Episteme, 9, 4 (2012) 361–376 © Cambridge University Press
doi:10.1017/epi.2012.22

JUSTIFICATION AS 'WOULD-BE' KNOWLEDGE

AIDAN MCGLYNN
aidan.mcglynn@gmail.com

ABSTRACT

In light of the failure of attempts to analyse knowledge as a species of justified belief, a number of epistemologists have suggested that we should instead understand justification in terms of knowledge. This paper focuses on accounts of justification as a kind of 'would-be' knowledge. According to such accounts a belief is justified just in case any failure to know is due to uncooperative external circumstances. I argue against two recent accounts of this sort due to Alexander Bird and Martin Smith. A further aim is to defend a more traditional conception, according to which justification is a matter of sufficiently high evidential likelihood. In particular, I suggest that this conception of justification offers a plausible account of lottery cases: cases in which one believes a true proposition – for example that one's lottery ticket will lose – on the basis of probabilistic evidence.

Most epistemologists nowadays are willing to concede that the project of offering a reductive analysis of knowledge has not been a success, and for reasons of principle rather than lack of imagination, ingenuity or application on the part of those who engaged in that project. The prevailing scepticism about this project seems to me to be to some degree merited, though I also think that much recent epistemology has recoiled much too far in the opposite direction. Proponents of knowledge first epistemology typically agree with knowledge's would-be analysers that belief and justification are necessary for knowledge, but they have proposed radically rethinking how they relate to knowledge. Belief and justification are to be illuminated in terms of our grasp of knowledge, not vice versa – which is not, of course, to suggest that reductive analyses of belief and justification in terms of knowledge are possible.

On the traditional picture, presupposed by the project of analysing knowledge, knowledge is built out of more fundamental elements. An item of knowledge is to be understood as a true belief that has sufficiently impressive epistemic credentials, where this involves justification and the right kind of connection between the fact that P and the subject's belief that P. Given this conception of knowledge, most of the action involves spelling out what it takes for a belief to be sufficiently justified, what it takes for it to be hooked up to the relevant facts in the right way, and determining the extent to which buying into these conditions involves compromises of internalism – whatever these amount to. Knowledge first epistemology promises an end to politics as usual. The recurring frustrations encountered by attempts to analyse knowledge in terms of belief and justification are diagnosed as symptoms of having adopted a fundamentally misguided project, one which we need not pursue any further. And knowledge first accounts of justification tend to be unabashedly externalist in character, in the sense that one lacks any kind of special access to the factors in virtue of which propositions or beliefs are justified or unjustified, or to whether a given proposition or belief is justified or not. On such accounts, the externalist

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

character of justification trickles down from the externalist character of knowledge, leaving no room for internalism.

In earlier work (McGlynn 2011), I have argued against this approach to theorising about belief. Belief should not be thought of as a state essentially aimed at knowledge or governed by a knowledge norm, nor is to believe something to treat it as something one knows, or to be committed on pain of irrationality to so treating it.[1] I won't rehearse that discussion here (though I will consider the claim that belief aims at knowledge again below). Rather, I want to turn to knowledge first approaches to justification.

Such accounts provide too broad and varied a target to be treated in single paper. There are no fewer than three distinct proposals to be considered. First we have Timothy Williamson's (2000) influential suggestion that one's total evidence just is one's knowledge, with a proposition being justified to the degree that it's likely in light of one's evidence. Second, Jonathan Sutton (2007) has proposed that in the sense of primary interest to epistemologists, P is justified for S just in case S knows that P. And lastly, Alexander Bird (2007) and Martin Smith (2010) have proposed that justification is a kind of 'would-be' knowledge, in a sense to be explained shortly.[2] Williamson's equation of one's evidence with what one knows demands extended discussion, and I won't attempt such discussion here.[3] Sutton's proposal is utterly implausible on the face of it, and the arguments in its favour have been well critiqued elsewhere (see Coffman 2010; Kelp 2011). Here my focus will be on Bird and Smith's accounts. My primary aim will be to show that these accounts are not well motivated, and that they are problematic. A secondary aim will be to defend a more traditional conception of justification, according to which one has justification for a proposition just when it is made sufficiently likely by the evidence one has. Following Smith, I'll call this the *risk-minimisation* conception. My discussion will fall short of putting us in a position to respond to Smith's objections to this conception once and for all – its defenders still have much work ahead of them. But I do want to argue that this conception of justification looks rather more promising than the alternatives offered by Bird and Smith.

This paper divides into four sections. §1 introduces Bird's account of justification, and argues that Bird fails to motivate the key thought underlying it. §2 considers Bird's claim that his account explains one of the key intuitions that motivates internalism, and finds it problematic. §3 looks at the motivations for Smith's account, which turn on contrasting it with the risk-minimisation account that I favour. I argue that Smith's objections to the risk-minimisation account are at best inconclusive. Finally, in §4 I offer two objections to Smith's own account: one to the claim that he has offered a plausible sufficient condition for having justification, and the second to the more important claim that he has offered a plausible necessary condition for having justification.

1 See also Whiting forthcoming, which independently makes a number of related points against these theses.

2 The phrase originally comes from Sutton (2007: 10), but it is co-opted by Bird (2007: 83). I should note that while there's a sense in which Smith puts forward a knowledge first account of justification, it is unclear to what extent he would endorse the main tenets of knowledge first philosophy as I have characterised them here. I will discuss this point more fully when I introduce his account below.

