A justification for excuses: Brown's discussion of the knowledge view of justification and the excuse manoeuvre

Clayton Littlejohn

Introduction

In *Fallibilism: Evidence and Knowledge*, Jessica Brown identifies a number of problems for knowledge-firsters. After considering her arguments, many will undoubtedly agree that things don't look great for knowledge infallibilists who want to free themselves from their view's (apparent) sceptical consequences by identifying a subject's evidence with her knowledge (E=K). I'm sure most readers will be sympathetic to her criticisms of other views associated with the knowledge-first movement, particularly her criticism of the knowledge view of justification (i.e., the view that our beliefs are justified iff they constitute knowledge (J=K).[1] Here, however, knowledge-firsters have the resources to mount a defence of their view.

Readers don't need to be told that J=K is a highly unorthodox view. While J=K still treated as a curiosity (or worse) in the literature, I think there's a perfectly good sense in which justified beliefs are justified by virtue of the fact that they constitute knowledge.[2] The recognition that we're not in a position to know makes it unreasonable to continue to believe.[3] This is the easy part of defending the normative significance of knowledge. The hard part is defending the idea that the conditions we couldn't have known about that prevent us from knowing prevent us

---

[1] This view is controversial even amongst knowledge-firsters. While Littlejohn (2013), McDowell (1998), Sutton (2007), Williamson (forthcoming), defend it, we can find alternative knowledge-first approaches to justification (or rationality) in Bird (2007), Dutant and Littlejohn (2021), Ichikawa (2017), and Rosenkranz (2021). I should also add that some knowledge-firsters see knowledge as a goal or aim even if they deny that it's a norm. See Ghijsen, Kelp, and Simion (2016),

[2] To be sure, there might be other more 'subjective' senses of justification just as we might distinguish between different readings of 'ought'. My impression is that most critics of J=K think that the view isn't true of any notion of justification. If they were pluralists who were content to admit that there were different readings of 'justified' that, say, connected to different readings of 'ought', they might acknowledge that there is *some* legitimate notion of justification that is identified with knowledge and say that they *choose* to discuss some other notion. Their attitude instead seems to be that we can see the view is false without considering different readings of the normative language.

[3] Suppose I concede that I don't know whether it was you who spoiled the surprise party I've been planning. Having conceded this, I cannot reasonably continue to be upset with you for doing that. Why not? It's not because doing such a thing wouldn't be upsetting, but precisely because I cannot reasonably continue to take *you* to be responsible for this. Acknowledged ignorance enjoins suspension.

from believing with justification. The case where, say, a BIV mistakenly believes that it has hands seems similar to the case in which a reasonable, mature, but regrettably ignorant agent uses force against someone who isn't a threat, trespasses on somebody else's land, or uses or sells something that doesn't belong to them. In each instance, I'd say that a normative standard was violated (e.g., a legal standard having to do with property rights or the rights of people not to be attacked or an epistemic standard having to do with knowledge and belief) and that the best the agent can reasonably hope for is an excuse. It seems like a mistake to apply justification to, say, my decision to drive off in a car that belongs to you when I have neither your consent for taking it nor an overriding reason to take it without your consent. Just as it seems to be a mistake to say that my decision to take your property without consent could be justified because I didn't understand what was happening, it seems like a mistake to say that my belief that I should take this car is one that I should have and should be guided by. Indeed, it seems like pretty much the same mistake. If (at best) I can be excused for taking what's not mine, it seems natural that we should draw the line between justification and excuse in such a way that the belief that led me to take what's not mine is (at best) one that I can be excused for relying on in practical reasoning and one that (at best) I can be excused for holding.

This struck me as a promising methodology. If we want to draw the distinction between justification and excuse in the right way, we should crib the best arguments they can find in the legal literature (Littlejohn 2009, 2012). If we do so, we'll end up saying the kinds of things that defenders of J=K have been saying in response to this sort of challenge:

> Intuitively, both the person [who knows they have limbs]
> and the BIV are equally justified in believing that they
> have limbs, and have the same strength of support for the
> claim that they have limbs (Brown 2018: 68).

We say that intuitions like this don't threaten our view and that those who use them as evidence against J=K are mistaken about the normative significance of the BIV's robust ignorance. We know that people will say that:

> … the excuse strategy … misidentifies the positive
> epistemic property had by key cases of non-knowledge
> belief. For example, some complain that a BIV doesn't
> need any excuse for her belief that she has hands; rather,
> her belief merits a positive epistemic evaluation which
> goes beyond mere excuse (Brown 2018: 79).

