preserve the intuitive idea that a great deal, or at least some, of our perceptual knowledge is such that we know when we have it.

The conclusion we can draw from this section is that McDowell either fails the first anti-luck desideratum, or avoids failing it only at the expense of either contradicting what he says about testimonial knowledge or failing to preserve the platitude that we can know, at least sometimes, when we have perceptual knowledge.

4. McDowell and Credit: The Second Anti-Luck Desideratum

We shall return, in this section, to what I described previously as the second anti-luck desideratum.

This states that:

[D2] If S knows p , then S is credit-worthy for her true belief that p .

[D2] will be satisfied by a theory that allows an agent to be credit-worthy for her true belief only in those cases in which it is her agency that is the salient feature of the causal explanation we give for reaching a true belief. As I mentioned in the first section, the idea of ‘credit-worthiness’ with which I use here is the one in use by John Greco’s recent work. Greco summarizes the intuitive idea behind [D2]:

When we attribute knowledge to someone we imply (at least pragmatically) that they are responsible for getting things right. Put another way, we imply that it is to their credit that they have got things right—that their believing the truth can be explained by their own efforts and abilities, rather than by dumb luck, or blind chance, or something else.²³

While Greco’s idea here is indeed intuitive, what might not be intuitive is that it is correct to present [D2] as a criterion over and above the anti-luck criterion at work in the previous section.²⁴ There is an explanation owed to the reader at this point, which is as follows: I have, in the first section, mentioned only in passing that the satisfaction of [D2] would not entail the satisfaction of the modal-luck condition (first desiderata) and vice-versa. I did not, though, provide a defense of this claim, which I think is now due. To be clear, it is due so that the following idea is reinforced: a theory can fail to preserve what is intuitive about the relationship between luck and knowledge by failing *either* or *both* desiderata. The reinforcement of this idea is important to the purpose of this paper, which is to bring the stronger of the two charges against McDowell. To defend the stronger claim, the desiderata must be shown to be clearly logically distinct. I will return to my argument that McDowell fails [D2] after first having vindicated the importance of satisfying [D2] *via* demonstrating it as a clearly distinct anti-luck desideratum.

This distinction, which I have defended elsewhere²⁵, rests on the assumption that one could hold a safe belief (and thus, satiate the first desideratum) and not be credit-worthy for holding this belief.

Before delineating a particular case in which the two ideas in question could, as it were, ‘come apart’, I want to note that, although I have been referring to a ‘safe’ belief and a belief that satisfies the first desideratum (i.e. it couldn’t have easily been wrong) interchangeably in this essay, I should point out that this has only been for convenience. The idea that ‘if S knows p , then S couldn’t easily have been wrong that p ’ is a less technical, and generalized version of the *safety principle*. I’ve described the safety principle in the less technical language, to this point, only to avoid tediousness.

Now, with the goal at hand of demonstrating that D2 is separate from our modal luck condition, I will present the latter in ‘all its glory.’

The safety principle

[S1] For all agents, ϕ , if an agent knows a contingent proposition ϕ , then, in most nearby possible worlds, that agent only believes that ϕ when ϕ is true.²⁶

I’ll offer some reasons which, taken together, will offer a presumptive case for thinking that the meeting [S1] does not entail a satisfaction of [D2]. Firstly, I will offer a counterexample to the idea that, if $S(x)$ represents the satisfaction of x , then $S[S1] \rightarrow S [D2]$. The counterexample turns on the distinction between *animal knowledge* and *reflective knowledge*. Ernest Sosa²⁷ distinguishes between the two by pointing out that animal knowledge, though ‘apt’, is not *defensively* apt; the agent will not be able to provide good reasons for why she knows what she does. Animal knowledge, which is usually attributed to animals and small children, is such that the possessor could satisfy the external condition in [S1], the satisfaction of which would not be undermined by the possessor’s inability to give good reasons. It seems dubious to think, though, that satisfaction of [S1] in cases of animal knowledge would be such that it is the *agency* of the knowledge possessor that is most explanatorily salient in explaining why she reaches the beliefs she does. An equally compelling idea for why such agents know what they do is that they have been—luckily for them—wired to form true beliefs rather than false ones. It seems to be, at minimum, an open question as to whether such knowledge constitutes an achievement on the part of the agent.