3 I will address Williamson's position in detail in McGlynn in progress.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

1. JUJU

Bird's core idea is that a belief or judgement is justified when it fails to be knowledge only because of non-mental confounding factors, of the sort present in standard Gettier cases.[4] Bird formulates this view as follows (2007: 84):

(*JuJu*) If in world w_1 S has mental states M and then forms a judgment, that judgment is justified if and only if there is some world w_2 where, with the same mental states M, S forms a corresponding judgment and that judgment yields knowledge.[5]

We can get a feel for this account by considering three test cases. The first is a standard Gettier case (though I adapt this particular example from Russell 1948, rather than from Gettier 1963). Ashley looks at her hitherto impeccably reliable watch, and sees that it reads two o'clock. She forms the belief that it is two o'clock on this basis. As it happens, Ashley's watch stopped exactly twelve hours prior, and she just happened to look at it during a minute when trusting it would give her a true belief. The second test case we can consider is a variant, in which Bryan correctly guesses that it is two o'clock. Lastly, we can consider a lottery case, in which Clare believes a 'lottery proposition' (Vogel 1990); she has a true belief that her lottery ticket – one of one million in a draw with a single guaranteed winner – will lose, basing this belief on the probabilistic grounds available to her before the outcome of the draw has been announced.

As Bird interprets (*JuJu*), it entails that Ashley is justified but Bryan is not. Ashley's belief is justified because she forms it in such a way that it would have been knowledge had her environment been suitably cooperative (2007: 85). Bryan's guess isn't like this at all; since Bryan has no good evidence that it's two o'clock, there's no world in which someone with the same mental states as Bryan reaches knowledge. Much more controversially, (*JuJu*) entails that Clare is no better off than Bryan.[6] According to Bird, Clare's belief also fails to be knowledge for reasons internal to her mental life, and so for which she, rather than an uncooperative world, is responsible. That's because Clare's belief is based on what Bird (following Nelkin 2000) calls a *P-inference*: an inference of the form 'P has a statistical probability of n (where n is very close to 1) therefore P' (2007: 102).

4 I won't worry about the distinction between belief and judgment here. Bird does distinguish these (2007: 96–7), but the distinction doesn't play a significant role in his discussion, and generally he seems pretty happy with moving back and forth between them fairly freely.

5 Notice that there's no requirement that w_1 and w_2 be distinct, and so justification remains a necessary condition for knowledge on Bird's account. Notice too that (*JuJu*) does not require that there is a world in which one's judgement yields knowledge, but only that some 'corresponding' judgement does. This enables Bird to leave room for justified judgements and beliefs in necessary falsehoods. It is the counterpart of a standard move made by defenders of a safety condition on knowledge in order to leave room for unsafe beliefs in modally robust propositions, such as necessary truths. Several people have suggested to me that this complication in Bird's account will lead to trouble when it comes to knowledge of one's own mental states. I haven't been able to construct a convincing objection of this sort, which is why I do not pursue the point in the text.

6 I'm assuming here that Clare cannot know that her ticket has lost. Not everyone accepts that lottery propositions cannot be known, and Bird notes that (*JuJu*) is compatible with one both knowing and being justified in believing that one's ticket has lost. However, he seems strongly inclined to think that one can't have such knowledge, and both Smith and myself agree, so I'll take this as a piece of common ground throughout.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

(*JuJu*) is primarily supported by the idea that knowledge is the aim of belief. Bird is quite explicit about this (2007: 93):

> Justification is a kind of near miss. Knowledge is the epistemic bull's-eye. Justification occurs when an unexpected gust of wind nudges the arrow off-centre. Unjustified belief occurs when one's aim was poor or one fumbled the release. The key element of this view is that belief aims or ought to aim at knowledge; knowledge is thus the primary norm for belief.

Consequently, Bird's principal motivation for (*JuJu*) consists of an argument for the claim that knowledge is the aim of belief. The argument comes in the following passage, which I quote in full (2007: 93–4; italics in original):

> The question amounts to this, is the truth of a belief enough for it to have achieved its constitutive aim, for it to have fulfilled its function fully? Or is mere truth not alone sufficient for the function to be fulfilled? Certainly falsity is a good reason to give up a belief – its presence marks a lack of achievement on the part of the believer. So truth is *part* of the aim of belief. But, I claim, it is not enough. If truth were sufficient for the fulfilment of its constitutive aim then one could have no complaint with someone whose beliefs are true. But clearly this is not right. A belief that is merely accidentally true and lacks justification may rightly be criticised on this score (even a self-deceiver may be accidentally right). So, I suggest, justification is required as well. Are justification and truth jointly sufficient for the success of belief? It would be odd if the norm of belief were the conjunction of two independent factors. As Gettier-style examples show, there is nothing special about a belief that is both justified and true. The best explanation of all this data is, I propose, that *knowledge* is the aim of belief.

This argument is questionable on a number of fronts, even if one concedes that belief aims at truth.[7] The main defect of the argument is its reliance on the deeply implausible assumption that, if truth were sufficient to fulfil the constitutive aim of belief, one could have no complaint with any true believer. Complaints can be levelled at failures to fulfil aims that are not constitutive of the act being performed – or so one naturally supposes in the absence of any argument to the contrary.[8] For example, Williamson (2000: 238) points out that assertions can be criticised as rude, but no one would suggest that avoiding rudeness is a constitutive aim of assertion.[9]

In a footnote (2007: 93, n. 8) Bird suggests that his account 'does not strictly depend' on his argument for the claim that knowledge is the aim of belief, since (*JuJu*) might still be

7 See McGlynn 2011: §6 and Whiting forthcoming: §5.

8 Whiting (forthcoming) offers a similar criticism, noting that defenders of a truth norm of belief typically take it to generate secondary norms, including a justification norm. Allan Hazlett has pointed out a closely related way to cast doubt on Bird's assumption. The defender of a truth norm of belief might naturally distinguish evaluations of beliefs from evaluations of believers, just as defenders of a truth norm of assertion will distinguish between evaluations of assertions and evaluations of asserters (Weiner 2007: 193). Then the fact that it can be appropriate to criticise a true *believer* won't show that is something defective about the *belief*, contrary to Bird's claim.