We of course agree that someone is mistaken about the significance of the 'positive epistemic property'. We think there are powerful arguments for thinking that the positive epistemic property that false beliefs can have is necessary for excusable norm violation and not something that indicates justification.

Most discussions of externalist views of justification that appear to clash with the intuition just mentioned say very little about the case that's been made for those views (e.g., about why people might think that knowledge is the norm of belief) or about the arguments for thinking that it matters to justification whether

beliefs are true.[4] I think it's strange that people think we can evaluate the merits of the claims of those who say that beliefs that violate J=K couldn't be justified but might be excused without addressing these arguments. While I wish Brown said more about why she didn't find these arguments persuasive, a virtue of her discussion is that it transcends bare appeal to intuition. She evaluates various attempts to give an account as to why the relevant beliefs should be seen as excused and finds each of them suffers from genuine problems.

While her discussion presents some knowledge-firsters with a formidable challenge, I will defend three points in response. First, J=K explains the accuracy of an intuition (mistakenly?) used against it. Second, there are persuasive arguments for thinking that in cases very similar to the BIV case, heroic but tragic agents lack justification for their actions but should be excused. This point can be used to argue that the conditions that differ between our situation and the BIV's situation *can* matter to justification and, if we assume some bridge principles, that some false beliefs about the agent's situation cannot be justified. This wouldn't be offered as a full defence of J=K, but it would support some of its more controversial elements. Third, we can make progress towards giving Brown the account of excusable norm violation she seeks. While I don't think that defenders of J=K owe an account if they can argue that some beliefs taken to be justified are really just excusably held, but it would be a nice account to have. Brown's critique of extant proposals pushes us in a helpful direction.[5]

1. Beliefs and believers

Those who defend J=K agree that belief is governed by the knowledge norm and that we shouldn't believe what we don't know. In personal communication, some have said they don't understand what we could even mean when we say justification is knowledge. I'll try to clarify.

You will think J=K is bizarre if you think 'justification' is just a term we use to talk about a thing that plays a certain role in the theory of knowledge (e.g., that of being a necessary but insufficient condition for turning true belief closer to knowledge), but I don't think justification is like this. What seems fundamental to justification is its normative role. Remember that justification is often characterised in normative terms as what's permitted, proper, or allowed. A theory of justification should fill in these blanks:

---

[4] This issue is discussed in length in Littlejohn (2012: 202-22, 2014) where I argue that given a standard way of drawing the line between justified and merely excused action, we must recognise that some beliefs (e.g., about things in the external world relevant to discharging our duties, about moral facts or principles) cannot be justified unless they are true.

[5] In thinking about whether this challenge can be met, I found Boult's (2017) discussion of a virtue-theoretic approach as an alternative to the more familiar dispositional approach of Lasonen-Aarnio (forthcoming) and Williamson (forthcoming) particularly helpful.

Justification: A may believe *p* if [_____] but A shouldn't
believe if [_____].

If you wanted to fill the blanks, you could say something here about truth, evidential support, or knowledge. We think we need to make reference to knowledge.

This understanding of the justification of responses in terms of norm conformity is more familiar in legal or moral settings where we might encounter norms like this:

Battery: You should not use force or violence on another.

If someone violates Battery, it's difficult to justify their behaviour. To do so, we'd need to identify reasons that conflict with the reasons to conform to Battery that would, in full awareness of the conflict, show that force or violence is permitted (e.g., in the defence of others).

It seems natural to characterise the justification of all responses (whether attitudes or actions) in terms norm conformity. There are substantive questions about *which* norms govern belief (e.g., whether they're concerned with knowledge, truth, or something else). It's fair to ask why we should think that there is a norm that's concerned with knowledge, for example. One quick and dirty argument goes like this. The views that take the *prospect* of knowing to be necessary for justification better account for our intuitions (e.g., about mundane cases, about lotteries, about preface cases, about negative self-appraisals) than any other alternative approach (Dutant and Littlejohn 2021). Moreover, the views that deny that truth is necessary for justification either fail to vindicate our intuitions about practical cases or sever the normative connection between the justification of belief and the justification of the things we'd be rationally compelled to do in light of our beliefs (Littlejohn 2012, 2014). J=K is the only view that does justice to both points.