A second reason for doubting that $S[S1] \rightarrow S [D2]$ turns on the idea that epistemology is a *normative discipline*. As Greco thinks, “This claim is hardly controversial—almost everyone in the field would accept it²⁸.” A consequence of Greco’s premise is that knowledge is not analyzable without recourse to normative evaluations of agents. It is Greco’s thought that a virtue-theoretic approach to epistemology is the most promising strategy for normative evaluations. Without going into unnecessary detail, or making commitments to the virtue-theoretic project, the idea I want to suggest is: If Greco is correct, then a defense of $S[S1] \rightarrow S [D2]$ would require that we commit the ‘illegal’ move of deriving an evaluative conclusion from a factual premise. That is, we would derive a normative claim from a descriptive fact about the satisfaction of [S1].

I offer, then, that these two independent reasons for denying that $S[S1] \rightarrow S [D2]$ provide a presumptive case against it. And so, in returning now to the original task of assessing McDowell’s theory with respect to [D2], we can see that the failure to satisfy [D2] provides an additional problem to his view, over and above the claim in the previous section that he fails the modal-luck condition.

Against this background, let us consider what McDowell has to say about the persistence of knowledge over time. In “Knowledge by Hearsay.” To borrow from his example, suppose you are watching a television broadcast in which you see Blair giving a speech as prime minister. McDowell thinks that you can *know* that Blair is prime minister (i.e. you are seeing *that* he is prime minister) and

that your knowledge does not ‘disappear’ once your television screen goes black. McDowell thinks, therefore, that *remembering that* something is the case constitutes a factive state. You can cease to know that Blair is prime minister, after remembering that he is, for one of two reasons. First, Blair could fail to be prime minister (for example, he could be assassinated immediately after you turn off your television). Secondly, you could fail to know that he is prime minister if you believe him to be such (even if he remains to be) whilst neglecting to confirm that he is in doxastically responsible intervals. For example, if he is prime minister five minutes after you turn off your TV and see that he is, then you can meet the demands of doxastic responsibility without requiring another immediate confirmation interval. However, if you secluded yourself from confirmation possibilities for a year and continued to believe he was prime minister, then even if he were, you would fail to know. Here’s McDowell:

Suppose one has become informed of some impermanent but durable state of affairs, and goes on taking it to obtain after one’s original epistemic access to it has lapsed. I do not claim that if the state of affairs still obtains then, come what may, one’s continued taking it to be amounts to knowledge. On the contrary: one’s status as a knower is undermined, even if things still are as one takes them to be, if one’s taking things to be that way is, as I put it, doxastically irresponsible. It is doxastically irresponsible to go on taking it that some state of affairs of the right kind still obtains if the interval since one’s last confirmation is too long, or more generally if the intervals between confirmations are too long, say if one has missed the news for an excessive period.²⁹

I am sympathetic to the idea that knowers must be doxastically responsible, and that this will require some degree of continued confirmation in cases of testimonial knowledge. I think, however, that more careful scrutiny of the view will uncover consequences which render his view ill-pressed to preserve [D2].

Suppose I bury a silver coin underneath a tree in my back yard. I see that the coin is there, and suppose further, that the next day, I remember that the coin is there. Suppose I never again check again under the rock to see if it is there. On McDowell’s view, there is some point in time t_n at which I no longer know that the coin is there, even if the coin is still there. Suppose that, seconds before t_n I form the belief “the coin is there.” Seconds later, and infinitesimally after t_n , I form the belief “the coin is there.” McDowell seems to rule that I know in the first case, but not in the second. If McDowell is to preserve the intuitive idea that knowers are credit worthy, then he must allow that my belief seconds before t_n constitutes a creditworthy achievement, but not the belief I form seconds after t_n . This consequence seems rather hard to swallow.

My reason for not swallowing it is that it is at odds with the idea that if creditworthiness is attributable to an agent for her belief at time t , then it will also be attributable to the agent at t_n if t_n is indistinguishable to the agent from t . McDowell, then, seems to run up against the same sort of ‘paradox of the heap’ problem that Timothy Williamson articulates in his anti-luminosity argument. If the agent, and not luck, is explanatorily salient in the causal explanation of her forming a true belief at t , then it seems bizarre to deny that her agency would be explanatorily salient in the causal explanation

of her forming a true belief at some time m individuated so closely to t that the agent cannot distinguish between the two.