9 Declan Smithies (2012) also argues that belief aims at knowledge. His first argument turns on Williamson's 'observation' (2012: 284) that a belief that P is defective if one is not in a position to know that P. This is hardly a datum. His second rests on his independently implausible claim that one has justification to believe P if and only if one has justification to believe that one is in a position to know that P. See Smith 2012b for, to my mind, convincing criticism of this claim.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

motivated on the grounds that it is 'a component of a package that overall provides a more satisfactory explanation of various phenomena than its competitors'. This is a little puzzling; as Bird notes in the passage to which this footnote is attached (quoted above), that belief aims at knowledge is 'the key element' of his account, and if its key element is left unmotivated, surely that is some reason not to accept it.

But let us set that concern aside. Bird's claim that his account gains motivation from its explanatory power rests primarily on two points. First, Bird suggests that it offers an improvement on Dana Nelkin's (2000) unified response to the knowledge and justification versions of the lottery paradox (2007: 100–3). Second, he argues that it explains the truth of an intuition typically taken to support internalist conceptions of justification over externalist ones (2007: 97–100). Bird's treatment of the lottery paradox rests largely on his contention that belief aims at knowledge (2007: 101), and I will discuss lotteries at some length below when considering Smith's account, so I won't consider this further here. In the next section I will evaluate Bird's claim to be able to explain a key internalist intuition.

2. MENTALISM

Let *mentalism* be the thesis that justification supervenes on mental states alone. Mentalism is usually associated with internalism; indeed, mentalism is sometimes taken to be the core claim of internalism (e.g. Conee and Feldman 2001). Bird presents the truth of mentalism as a 'central internalist intuition' (2007: 83), but notes that it is entailed by (*JuJu*). And so Bird claims that an advantage of his (externalist) account is that it can explain the truth of this supposedly internalist intuition.

However, I think it is a mistake to describe the truth of mentalism as an internalist intuition. Rather, it is a thesis that some internalists have tried to motivate on the grounds that it best explains the intuitions we have about certain cases. First of all, we can construct a series of cases in which one subject is justified while another is not (or the first subject is better justified than the second), while the only relevant differences between them seem to be in their mental states (Conee and Feldman 2001). Second, there is a powerful intuition that the victim of a Cartesian demon can be just as justified as we are, and one explanation of this is that such victims might be mental duplicates of ordinary epistemically successful subjects like ourselves.[10]

Bird follows Williamson in taking knowledge to be a mental state (2007: 82), and so he argues that factive mental states should be allowed into the supervenience base of justification on the grounds that otherwise the internalist 'intuition' will be question-begging (2007: 97), and since 'when it is qualified by reference to the semi-technical idea of a non-factive mental state, it begins to look less like a ubiquitous intuition and more like a substantive thesis in need of a defence' (2007: 98). But mentalism *is* a substantive thesis in need of defence. Moreover, this remains as true when the thesis is in Bird's hands as it is when it is in the hands of the internalist. And it is not obvious that Bird's externalist version of mentalism inherits either of the standard internalist motivations. Perhaps an argument parallel to that offered by Conee and Feldman can be constructed. This

10 An example is provided in the passage from Ralph Wedgewood that Bird quotes (2007: 97–8), though tellingly Bird ignores this aspect of it. I should note that I'm not endorsing these internalist arguments for mentalism – I'm only noting that internalists have supplied such arguments.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

would involve assembling sufficiently many examples in which it is plausible that a difference in justificatory status is turning on differences in the mental states of the subjects, *including their factive mental states*.[11] Bird offers no such argument. As for victims of the 'new evil demon', accepting Williamson's thesis that there are factive mental states has the consequence that no such victim can be the mental duplicate of any ordinary epistemically successful subject, since such subjects will have knowledge that the demon victim lacks. This consequence may have its advantages, particularly when it comes to combating scepticism (see Williamson 2000: ch. 8). But it is easy to see that it renders the comparison between such victims and ordinary successful subjects useless for motivating mentalism.

Bird's account of justification entails mentalism, but only in a form that tears it completely free from its moorings. In the end, it is simply another unmotivated aspect of his account.

3. RISK-MINIMISATION AND NORMIC SUPPORT

Smith offers an alternative way of making more precise the thought underlying Bird's account, namely that my belief is justified just I case I have done everything I can to ensure that it is knowledge, any failure being down to the world failing to hold up its end of the bargain (Smith 2010: 12, 17). This is the sense in which Smith too holds that justification is 'would-be' knowledge. He also says that Bird's (*JuJu*) 'comes close to being right' (2010: 27–8, n. 2). But we should note some differences, even before we get into the details of Smith's account. First, while Bird's account is an account of *doxastic* justification (what it takes for a belief or judgement to be justified), Smith is more concerned with *propositional* justification (what it takes for one to have justification for some proposition, which one may or may not avail oneself of).

Another difference is that, while Bird explicitly identifies his account as a contribution to knowledge first philosophy, Smith does not. Indeed, he expresses some misgivings about theses that, while not quite constitutive of knowledge first philosophy as I characterised it in the introduction, tend to be held by its proponents (for instance, that one's total evidence is one's total knowledge, and that knowledge is the norm of assertion). Nonetheless, Smith's account is naturally taken as one in which knowledge has an important kind of priority over justification: like Bird, Smith sees justification as a kind of 'would-be' knowledge. On issues about justification, Smith's aims and views align very closely with proponents of knowledge first philosophy, even if they diverge significantly elsewhere.

Smith develops and motivates his account by contrasting it with a more traditional conception of justification, which he calls the *risk-minimisation* conception. According to this conception, one's having justification for P is a matter of P being sufficiently likely given one's evidence, and the more likely P is given one's evidence, the more justification one has for P. The evidential likelihood of a proposition P is neither the degree to which one actually believes it, given that evidence, nor the objective likelihood that it is true. Williamson (2000: 209–11) argues persuasively that the evidential likelihood of P cannot be identified as the credence a perfectly rational agent in possession of our current evidence would have

11 The thought here is that if the italicised condition is not met, then the cases cannot provide any motivation to broaden the supervenience base to include factive mental states.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

in P, though Smith suggests (2010: 12) that this comes close to the intended notion, and so perhaps this identification will suffice to fix ideas enough for the discussion to follow.