What can we say in response to Brown's observation that, "Intuitively, both the person and the BIV are equally justified in believing that they have limbs" (2018: 68)? We can take a cue from the reliabilists and appeal to the distinction between *personal* and *doxastic* justification (Bach (1985) and Engel (1992)). We should distinguish between a person being justified when they've responded in some way and their response being justified. When we say that x is justified, we're saying that x does well when appraised or evaluated. When appraising or evaluating a response (e.g., a belief, decision, action), we check to see if norms are violated. This is a *deontic* evaluation. When appraising or evaluating a responder (e.g., a believer, a decider), we must decide what to think of them and whether they escape blame. This is a *hypological* evaluation. It would be a mistake to think that a *responder* doesn't hold up well when evaluated simply because their *response* doesn't hold up well—even the most virtuous person on their best day might fail to discharge their obligations. This is why we need excuses to inform our appraisals of persons when we're not convinced that their responses are suitable to their situation.

When we apply this distinction to the BIV case, defenders of J=K must say that our *response* is justified and theirs is not. (We don't violate the knowledge norm, but they do.) When it comes to the *responders*, however, we can say that our *responders* are both justified. We know that both have done everything that can be expected to try to acquire knowledge without violating knowledge norm.

In this way, then, J=K, predicts and explains why BIVs would be just as justified as we are. If we run the distinction between justified responders and responses, we get the appearance that J=K is counterintuitive. Once we're clear on distinction it's not clear that there's any clash between the view's verdicts and intuition.[6] It's hard to object to J=K on the grounds that it conflicts with intuitions if they explain why they're correct.

This approach to the BIV case paralleled an influential line on 'imperfect' self-defence (Moore: 1997). Contrast the case in which, say, Agnes knows that she's justifiably using force to defend herself from an aggressor with the case in which Bertrand with similar information, abilities, and mental states uses force against a merely apparent threat. Moore said that the use of force on a merely apparent aggressor is *not* justified. However, he added, we should say that the *agent* was justified. This distinction between the justification of agent and act, he said, was necessary for making sense of what's going on in cases of excusable wrongdoing. This description was supposed to do justice to the idea that someone like Bertrand might do everything that can be expected of him and still violate Battery while registering an important normative difference between a response that justifiably contravenes that norm (i.e., Agnes's defence against an aggressor) and unjustified violations of it (i.e., Bertrand's use of force against an innocent person).

Some readers will ask whether we should say that Bertrand's actions as unjustified? There's an important deontic difference between Bertrand's actions and Agnes's. In Agnes's case, if third-parties got involved, it should be to assist Agnes. In Bertrand's case, third parties who get involved should stop him. Given the plausible assumption that a person who isn't a threat retains the right not to be attacked and rights are correlative with duties, we get the difference that Bertrand is under a duty Agnes isn't and that Bertrand hasn't fulfilled this duty.[7]

---

[6] The distinction also appears early in the Gettier literature. See Lowy (1978).

[7] This fits with Fuller's take on the case when he says, 'Whether a wrongful actor is excused does not affect the rights of other persons to resist or to assist the wrongful actor. But claims of justification do' (2000: 760). This is, of course, controversial. Baron (2007), for example, takes the view that objective features of the situation don't have any bearing on justification. To my mind, this comes with a cost. As Zimmerman (2008) notes, if you take this line, you probably have to deny that the innocent person Bertrand assaults, shoots, stabs, etc. didn't have any right not to be assaulted, shot, stabbed, etc. Rights have correlative duties. On the view that Baron and Zimmerman defend, since there's no duty not to shoot the person who has the misfortune of looking like a dangerous person, there's no right here to violate. Explaining intuitions about justified interference and reparative duties gets tricky. This, in my view, is the clear sign that they've made a mistake. I don't think people critical of J=K typically think that they're embracing this kind of view. Most of them would probably complain that being shot by someone who mistakes them for a threat isn't just one of those bad things that happened. They'd probably think that, at the very least, they are owed something by the shooter and it's hard to see how

We can take two lessons from this. First, when epistemologists say that the 'positive property' in the BIV case is sufficient for justification, remember that this positive property (or something analogous to it) was present in Bertrand's case, too. He couldn't be excused if it weren't. And since his use of force isn't justified, that property isn't the mark of justified response. It's strange to think that the positive properties suffice for the justification of a response in the epistemic cases but not the practical cases. Second, if there's a principled link between the justification of belief and action according to which justified beliefs justify the responses that they rationalise, we can infer that because Bertrand's use of force wasn't justified, his belief that the use of force would be necessary *likewise* wouldn't be justified. Classifying Bertrand's beliefs about the need to use force as justified or the belief he ought to have seems to have implications for claims about who he gets to use force against that clash with important intuitions about differences between Bertrand's actions and Agnes's (Littlejohn 2012: 210).