McDowell notes that “the topic of doxastic responsibility is clearly complex³⁰.” He describes it metaphorically, in fact: confirmation intervals nourish our knowledge, he says, like intervals of water nourish plants. Unlike with plants, though, it is not clear how regularly confirmation intervals are required to preserve knowledge of what we remember. If McDowell attempted to patch up this problem by supposing we could be in factive states with regard to the extent to which confirmation intervals are demanded by doxastic responsibility, then we are led to a regress problem, as doxastic responsibility is then required for the upkeep of this ‘meta’ knowledge about the requirements of doxastic responsibility.

I wish to mention now a separate sort of problem for McDowell’s prospects of preserving [D2]. Continuing with the ‘coin under the tree’ example: suppose that you and I hide silver coins under separate trees. Unlike in the last case, however, we *do* meet whatever should be said to be required by doxastic responsibility in our confirming these beliefs; we check them every day, and confirm that they are still there. Suppose that yours is lifted by a thief hours after you check it one day, and mine is not. Later that day, we both form the belief “My coin is under the (respective) tree.” McDowell must allow that I am a knower (for I remember that it is there, and I am not irresponsibly in continuing to believe what I remember) and you are not. Is what separates us, as he says in “Knowledge and the Internal” *epistemically significant*? What separates us is certainly not epistemically significant in any sense relevant to any evaluative assessment of the salience of our agency that would be, itself, more important than luck in causal explanations of our forming our beliefs.

Even more dangerous an idea: Suppose that we hold this story the same, but suppose the thief in the night felt guilty and minutes later, re-buried your coin. His verdict must be that I can remember that the coin is still there. If he rules that you can as well, then his view immediately fails our modal luck condition. If he rules that you do not know, then the demands of doxastic responsibility on your part (demands relevant to whether you know) are not demands you could ever know you have met when you have. And so, McDowell could preserve [D2] only if what separates you and I as creditworthy is the action of a thief in the night.

I want to confess, at this point, the following: My commitment to the importance of satisfying [D2] also commits me to the idea that we must define central epistemic concepts by recourse to normative language. The requisite use of such language separates epistemology as a normative discipline from a naturalized discipline, and as a result, into a discipline that must be wed to relatively ‘thinner’ concepts. The conditions under which thin concepts are satisfied on a view will be murkier than others. And so, I do not wish to criticize McDowell in a way in which I say he first must meet the

demands of 'credit-worthiness' apposite to normative epistemology, and then charge him for failing to meet demands that it is unreasonable to think such a project should meet.

Rather, I take it as a commitment to the idea that the project of explaining when one knows requires evaluation of agents that such evaluations must be conducted in a principled way. McDowell's view not only appears to fail [D2], but his position does not give a principled idea for distinguishing conditions under which agents are creditable for what they know.

5. Conclusion

I have argued in this paper that preserving what we want to say about the relationship between knowledge and luck requires the satisfaction of two anti-luck desiderata, and that McDowell's theory of knowledge, though an interesting and I think promising alternative, doesn't adequately meet them. McDowell's attempt to eliminate favors from the world as epistemically significant to one's standing in the space of reasons shows that he takes challenge of luck seriously. His account, I hope to have shown, does not go far enough in its attempt to meet this challenge.

¹ This idea has been expressed by Duncan Pritchard (2005b) in the terminology that knowledge “excludes” luck.

² Beliefs can be true, and the fact that they are can depend on luck, and yet, knowledge is not undermined. In such cases, it is by luck that the agent formed a true belief, however, it will not be by luck in such cases that the agent’s belief is true. For example, suppose a computer glitch in the registrar’s office results in your being in rolled in European History class rather than American History. Suppose further that you learn five facts about Robespierre’s involvement in the French Revolution. Although your having these true beliefs depends on luck—(i.e. on the chancy computer glitch)—it is not by luck that the beliefs you formed are true. I mention this to wax the idea that luck merely *can* undermine knowledge, but need not always do so.

³ Pritchard (2005b) notes that William Harper (1996), Carolyn Morillo (1984) and Thomas Nagel (1979) are among those who have offered such accounts.

⁴ Pritchard, 2005b, p. 128. Pritchard, further in his account, offers a corollary that addresses the subjective significance of luck. He offers: “(L2) If an event is lucky, then it is an event that is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts)” (132).