So the risk-minimisation conception sees justification as a matter of high evidential likelihood, and having more or less justification to believe a proposition as a matter of that proposition's being more or less evidentially likely.[12] These are the claims that Smith attempts to cast doubt on, and he does so using the following case (2010: 13–41, adapted from Nelkin 2000):

> Suppose that I have set up my computer such that, whenever I turn it on, the colour of the background is determined by a random number generator. For one value out of one million possible values the background will be red. For the remaining 999 999 values, the background will be blue. One day I turn on my computer and then go into the next room to attend to something else. In the meantime Bruce, who knows nothing about how my computer's background colour is determined, wanders into the computer room and *sees* that the computer is displaying a blue background. He comes to believe that it is. Let's suppose, for the time being, that my relevant evidence consists of the proposition that (E_1) it is 99.9999% likely that my computer is displaying a blue background, while Bruce's relevant evidence consists of the proposition that (E_2) the computer visually appears to him to be displaying a blue background.

The conclusion we are invited to draw is that Bruce has justification to believe that the computer background is blue, while Martin does not.[13] But, so Smith argues, the likelihood that the computer screen is displaying a blue background is lower given Bruce's evidence (E_2) than given Martin's evidence (E_1). However unlikely it is that Bruce is hallucinating, being subject to a sudden bout of colour-blindness, being taken in by a cunning illusion which makes it appear that the background is blue when it is in fact red, and so on, 'the likelihood, given Bruce's evidence E_2, that the computer is displaying a blue background would be no where near as high as 99.9999%' (Smith 2010: 14). And so the worry is that the risk minimisation conception of justification gets this case completely wrong, since it entails that Martin has more justification than Bruce, when in fact Bruce has justification and Martin does not.

Smith's alternative proposal is that to have justification to believe a proposition P it is necessary that one's body of evidence E *normically support* P, where for E to normically support P is for the most normal worlds in which E is true to be worlds in which P is true (2010: 17). Consider how this works in the above case. The idea is that there is nothing abnormal about a world in which Martin's evidence (E_1) is true, but in which the screen background is red. It is statistically highly unlikely, but it requires no special explanation. In contrast, the most normal worlds in which Bruce's evidence (E_2) is true are all ones in which the computer screen displays a blue background; the worlds in which the computer visually appears to him to be displaying a blue background when it is not involve abnormalities like hallucination, random bouts of colour-blindness and the like.

12 This is obviously just a sketch of the risk-minimisation conception. The basic idea can be elaborated in a number of different ways, and I intend to stay neutral on these further issues about what shape it should take here. The core idea should be clear and familiar enough for the discussion to follow.

13 As is clear from the quote, Smith features in his own example. For clarity I will use 'Smith' to refer to the philosopher, and 'Martin' to refer to the subject in the example. This follows a precedent set by discussions of Keith DeRose's 'bank-cases' in the literature on epistemic contextualism.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

In the remainder of this section I will consider whether this example really does favour Smith's normic support account over any version of the risk-minimisation conception, concluding that it does not. In the following section I will develop two objections to Smith's account.

Clearly an important aspect of Smith's attempt to motivate an alternative to the risk-minimisation conception is his claim that Martin lacks justification in the example he describes. The truth of this claim is hardly obvious; indeed, as Smith notes, a number of philosophers have found it natural to assume that such examples *are* cases of justification (see, for instance, Lewis 1996: 551; Hawthorne 2004: 9; Pritchard 2008: 440; and perhaps even Williamson 2000).[14] Against this assumption Smith (2010: 14) offers four considerations. First, that it is 'natural' to describe Martin's belief, should he form one on the evidence available to him, as a presumption, but this is not a natural description of Bruce's belief. Second, that Bruce's belief is a very promising candidate for knowledge, and in fact will be knowledge if the case is filled out in the most natural way, while in contrast Martin's belief isn't a good candidate for knowledge at all. Third, that Bruce can outright assert that the screen is blue, while Martin should say something more qualified, like 'It's overwhelmingly likely that the background is blue'. And finally, that presumably Martin should only believe the qualified thing too. Smith concludes that 'the implication of these considerations is clear enough: Bruce has *justification* for believing that the computer is displaying a blue background while [Martin] does not' (2010: 14).

I don't think this supposed implication is clear at all. Consider the first point. We naturally call something a presumption when one's evidence for it is manifestly circumstantial: that is, when it clearly involves ampliative inference from one's evidence. Calling one's belief a presumption carries no implication that one lacks justification. Knowing that Dave is a naturally reclusive character, and that he hasn't come to any of the last thirty dinner parties I have invited him to, I believe that he will not come to my next dinner party. It is natural to describe this as a presumption, so long as Dave hasn't actually explicitly turned down this particular invitation. I can say to my partner, 'I'm presuming that Dave isn't coming this time either, so I won't bother setting a place for him', though again, only so long as Dave hasn't explicitly told me he is not coming to this particular party. Still, there's no temptation on that score to say that my belief is unjustified. Sometimes circumstantial evidence can be good enough for justification or even knowledge. All that we can conclude from Smith's first point is that it is natural to think Martin's evidence for believing that the computer background is blue is obviously circumstantial, but that Bruce's is not. But that's pretty uncontroversial; Martin's belief that the computer background is blue is the conclusion of an inference from the probabilities involved, while Bruce is just looking directly at the blue screen.