2. Excuses, justifications, and (mere) blamelessness

Brown is sympathetic to the view that the defenders of J=K are wrong about the positive epistemic property in the BIV case. We can agree that she's actually right (and in agreement with the critics of J=K) in thinking that the *believer* might be justified even if J=K is right that the *belief* isn't justified. I suspect that some resistance to J=K stems from a disagreement about what excuses are and what they're meant to do.

One difficulty we face in reaching agreement on how to think about these cases is that we seem to disagree about the significance of norms and we seem to disagree about the connection between excuse and blamelessness. Let me briefly take up these points.

Brown says that she questions J=K but won't question the knowledge norm (2018: 71). She's not alone in thinking that justified beliefs might violate this norm even if this norm is genuine.[8] I find this view puzzling. Someone who insists that the BIVs beliefs are justified but acknowledges that there might be a knowledge norm that says, in effect, that we shouldn't believe unless we know is saying that the BIVs shouldn't believe what they justifiably believe. At this point, I struggle to understand what they're theorising. Suppose I say that while the shooting wasn't justified, Bertrand should be excused. Now imagine some philosophers thought that this is a mistake and a mistake that's important to correct. They agree, let's say, Bertrand shouldn't have pulled the trigger, but they insist that his shooting merited something that 'goes beyond mere excuse'. What could that be? I don't know. I've run out of normative categories.

I can imagine critics of J=K saying that there are objective and subjective readings of 'ought' and that the notion that they're really interested in is the

---

such a rich normative relation could obtain just because someone out there permissibly did something that was bad in some way.

[8] Bird (2007) and Ichikawa (2017) defend similar views.

subjective one. If that's the game, they shouldn't present themselves as *objecting* to J=K by pointing to counterexamples involving BIVs. They should instead *clarify* that they're concerned with subjective rightness, something that nobody thinks J=K is concerned with.

I'm not sure how to resolve the first obstacle to understanding (i.e., that of understanding the connection between norms and justification) except to say that I think that *if* it makes sense to distinguish objective epistemic rightness from some more subjective notion, J=K is concerned with the objective notion. If someone wants to *argue* that the normative isn't concerned with anything that has an objective dimension, we can argue about that and talk through the cases of imperfect self-defence or defence of others again. I hope it's clear that I defend J=K, in part, because I accept the view that a proper description of our duties can make reference to objective features of the situation.

There's a second obstacle to reaching agreement. Many epistemologists treat the excused and the blameless interchangeably. I don't think Brown explicitly endorses the view that A's Xing should be excused iff A is blameless when she Xs, but she describes attempts to explain why someone should be excused as attempts to show that they're blameless (Brown 2018: 80). While excuses remove blame, I would deny that mere blamelessness is sufficient for excuse. There are different ways of removing blame and excuses differ from other ways of doing so. I think it's enlightening to think about the differences and worry about discussions of excuses that don't consider the distinction between excuses and exemptions.

Critics of J=K often give as their examples of excuses things that either incapacitate an agent or interfere with the capacities they need to assume responsibility for their responses (e.g., drugs, brainwashing) or agents who cannot be held responsible because of infancy or insanity. If we think of excuse in terms of 'mere blamelessness', there's nothing wrong with thinking of these as excuses, but it would be a mistake to think that all excuses are like this. Almost none of Austin's (1956) examples were like this. Bertrand might have been a perfectly virtuous adult who hasn't been drugged or brainwashed. He might have had eight hours of sleep. It's outrageous to blame him in the way we'd blame some gunslinger who shoots someone for snoring. It's not right to say that he's blameless because he's like someone who isn't fit to stand trial. And it's a mistake to think that his action is just as justified as Agnes's because we should help her and stop him. We lose importance distinctions if we treat the excused and the merely blameless as interchangeable and we run the risk of treating something as 'more' than a mere excuse if we take our paradigmatic cases of excuse to be ways of showing that the agent cannot even bear responsibility for how they conduct themselves. Our view is that BIVs are similar in crucial respects to Bertrand. The real issue is about whether there's a knowledge norm in addition to norms like Battery.

To make progress in reaching a shared conception of what excuses do, let's think about their function. The most obvious two:

> F1. An excuse should remove responsibility and show that
> the individual isn't (fully) blameworthy.