⁵ This example, taken from Gettier, E., “Is Justified True Belief Knowledge” (1967), exploits the fact that if one is justified in *p*, then one is justified in *p* or *q* or *z* or *n*... etc. The ‘Barelona’ Gettier case is generated when one is justified in believing *p*—and thus—any disjunctive combination of *p* and other disjuncts. The disjunctive combination is true not because *p* is true, however, but because some other disjunct—for which the agent was not justified in believing—happens to be true.

⁶ To avoid confusion that arises when considering whether crediting knowledge of necessary truths violates the “ought implies can” maxim, I’ll limit the scope of this claim to knowledge of contingent propositions.

⁷ This argument has been advanced in Greco’s 2006 manuscript “What’s Wrong with Contextualism?” to be presented at the 2007 Epistemology and Linguistics Conference at the University of Aberdeen.

⁸ Consider that, even if an agent holds a ‘safe’ belief, and hence satisfies the first desideratum, it will remain an open question whether the agent is credit-worthy for her forming the safe belief.

⁹ I have advanced a proposal along these lines in my 2006 manuscript “A Happy Triumvirate: Virtue Epistemology, Contextualism and the Safety Principle.” What is required for satisfying the second desideratum, on my view, is captured best within a virtue responsibilist framework.

¹⁰ I borrow this handy locution from McDowell, 1998a.

¹¹ Cite contextualism paper

¹² Cohen’s contextualism endorses the fallibilist program in low-stakes contexts, a move that allows him to preserve the neo-Moorean motivation behind P3. However, Cohen would allow that in skeptical conversational contexts, the standards for knowing could be raised to (near) infallibilist criteria, such as would be the case when the denial of BIV scenarios are taken seriously.

¹³ McDowell, J. (1998), “Knowledge by Hearsay.”

¹⁴ McDowell, J. (1998a) “Knowledge and the Internal” p. 403.

¹⁵ McDowell suggests in “Knowledge by Hearsay” that doxastic responsibility will also be relevant in assessing an agent’s standing.

¹⁶ McDowell, J. 1998a “Knowledge and the Internal” p. 403.

¹⁷ From conversation, University of Stirling, Fall 2006.

¹⁸ Although it wouldn’t be likely that this would happen, we can nonetheless build it into the case that it ‘would’ happen.

¹⁹ McDowell, J. “Knowledge by Hearsay” p. 421.

²⁰ McDowell, J. “Knowledge and the Internal” and “Knowledge by Hearsay”, 1998b.

²¹ Millar, A. “Perceptual-Recognitional Abilities and Perceptual Knowledge” (forthcoming).

²² I wish to note here that it is not clear to me how ‘environments’ are supposed to be demarcated. Suppose, for example, that we take on one hand Goldman’s valley salted with barn facades and determine that that valley constitutes an environment. Suppose that you are lost while driving in the general area. You come to a fork in the road; one road leads down Goldman’s barn façade valley, the other road leads down a valley of all genuine barns. You flip a coin and happen to take the road that leads down the road with all genuine barns. It is unclear to me whether this road constitutes a ‘good’ environment simply by virtue of the fact that there are no barn facades in the environment. Or, should we say that it is part of a ‘bad’ environment (i.e. one in which you do not have barn-recognizing perceptual recognitional abilities) by virtue of the fact that you so easily could have driven down Goldman lane. Alan Millar, in conversation, has suggested to me that—at least on his view—it will be inevitable that there will be some indeterminate cases in which the issue of whether one has exercised perceptual-recognitional abilities in a certain environment will not be clear given that the delineation of the environment relative to which one would have such abilities is obfuscated.

²³ Greco, J. (forthcoming 2007) *Epistemic Evaluation: A Virtue-Theoretic Approach*, p. 2.

²⁴ Thanks to Alan Millar for noting that this point deserves clarification.

²⁵ Carter, J.A. “A Happy Triumvirate: Virtue Epistemology, Contextualism and the Safety Principle” (manuscript).

²⁶ Pritchard, D. (2005) *Epistemic Luck*. p. 71.

²⁷ Sosa, E. "A Virtue Epistemology" (manuscript)

²⁸ Greco, J. "Greco, J. (forthcoming 2007) *Epistemic Evaluation: A Virtue-Theoretic Approach*, p. 1.

²⁹ McDowell, J. "Knowledge by Hearsay" p. 429.

³⁰ *Ibid.* p. 429. *fn.*