14 One then needs to qualify the principle of multi-premise closure for justification if one is to avoid the lottery paradox, but as Smith notes, the risk-minimisation conception commits one to qualifying that principle anyway, due to familiar considerations about risk-aggregation. A more worrying possibility is that the risk-minimisation conception also commits one to qualifying even single premise closure for justification, since it has been argued that we find a form of risk-aggregation even with supposedly single-premise inferences (DeRose 1999: 23, n. 14, Lasonen-Aarnio 2008). This is not the place to engage with this challenge (though see Smith 2012a). I will argue below that Smith's conception of justification is no better placed to preserve standard closure principles for justification, and so the issue, while important, will turn out to be orthogonal to the present debate.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

Smith's second point rests on his assumption that justification is what makes a belief a good candidate for knowledge (2010: 12). This can be read in at least two importantly different ways. Interpreted as a kind of generic claim or as a functional characterisation of justification, it seems hard to contest. However, so read, it provides no support for Smith's conclusion. Justification being the kind of thing that makes a belief a good candidate for knowledge doesn't entail that if a belief isn't a good candidate for knowledge, then it isn't justified. Smith needs the stronger interpretation, according to which any belief that is not a good candidate for knowledge isn't justified. But the stronger interpretation lacks the intuitive appeal of the weaker claim, and Smith provides no argument for it.[15]

Let's turn to Smith's third observation, that while Bruce can outright assert that the screen is blue, Martin can only assert something more qualified. Now, defenders of the knowledge account of assertion, according to which one ought to assert P only if one knows P, will wonder why his observation isn't adequately explained by observing that Bruce knows that the screen is blue, while Martin doesn't. Since I'm not inclined to accept the knowledge account, and Smith doesn't seem to be either, I'll leave this response aside here.

Even without the knowledge account of assertion, the inference from the claim that Martin cannot outright assert that the screen is blue to the conclusion that he lacks justification looks suspect. I believe Dave won't come to my next dinner party on the grounds that he hasn't come the previous thirty times I have invited him to such a gathering, rather than because he's responded to this particular invitation. It is plausible that I should not flat-out assert that he's not coming; rather I should say 'Dave's probably not coming', or 'Presumably Dave's not coming', or something like that. But again, the case should strike most non-sceptics as a routine instance of an inductively justified belief. So there's no general implication from fact that one should make assertions that are qualified in this manner to the conclusion that one lacks justification to believe the unqualified claim.[16]

The bearing of Smith's fourth consideration, that Martin should only believe the qualified claim until he has investigated further, is more immediate. But in fact, it is too immediate. Smith (2010: 14) assumes a 'minimal conception' of justification, which is supposed to be common ground between proponents and opponents of the risk-minimisation conception:

> What I mean when I say that one has justification for believing P is simply that it would be epistemically or intellectually appropriate for one to believe or accept that P is true, given the evidence at one's disposal.

This thin characterisation of justification does seem unobjectionable. However, once we adopt it, Smith's presumption that Martin ought not believe the unqualified claim that the screen is blue given the evidence available to him looks tantamount to a presumption

15 See McGlynn 2011 for criticism of a number of arguments that one might appeal to in support of the stronger reading of the claim.

16 Of course, to say this much is not to explain why Martin cannot outright assert that the computer screen is blue. One strategy is to appeal to the idea that there's some kind of pragmatic impropriety here (Weiner 2005; Lackey 2007; thanks to Allan Hazlett for suggesting this). I'm sceptical that this kind of explanation can be made to work (compare McGlynn 2011: §8), and so I need to offer an alternative. I leave that as a task for another occasion (McGlynn in progress).

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

that he does not have justification for that claim. But that's just the conclusion that Smith is working for.

I conclude that none of Smith's four considerations should persuade us that in the case described Martin lacks justification to believe the unqualified claim that the computer screen background is blue.[17] And to the extent that we find it plausible that Martin does have justification in the example, we have reason to favour the risk-minimisation conception over Smith's alternative.

There is, however, a residual reason to think that this example favours Smith's account. Even if it is not clear that the risk-minimisation conception gets the wrong verdict on whether Martin *has* justification, one might still contend that it mistakenly implies that Martin has *more* justification than Bruce. Smith's account doesn't face this worry. It is part of Smith's account that 'in order for one to have *more* justification for believing a proposition P than a proposition Q, it is necessary that one's evidence E normically support P more strongly than Q' (2010: 17). And plausibly Bruce's evidence normically supports the proposition that the computer background is blue more strongly than Martin's does.

How might a defender of the risk-minimisation conception respond to this challenge? Here I will have to content myself with outlining one strategy that I think is promising, though the issue deserves a more extended treatment.[18]

The simplest proposal is just to bite the bullet, and insist that in the example described Martin in fact has more justification than Bruce. Now, I do think that it should be conceded that it is plausible that, in some sense, Martin is epistemically worse off with his probabilistic evidence than Bruce is with his non-probabilistic evidence. But we can explain this without conceding that Bruce has more justification than Martin. On the account of knowledge I favour, knowing requires that one's belief meet a safety condition; very roughly, a belief is safe just in case in all the close worlds in which one forms that belief, it is true.[19] Bruce's belief is safe, while Martin's is not. In this sense I think Smith is right that Bruce's belief is a good candidate for knowledge while Martin's is not. But Martin's epistemic situation being inferior in this respect is perfectly compatible with it being superior in others, including those that determine degree of justification. The resulting picture will seem rather strange from the perspective of one who thinks of belief

17 In a recent paper Declan Smithies (2012) argues that one isn't justified in believing lottery propositions and the like, but only that it is very likely that (say) one's lottery ticket has lost. On Smithies's view, belief is distinguished from a very high degree of confidence by one's willingness to take P as settled, and so to take P for granted in one's practical reasoning. By these lights, it is plausible that one doesn't really believe that one's ticket will lose after all, and so it is not a big concession to hold that such a belief cannot be justified. But what Smithies treats as the crucial difference between believing P and mere high confidence that P can also be seen as the difference between believing P while believing that one knows P and believing P while recognising (or at least suspecting) that one does not know P, and Smithies offers no argument for preferring his view over this alternative. (It is no part of my alternative proposal to deny that there is a difference between belief and mere high confidence, of course. It is only where and how to draw the line that is at issue.) Moreover, I worry that on Smithies's proposal we will count as believing very little, particularly about the future.