> F2: The excuse can remove responsibility for a response
> that is wrong all things considered.

This much is uncontroversial. As Austin (1956) notes, we offer justifications by showing that something that was undesirable or wrongful was nevertheless something the agent had sufficient reason to do. We offer excuses when we want to show that the nonexistence of such reasons nevertheless doesn't reflect badly on the agent.

If we stop here, we won't make sense of the difference between excuses and exemptions. Exemptions (e.g., insanity, infancy, incapacitation) show that the agent lacked the rational capacities needed to bear responsibility and be answerable for their responses. Excuses (in the narrow sense) assume that the subject is answerable:

> F3: An excuse should not negate the claim that the agent
> is answerable for their responses.

This might seem puzzling. If an agent commits a wrong and this wrong isn't justified why wouldn't the agent be held responsible if they can bear responsibility? This is Gardner's answer:

> F4: An excuse should establish that the agent hasn't failed
> to meet our expectations.

We blame when we judge that someone doesn't meet our expectations. Excuses show that someone who can be held to account shouldn't be.

If we equate excusable violations of norms with blameless responses, we aren't doing justice to important distinctions between (a) wrongful responses that manifest virtue and rational excellence and (b) avoiding blame by lacking the capacities necessary for assuming responsibility. If we take the defenders of J=K to be saying that BIVs are somehow importantly similar to someone who is cut off from reality because they are brainwashed or insane, then the defenders of J=K would be wrong. This isn't our view. Our view is that the BIVs are like Bertrand and Agnes. When we know why they respond as they do, we don't see any indication that Bertrand or the BIV differs from the perfectly virtuous ideal reasoner or actor. And when we see that some of them violated norms, we can see why Bertrand and the BIV shouldn't be blamed for violating the norms that Agnes conformed to. Each lived up to our expectations showing insight, care, and good reasoning.

To sum up, if we want to show that the BIVs are excused for believing without justification, we need to do a number of things. We need to give show that there are norms we'd violate just by virtue of the fact that, say, we believe falsehoods. My argument for that is (roughly) that we shouldn't violate the rights of others and should be guided by our beliefs. It's impossible to see how both of these things could be true if false beliefs don't violate norms (Littlejohn 2012). From here, we argue that the violation of norms matters to justification since both have to do with permissibility, suitability, or propriety. There's no familiar normative kind that corresponds to the justified but forbidden response.[9] Once we clarify that there are

---

[9] Fans of the distinction between evaluating acts and agents or beliefs and believers can say that there are distinct evaluations here, but not between what's permitted by virtue of conforming to norms and what's permitted by virtue of being justified.

different ways of removing blame (i.e., by denying that someone can be held to account (offering an exemption) or by denying that someone that can be held to account should be (offering an excuse)), we see the kind of positive properties that someone like Bertrand needs for an excuse falls short of conferring justification. Given the obvious parallels between Bertrand's case and the BIV's case, they should be given the same treatment. My suggestion is that anyone critical of J=K should focus their efforts on undermining the case for believing there is a knowledge norm to violate. They shouldn't try to argue that the BIVs responses shouldn't be classified as excusable wrongs without challenging the claim that norms have been violated. Our substantive view of justification is unorthodox in epistemology, but our approach to excuses is textbook.

3. When and why we should be excused

In explaining why I think the BIVs beliefs aren't justified and why it's not contrary to intuition to say that they should be excused for violating the knowledge norm, I haven't offered an account of excusable norm violation. Instead, I've drawn on arguments familiar from the legal literature for thinking that defenders of J=K are entitled to say that the BIVs could be excused if they violate a norm by drawing parallels between the BIV's case and Bertrand's. I think Brown wants something different. She wants a positive story about why norm violations are excused.

I have concerns about this request. First, different excuses might have little in common. Fear, for example, might excuse violating a norm, but I doubt there's some simple formula that explains why fear and mistaken belief excuse when they do. We might need to try to give a long list of sufficient conditions and hope that that's satisfying. Second, we should be able to argue about whether something is known, right, justified, etc. when none of us have an account. I didn't have an account of excusable norm violation when I argued that Bertrand should be excused and I think that that's fine if our aim was simply to show that his use of force on someone who wasn't a threat wasn't justified. Still, an explanation as to why agents like Bertrand or like our BIVs should be excused would be nice to have. We might improve upon the accounts Brown criticises by thinking about her insightful objections to familiar accounts.