18 It is worth mentioning another strategy that has been defended recently, due to Comesaña (2009), who responds to a similar case offered as an objection to reliabilism in Adler 2005. Those who find the proposal in the main text unsatisfactory may want to consider Comesaña's proposal more carefully than I have the space to do here.

19 Smith is also a proponent of a safety condition on knowledge (see in particular Smith 2009). Indeed, Smith's normic support condition on justification is modelled on Safety (Smith 2009: 176).

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

as aiming at knowledge, and of degrees of justification as relating to how good one's attempt at hitting that target is. But that is precisely the perspective urged by Bird, and above we found it lacking motivation.

It is worth considering a variant of the example also discussed by Smith (2010: 18–19).[20] Let us take things to be as in the earlier case, but this time let us suppose that Bruce *is* hallucinating, and that despite the odds the computer background is red rather than blue. Smith holds that this does nothing to shake our sense that Bruce has more justification to believe that the screen is blue than Martin does – indeed, that Bruce has justification while Martin does not (2010: 19). But obviously in this variant Bruce is no more able to form a safe belief than Martin is, and so one cannot appeal to this difference between them to explain any sense we have that, at least in some respect, Bruce's belief has more epistemically going for it than Martin's. However, I no longer find myself particularly tempted to say there is any such respect. I agree that, given the way the variant is described, Bruce remains justified in believing that the computer background is blue. And I am willing (for present purposes at least) to grant that Bruce has the same evidence in the variant as he had in the original example. But despite this, I struggle to see any respect in which Bruce's false belief has more going for it than Martin's false belief.

A second variant also appears to unsettle the strategy I have suggested.[21] Suppose that Martin believes that the computer background is blue on the basis of his knowledge that this is overwhelmingly likely, given the setup of the computer. As in the variant discussed in the previous paragraph, the computer background is in fact red. However, this time Bruce's perceptual faculties are working just fine, and on the basis of the computer background's visual appearance, he comes to believe that it is red rather than blue. The risk-minimisation conception seems to commit us to saying that Martin has more justification for his false belief than Bruce has for his true one, and this might seem intolerable. In fact, the only real difference between this variant and the original example described by Smith is that, against the odds, Martin has a false belief. As a matter of psychology, knowing this may make us more tempted to say that Martin cannot have more justification than Bruce. But on the risk-minimisation conception, as on most accounts of justification, the truth-value of Martin's belief is irrelevant to how much justification he has. And so the clear-headed defender of the risk-minimisation should insist that this second variant raises no new issues. As in Smith's original example, Martin has more justification than Bruce, though Bruce's belief has something epistemically going for it that Martin's does not, namely that it is safe.

I have stressed that how best to respond to these kinds of objections is a large and complicated issue, one meriting further discussion. It may be that further investigation will reveal that the risk-minimisation conception really does flounder on lottery cases, or that one can tinker with it to get results that fit better with those that Smith thinks are correct. But for now the conclusion I want to draw is just that no strong argument against the risk-minimisation conception of justification seems to be in the offing here.

20 Thanks to Martin Smith for pointing out that this variant might reveal a problem for my proposal.

21 Thanks to Filippo Ferrari.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

4. NORMICAL SUPPORT, DEGREES OF JUSTIFICATION AND GETTIER CASES

In this section I will offer two objections to Smith's normical support account. The first will be to the claim that normical support is sufficient for having justification, and the second to the more important claim that it is necessary for having justification.

Smith (2010: 28, n. 6) notes that his account implies that if one has more justification for P than for Q, then one has justification for P. If one has more justification for P than for Q, then any world in which E is true and P false is less normal than the most normal worlds in which E is true and Q false (2010: 17). But then, assuming that there are maximally normal worlds in which E is true, it follows that the most normal worlds in which E is true are worlds in which P is true, which is just to say that E normically supports P.

Smith isn't worried by this implication, since he suggests that certain other gradable adjectives work the same way: for example, if S is wetter than T, then S is wet. And just as it is weird to say 'S is wetter than T, but S is not wet', Smith thinks it sounds odd to say 'One has more justification for believing P than Q, but one does not have justification for P'.

But there's a crucial ambiguity in 'having justification' that Smith doesn't take into account. It can mean that one has *some* justification (however slight). This isn't the interpretation that is relevant to Smith's discussion, since as we noted earlier, he operates with a minimal working conception according to which if one has justification then it is 'epistemically or intellectually appropriate for one to believe or accept that P is true, given the evidence at one's disposal' (2010: 12). But notice that the kinds of claims Smith thinks are odd are only odd on the first, irrelevant interpretation:

> One has more justification for believing P than Q, but one has no/hasn't any/hasn't the slightest justification for believing P.

If we disambiguate in favour of the second, relevant interpretation, the oddness disappears:

> Given the evidence at one's disposal, it's more epistemically or intellectually appropriate for one to believe P than for one to believe Q, but it's not epistemically or intellectually appropriate for one to believe P.

So, contrary to Smith's suggestion, it should be reckoned a cost of his account that it entails the claim that, if one has more justification for P than for Q, then one has justification for P. For in the sense of 'having justification' relevant to Smith's discussion, that claim is genuinely counterintuitive.

We must note that Smith's argument that his account has the consequence just discussed implicitly assumes that normical support is to be taken as a *sufficient* condition for having justification; if it were not taken to be sufficient, all that we would be able to conclude is that if one has more justification for P than for Q, then one's evidence normically supports P. Now, Smith does argue for the sufficiency claim towards the end of his paper, on the grounds that doing so enables him to preserve a multi-premise closure principle for justification (2010: 26–7). However, he also considers a hybrid account of having justification for P, which in addition to E normically supporting P requires that P be sufficiently evidentially likely given E (2010: 26). It is clear that Smith intends his official

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

position to be consistent with either of these options, and so while this first objection suggests that Smith's preferred version of the normical support account is more problematic than he reckoned, the more interesting issue concerns the defensibility of the necessity claim.[22] I turn to that issue now.

Smith wants to embrace the conclusion that one cannot have justification for lottery propositions and certain lottery-like propositions (e.g. that the computer screen in the example above is displaying a blue background), but he wants to resist the conclusion that subjects in standard Gettier cases lack justification (2010: 10–2).[23] But this puts an enormous strain on our grip of when normical support is present and when it is absent. Consider again the case of Ashley, who forms the true belief that it is two o'clock by looking at a generally reliable but stopped watch. Is there really any greater abnormality in a reliable watch stopping than in one's lottery ticket winning? Does the former cry out for explanation in a way that the latter doesn't? I find it hard to see why it should; after all, even very reliable watches will stop from time to time. Now, the one Gettier case Smith explicitly discusses clearly does involve abnormality – one correctly believes that a wall is red on the basis of its red appearance, but the wall is bathed in very strong red light emanating from a cleverly hidden source (2010: 11). But the point here is just that it is not at all obvious that this is a feature of standard Gettier cases in general, as Smith needs it to be.[24]

How might we adjudicate this dispute over whether Smith's normical support condition is met in all standard Gettier cases? It is not really clear. In support of the claim that normical support is absent in lottery cases, Smith (2010: 20. 21) cites Vogel's plausible claim that while one's ticket is very unlikely to win, there's an 'intuitive sense' in which there would be no abnormality in that happening (1990: 16). That strikes me as fair enough, but the worry is that there seems equally to be an intuitive sense in which there's nothing abnormal about a very reliable watch stopping.

Smith also offers a more theoretical reason to think that lottery and lottery-like cases differ from standard Gettier cases in the way he suggests. He writes (2010: 17):

> If one believes that a proposition P is true, based upon evidence that normically supports it then, while one's belief is not assured to be true, this much is assured: If one's belief turns out to be false, then the error has to be attributable to *mitigating circumstances* – the error can be explained in terms of disobliging environmental conditions, or cognitive or perceptual malfunction or some such.

---

22 However, notice that the hybrid account that this first objection pushes Smith to is no better placed to respond to challenges to closure principles resting on considerations about risk-aggregation than the risk-minimisation conception is, since these challenges turn only on the claim that high evidential likelihood is *necessary* for justification. This substantiates my earlier claim (n. 14) that such challenges do not provide Smith with any ammunition against the risk-minimisation conception.

23 I make the contrast with *standard* Gettier cases since although I think that lottery cases are Gettier cases at some level of abstraction – they are cases of justified true belief that fail to be cases of knowledge – I also think that there are significant differences between lottery cases and standard Gettier cases; only in the latter is it a matter of luck that one's belief is true. See McGlynn 2011: §1 for details and defence.

24 I think a parallel point could be argued concerning the lottery variants discussed in Vogel 1990 and Hawthorne 2004. Smith anticipates this worry, and argues that these lottery variants are actually unlike lottery cases and cases such as Smith's belief that the computer background is blue (2010: 21–2). I don't find Smith's response entirely satisfactory, but I won't push this point further here.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

Like Bird, Smith takes one significant difference between lottery (and lottery-like) cases and standard Gettier cases to be that only in the latter can an error be chalked up to mitigating circumstances of this sort. This seems plausible enough. But the desired conclusion, that there is normical support present in standard Gettier cases, doesn't follow from any of this. For this to follow, we'd need the converse of the claim Smith makes in the passage just quoted; that whenever error is attributable to mitigating circumstances, then one's belief must be based on evidence that normically supports the proposition in question. However, the latter is precisely the claim called into question by cases like that of Ashley. Were Ashley's belief that it's two o'clock to prove wrong, that would indeed be attributable to mitigating external circumstances: specifically, to her watch having stopped. But the worry remains that it is just not clear what sense of 'normal' licenses the claim that worlds in which Ashley's watch reads two o'clock but in which it is not two o'clock cannot be just as normal as worlds in which one's lottery ticket is almost certain to lose and yet wins. And so I remain unconvinced that normical support is a plausible necessary condition on having justification.

I don't pretend that either of these worries offers a decisive objection to Smith's account. As we have seen, my first objection only targets the claim that normical support is sufficient for having justification, a claim that isn't an essential component of Smith's position. And perhaps some method for deciding when normical support is present and when it is absent will bear the weight Smith's contrasting treatment of lottery cases and standard Gettier cases places on this distinction. As things stand, however, Smith's account strikes me as rather problematic.

CONCLUSION

Three limitations of my discussion are worth reiterating. First, I have not tried to pretend that the risk-minimisation conception is problem free. There is a knot of issues concerning lotteries and closure principles that the risk-minimisation conception seems to feed, and while I have argued that abandoning the conception in the face of these issues is an overreaction, they clearly deserve careful further consideration. Second, I have done little or nothing here to defend internalism about justification. Now, part of the reason I take the task of engaging with such knowledge first accounts of justification to be of significance is its bearing on the kind of resolution – or perhaps better, dissolution – of the internalism–externalism debate discussed in the introduction. But none of the arguments in this paper have turned on intuitions or principles that are in any way the exclusive property of internalists, and the risk-minimisation conception I have defended is neutral on this issue. Third, as I stressed in the introduction, nothing I have said here constitutes any kind of argument against Williamson's knowledge first account of justification; indeed, it is naturally interpreted as an externalist version of the risk-minimisation conception, according to which one has justification for P just in case P is sufficiently likely in light of what one knows.