My thinking about excuses draws most heavily on Gardner's work. To start, we should locate excuses on Jareborg's ladder:

> … it is best of all if we commit no wrongs. If we cannot but commit wrongs, it is best if we commit them with justification. Failing justification, it is best if we have an excuse. The worst case is the one in which we must cast doubt on our own responsibility. When I say 'best' and 'worst' here I mean best and worst for us: for the course of our own lives and for our integrity as people (2007: 88).

Remember we said this about excuses:

> Expectation: A should be excused if she meets our
> expectations.[10]

If we can explicate this notion of meeting our expectations, we can move closer to the account we seek.

Expectations here aren't epistemic. My friends expect that I'll flake out and miss the next dinner, the next party, the next trip to Scotland, the next opening, etc., but this doesn't get me off the hook. It helps explain why I'm always back on it. Expectations are normative. They have to do with how well we respond to normative pressures. Idealising in various ways, we typically fail to meet normative expectations when we fall short of these ideals:

> Responsive: The agent is moved by the (possible) presence
> of right- and wrong-making features in suitable ways.
> Judicious: The agent makes the right trade-offs when there
> are conflicts between (possible) reasons.
> Risk-Sensitive: The agent handles well the risk of
> responding to (possible) right- and wrong-making
> features.

We meet expectations (or come close to doing so) if our responses align with the responses of someone who is responsive, judicious, and handles risk well and if the reasoning behind them doesn't show us to be unresponsive, less than judicious, or insensitive to risk.[11] Thus, in cases of imperfect defence of self or other, the relevant agent might not deserve an exemption (e.g., they might be a responsible adult), might not act with justification (e.g., they violate Battery), but might be excused if the response doesn't show that the agent's decisions are reckless, that they seem to be supported reasons that are (objectively speaking) stronger than the apparent reasons that seem to favour alternatives, and that the agent's reasoning doesn't reveal either the indifference or desire to engage in wrongdoing.

Let me note some potential differences between this approach and the dispositional account defended by Lasonen-Aarnio (forthcoming) and Williamson (forthcoming):

> Good Dispositions: A should be excused for violating a
> norm if she did what someone disposed to conform to this
> norm would do.

Someone might be unresponsive even if they do what someone disposed to a norm would do. The reasons matter. Remember that while a virtuous person, say, has good dispositions, they also wholeheartedly accept that the things that constitute good reasons are good reasons (Hursthouse 2016). We can use incentives to ensure that the less than fully virtuous are disposed to do what the virtuous person would. In theory, couldn't someone violate a norm whilst doing what someone disposed to conform to it would because doing this serves an ulterior and terrible motive? Brown

---

[10] As Gardner puts it, 'The gist of an excuse … is precisely that the person with the excuse lived up to our expectations' (2007: 124).

[11] I'll assume that Arpaly (2002) is right that the kind of responsiveness that matters is *de re* responsiveness, not *de dicto*.

imagines a case in which an agent tries to be rude but does something that's improbable on their evidence to have rude-making features (2018: 81). When someone's response shows a willingness to violate a norm, it's hard to see why they should be excused. I don't know if the good dispositions account is supposed to incorporate a responsiveness condition, but Brown's objection shows that excuses aren't appropriate if the relevant agent does what someone disposed to a norm would do if moved by bad reasons.

It's possible that a fully fleshed out dispositional account will include multiple dispositions to conform to multiple norms and further dispositions to make the right trade-offs (e.g., between fidelity and non-maleficence). It might also include a risk-sensitivity condition. If it doesn't, it should. Agnes promises to return Inge's book. She chose the less reliable courier for no good reason. Even if the book is delivered, she can be blamed for taking an unnecessary risk. This might seem like a difficult intuition for the dispositional view to account for because Agnes didn't violate the primary norm to keep promises, might have the disposition to keep promises, and might do what someone disposed to keep promises would do, but if risk-sensitivity is built into the disposition, this helps with this case and with Brown's examples where it seems an agent needlessly takes a risk of violating a norm whilst conforming to the primary norm and doing what someone disposed to so conform would do (2018: 87). If risk-sensitivity is built into our understanding of the promise-keeper, this should help with some of Brown's cases. If it's not built in, the expectation account might improve upon the dispositional account.

Let's apply the expectation approach to the case of knowledge and belief. Suppose that when someone is inquiring into whether *p* or has an interest in determining whether *p*, they ought to believe if they're in a position to know (K+) and should not believe if they're not (K-). If they are responsive, their doxastic dispositions are triggered by conditions that matter for knowing and they won't be moved to believe or suspend by the possibility that conditions that don't matter for knowing obtain. If they are risk-sensitive, the risks of violating K- or K+ will be properly weighed against each other. If we discover that, say, our beliefs are based on naked statistical evidence and that such beliefs cannot constitute knowledge, we ought to suspend. What if we just suspect that such evidence is responsible for our beliefs (e.g., I suspect my mother believes the defendant was guilty just because he was in the prison yard when most of the prisoners attacked a guard)? Here, it seems that the likelihood of the belief being based on such evidence determines whether I can be blamed for holding the attitude. If it's very likely not to be knowledge, I can be blamed for not suspending. If it's unlikely not to be knowledge, my suspension might be unreasonable.

This approach implies that if the risk of violating K- (or K+) sits below some threshold, this helps determine when the thinker should be blamed or excused for their attitudes. Brown is sceptical of this approach because:

> … it gives rise to the question of how probable it needs to
> be on one's evidence that one is in accord with the relevant
> norm for one to excusably violate it Intuitively, it is too
> weak to allow the excuse condition merely to be that it's

more probable than not on one's evidence that one meets the norm. For, this would excuse one when one's evidence only makes it slightly more probable than not that one meets the norm, say 0.52 probable. On the other hand, it would be too demanding to require for an excuse that it be highly probable on one's evidence that one meets the norm, say 0.99 (2018: 83).

This issue about how a responsible agent will respond to the risk of violating a norm is difficult for many reasons. We could imagine someone saying that some risks are sufficiently small and that they don't matter for this reason, but this seems unpromising. Instead, we should say that the threshold is determined by comparing the force of possible reasons that pull in different directions.[12] It seems natural to think that two acts might violate two different norms and both be wrong whilst differing in how wrong they are. They are wrong, we might say, to different degrees. If we can represent cardinally the degree of wrongfulness of, say, violating a norm that enjoins us to keep a promise to one that enjoins us not to cause injury, this would give us the tools we needed to determine whether some response minimised expected objective wrongfulness.[13] If an agent chooses an option that doesn't minimise this quantity, we can blame the agent.

If we apply this to the case of belief, we could start by asking whether we think violations of K- and K+ are wrong to some degree and whether they're wrong to the same degree. If they were wrong to the same degree, a responsible believer could believe if risk of violating were below .5. My inclination is to say that violations of K- are worse than violations of K+. On this view, we can be blamed for

---

[12] In some cases, norms can conflict. Think of the clash between non-maleficence and justice. K- and K+ aren't like this. There's no situation in which they both apply to an agent and, say, pressure her to believe and to suspend. Nevertheless, they pull in different directions and given uncertainty about which one applies, it still makes sense to approach this case in the way we might approach norms that can conflict in the same situation provided that we can compare the degree of wrongdoing that results from violating these norms.

[13] For discussion and defence of views of the subjective ought in the neighbourhood of this idea, see Lazar (2020) and Olsen (2018).

believing unless it's more likely than not that we know.[14] The threshold will be determined by the difference in the force of the reasons to conform to K- and K+.[15]

4. Conclusion

I've tried to show that the defenders of J=K can address Brown's concerns about the view's clash with commonly held intuitions and her concerns about the excuse manoeuvre. I hope that I've shown that the critics of J=K should address the arguments for thinking that knowledge is the norm of belief and concede that the defenders of J=K are well within their rights for thinking that if such a norm exists, the BIVs should be excused for violating it for pretty much the same reason we should recognise that an agent like Bertrand should be excused for violating Battery. I've noted some of the ways in which Brown's objections about the good dispositions account of excusable norm violation might point us towards a better understanding of why Bertrand or the BIVs excusably violate norms. None of this is intended to show that J=K is the right theory of justification, but I hope we can reach some agreement that it's not quite as implausible as some critics say. If I had more space and more imagination, I could try to address Brown's further concerns about the gradeability of justification, but my hope is that others will see the need to engage with Brown's objections and have something fruitful to say.

References

Arpaly, Nomy. 2002. *Unprincipled Virtue*. Oxford University Press.
Austin, J.L. 1956. A Plea for Excuses. *Proceedings of the Aristotelian Society* 57: 1–30.
Bach, Kent. 1985. A Rationale for Reliabilism. *The Monist* 68: 246-63.