What I have done here is challenge several important senses in which justification might be thought to be a subsidiary of knowledge. First, I have argued against two attempts to flesh out the thought that justification is 'would-be' knowledge, and I have tried to defend a more traditional conception of justification, according to which having justification for P is a matter of P being sufficiently likely given one's total evidence. Second, I have resisted

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

the claim that one cannot have justification for lottery propositions and other propositions that one recognises one cannot know on the basis of the evidence at one's disposal – 'known unknowns', in Sutton's (2007) terminology – and the associated claim that belief aims at knowledge. So while my discussion of Bird and Smith falls far short of putting us in a position to dismiss the idea that justification should be understood in terms of knowledge once and for all, I hope that it will nonetheless provide encouragement to those who continue to hold that justification possesses a significant degree of autonomy from knowledge.

ACKNOWLEDGEMENTS

My thanks to audiences at the Northern Institute of Philosophy and the Epistemology Research Group at the University of Edinburgh, and to the Northern Institute of Philosophy Work in Progress (NIP-WIP) group. Particular thanks are due to Dylan Dodd, Filippo Ferrari, Allan Hazlett, Jonathan Jenkins Ichikawa, Carrie Ichikawa Jenkins, Federico Luzzi, Duncan Pritchard, Daniel Whiting and Crispin Wright. A conversation with Martin Smith first got me thinking about these issues, and I'm very grateful to him for discussion and for helpful feedback on an earlier draft.

REFERENCES

Adler, Jonathan. 2005. 'Reliabilist Justification (Or Knowledge) as a Good Truth-Ratio.' *Pacific Philosophical Quarterly*, 86: 445–58.

Bird, Alexander. 2007. 'Justified Judging.' *Philosophy and Phenomenological Research*, 74: 81–110.

Coffman, E. J. 2010. 'Is Justified Belief Knowledge? Critical Notice of Jonathan Sutton, *Without Justification*.' *Philosophical Books*, 51: 1–21.

Comesaña, Juan. 2009. 'What Lottery Problem for Reliabilism?' *Pacific Philosophical Quarterly*, 90: 1–20.

Conee, Earl, and Feldman, Richard. 2001. 'Internalism Defended.' Reprinted in E. Conee and R. Feldman, *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.

DeRose, Keith. 1999. 'Introduction.' In K. DeRose and T. A, Warfield (eds), *Skepticism: A Contemporary Reader*. Oxford: Oxford University Press.

Gettier, Edmund. 1963. 'Is Justified True Belief Knowledge?' *Analysis*, 23: 121–3.

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

Kelp, Christoph. 2011. 'Not Without Justification.' *Dialectica*, 65: 581–95.

Lackey, Jennifer. 2007. 'Norms of Assertion.' *Noûs*, 41: 594–626.

Lasonen-Aarnio, Maria. 2008. 'Single Premise Deduction and Risk.' *Philosophical Studies*, 141: 157–73.

Lewis, David. 1996. 'Elusive Knowledge.' *Australasian Journal of Philosophy*, 74: 549–67.

McGlynn, Aidan. 2011. 'Believing Things Unknown.' *Noûs*, publ. online: 29 Dec., doi: 10.1111/j.1468-0068.2011.00843.x.

—— In progress. *Knowledge First?* Under contract with Palgrave-Macmillan.

Nelkin, Dana. K. 2000. 'The Lottery Paradox, Knowledge, and Rationality.' *Philosophical Review*, 109: 373–409.

Pritchard, Duncan. 2008. 'Sensitivity, Safety, and Antiluck Epistemology.' In J. Greco (ed.), *The Oxford Handbook of Skepticism*. Oxford: Oxford University Press.

Russell, Bertrand. 1948. *Human Knowledge: Its Scope and Limits*. London: Allen & Unwin.

Smith, Martin. 2009. 'Transmission Failure Explained.' *Philosophy and Phenomenological Research*, 79: 164–89.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press

—— 2010. 'What Else Justification Could Be.' *Noûs*, 44: 10–31.
—— 2012a. 'Two Notions of Epistemic Risk.' *Erkenntnis*, publ. online: 17 April, doi: 10.1007/s10670-012-9376-5.
—— 2012b. 'Some Thoughts on the JK-Rule.' *Noûs*, publ. online: 22 June, doi: 10.1111/j.1468-0068.2012.00866.x.
Smithies, Declan. 2012. 'The Normative Role of Knowledge.' *Noûs*, 46: 265–88.
Sutton, Jonathan. 2007. *Without Justification*. Cambridge, MA: MIT Press.
Vogel, Jonathan. 1990. 'Are There Counterexamples to the Closure Principle?' In M. Roth and G. Ross (eds), *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer.
Weiner, Matthew. 2005. 'Must We Know What We Say?' *Philosophical Review*, 114: 227–51.
—— 2007. 'Norms of Assertion.' *Philosophy Compass*, 2: 187–95.
Whiting, Daniel. Forthcoming. 'Nothing but the Truth: On the Norms and Aims of Belief.' In T. Chan (ed.), *The Aim of Belief*. Oxford: Oxford University Press.
Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

AIDAN MCGLYNN holds a Chancellor's Fellowship in the Philosophy Department at the University of Edinburgh. His research lies in the intersection of epistemology, philosophy of language, and philosophy of mind. He has papers appearing in *Analysis*, *Noûs*, and *Philosophical Studies*, and is currently writing a book on the knowledge first approach for Palgrave-Macmillan's *Palgrave Innovations in Philosophy* series. Philosophy, Dugald Stewart Building, 3 Charles Street, Edinburgh, EH8 9AD.

https://doi.org/10.1017/epi.2012.22 Published online by Cambridge University Press