---

[14] With sufficient patience, we might run through preface cases of varying sizes to get a sense of when we think it's reasonable for someone with n beliefs who is highly certain that some subset of those beliefs constitute knowledge to retain belief in each of the n-propositions. We won't ever locate the threshold, but the situation doesn't seem different from that of trying to determine the force of reasons to keep promises and refrain from killing strangers in deciding when it's irresponsible to drive somewhere (and thus increasing the chance of killing a stranger) to ensure that the promise is kept.

[15] See Dorst (2019) for a discussion of how the threshold can be set that focuses on a veritist value theory. If the above is right, there might be a further problem with the dispositional account. On Williamson's (2014) description of his unmarked clock cases, a thinker could know even when it's nearly certain on her evidence that she doesn't. If showing sensitivity to risk is part of meeting expectations, a thinker might believe what she knows (i.e., conforming to K+ and K-), have the disposition to conform to K- and K+ (if risk-sensitivity isn't part of the disposition), and respond in the way that someone disposed to conform to K- and K+ would. Intuitively, however, it seems unreasonable to believe in this case given the risk of violating K-. The problem is easily addressed by building risk-sensitivity into the relevant dispositions, but it's worth noting that the risk-sensitivity condition seems important for addressing a number of objections to the dispositional account.

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

Baron, Marcia. 2007. Excuses, Excuses. *Criminal Law and Philosophy* 1: 21–39.

Boult, Cameron. 2017. Epistemic Normativity and the Justification-Excuse Distinction. *Synthese* 194: 4065–81.

Brown, Jessica. 2018. *Fallibilism: Evidence and Knowledge*. Oxford University Press.

Cohen, Stewart. 1984. Justification and Truth. *Philosophical Studies* 46: 279-95.

Dorst, Kevin. 2019 Lockeans Maximize Expected Accuracy. *Mind* 128: 175–211.

Dutant, Julien and Littlejohn, Clayton. 2021. Defeaters as Indicators of Ignorance. In J. Brown and M. Simion (ed.), *Reasons, Justification, and Defeat*. Oxford University Press.

Engel, Mylan. 1992. Personal and Doxastic Justification in Epistemology. *Philosophical Studies* 67: 133-50.

Fletcher, George. 2000. *Rethinking Criminal Law*. Oxford University Press.

Gardner, John. 2008. *Offences and Defences*. Oxford University Press.

Ghijsen, Harmen, Kelp, Christopher, and Simion, Mona. 2016. Norms of Belief. *Philosophical Issues* 26: 375-92.

Hursthouse, Rosalind. 2016. Virtue Ethics. *Stanford Encyclopedia of Philosophy*.

Ichikawa, Jonathan. 2017. *Contextualising Knowledge*. Oxford University Press.

Lasonen-Aarnio, Maria. 2010. Unreasonable Knowledge. *Philosophical Perspectives* 24: 1-21.

Lasonen-Aarnio, Maria. Forthcoming. Virtuous Failure and Victims of Deceit. In J. Dutant and F. Dorsch (ed.) *The New Evil Demon*. Oxford University Press.

Lazar, Seth. 2020. Duty and Doubt. *Journal of Practical Ethics* 8, 28-55.

Lowy, Catherine. 1978. Gettier's Notion of Justification. *Mind* 87: 105-8.

Littlejohn, Clayton. 2009. The Externalist's Demon. *Canadian Journal of Philosophy* 39: 399-434.

Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press.

Littlejohn, Clayton. 2013. The Russellian Retreat. *Proceedings of the Aristotelian Society* 113: 293-320.

Littlejohn, Clayton. 2014. The Unity of Reason. In C. Littlejohn and J. Turri (ed.), *Epistemic Norms*. Oxford University Press.

Madison, B. 2014. Epistemological Disjunctivism and the New Evil Demon. *Acta Analytica* 29: 61-70.

McDowell, John. 1998. *Meaning, Knowledge, and Reality*. Harvard University Press.

Moore, Michael. 1997. *Placing Blame*. New York: Oxford University Press.

Olsen, Kristian. 2018. Subjective Rightness and Minimizing Expected Objective Wrongness. *Pacific Philosophical Quarterly* 99: 417-441.

Rosenkranz, Sven. 2021. *Justification as Ignorance*. Oxford University Press.

Sutton, Jonathan. 2007. *Without Justification*. MIT University Press.

Williamson, Timothy. 2014. Very Improbable Knowing. *Erkenntnis* 79: 971-999.

Williamson, Timothy. Forthcoming. Justifications, Excuses, and Skeptical Scenarios. In F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.