



Kelp, C. (2016) Justified belief: knowledge first-style. *Philosophy and Phenomenological Research*, 93(1), pp. 79-100.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/140967/>

Deposited on: 16 May 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Justified Belief: Knowledge First-Style

Christoph Kelp
Centre for Logic and Analytic Philosophy
KU Leuven

Abstract

Recent knowledge first epistemology features a number of different accounts of justified belief, including a knowledge first reductionism according to which to believe justifiably is to know [Sutton 2005, 2007, Littlejohn 2015a,b, Williamson 2000, 2010], a knowledge first version of accessibilism [Millar 2010] and a knowledge first version of mentalism [Bird 2007]. This paper offers a knowledge first version of virtue epistemology and argues that it is preferable to its knowledge first epistemological rivals: only knowledge first virtue epistemology manages to steer clear of a number of problems that its competition encounters.

Introduction

Knowledge first epistemology (KFE) has been on the rise in recent epistemological literature. For the purposes of this paper, I take its core theses to be the following. First, according to KFE, knowledge does not admit of reductive analysis in terms of justified belief. That is to say, no non-circular set of individually necessary and jointly sufficient conditions for knowledge including justified belief can be given. Second, KFE reverses the traditional direction of analysis in epistemology: rather than analysing knowledge in terms of justified belief, justified belief is analysed in terms of knowledge.¹

The main motivation for KFE is that the project of reductively analysing knowledge in terms of justified belief is marked by a series of failures. Since Gettier's seminal 1963 paper, there has been a flurry of proposed analyses of knowledge.

¹ Note that Timothy Williamson, the most prominent champion of KFE, also accepts a number of further theses, including that knowledge is a mental state in its own right and that belief is to be analysed in terms of knowledge [Williamson 2000, 2010]. In this way, Williamson's version of KFE makes epistemology part of the philosophy of mind. I'd like to flag that my project is not as ambitious as Williamson's. Rather, my version of KFE is restricted to the purely epistemological theses stated above. In fact, my views diverge from Williamson's on a number of issues. For that reason, I'd like to ask readers not to take the fact that I consider my view a version of KFE to mean that I consider it a view that Williamson would accept or even a view that is compatible with Williamson's.

Despite these efforts, no analysis is even widely accepted in contemporary literature. On the contrary, the debate is characterised by ever more sophisticated analyses and ever more sophisticated Gettier-style counterexamples.² Champions of KFE have taken this development to indicate that the traditional analytical project is a degenerating research programme and have proposed KFE as an alternative [e.g. Williamson 2000: 31, Bird 2007: 82]. I am inclined to agree. Moreover, for the purposes of this paper, I will simply assume that these considerations provide good enough reason to adopt KFE or at least to make the view worth exploring. I will not go into the motivations for KFE in any more detail.³

This paper focuses on KFE accounts of justified belief. To facilitate the discussion I would like to introduce a distinction between two variations of such accounts: strong and weak. An account of justified belief is strong if and only if it construes justified belief as requiring knowledge. An account of justified belief is weak if and only if it is not strong. Strong KFE accounts of justified belief are perhaps the most prominent ones in the literature. Their champions include Jonathan Sutton [2005, 2007], Clayton Littlejohn [2015a, 2015b] Alan Millar [2010], alongside Williamson [2000, 2010]. I will outline two strong KFE accounts of justified belief and argue that these accounts succumb, perhaps somewhat surprisingly, to a version of the Gettier problem (§1). §2 turns to weak accounts and, in particular, to a recent proposal due to Alexander Bird [2007]. I argue that Bird's account also remains unsatisfactory on two counts: first, it also falls prey to a version of the Gettier problem and, second, it fails to accommodate a plausible transmission principle for justified belief. In §3, I will outline a new kind of weak KFE account of justified belief, to wit, a knowledge first virtue epistemological one. Finally, I argue that this account steers clear of the problems on both sides (§4).

1 Strong Knowledge First Accounts of Justified Belief

This section introduces two strong KFE accounts of justified belief. The first has been defended by Sutton, Littlejohn and Williamson, whereas the second is due to Alan Millar. I will start with Sutton's, Littlejohn's and Williamson's view.

1.1 Knowledge First Reductionism

Sutton, Littlejohn and Williamson endorse an appealingly simple KFE account of justified belief. The idea here is that justified belief reduces to knowledge. They thus accept the following reductionist KFE account of justified belief:

² For overviews over the post-Gettier literature see [Slaght 1977, Shope 1983, Lycan 2006].

³ For a detailed argument that the knowledge first virtue epistemology is preferable to traditionalist virtue epistemology see [Kelp 2015b,d]. In a nutshell, key to my argument is that traditionalist versions of virtue epistemology encounter a version of what's become known as the creditworthiness dilemma [Lackey 2007, 2009, Pritchard 2008, Pritchard et al. 2010]. However, unlike standard versions of this dilemma my version of it has Frankfurt cases [Kelp 2009] on one horn and fake barn cases on the other. For arguments that traditionalist versions of virtue epistemology can escape other versions of the dilemma see [Kelp 2011, 2013a,b, 2014].

Knowledge First Reductionism (KFR). One justifiably believes p if and only if one knows p .

Crucially, of course, the direction of analysis proceeds from right to left: justified belief is analysed in terms of knowledge. In this way knowledge comes first.

It is easy to see that KFR entails that a belief is justified only if it qualifies as knowledge. Hence, KFR is a strong KFE account of justified belief in the sense specified above.

1.2 Knowledge First Accessibilism

A different strong account has been proposed by Alan Millar who endorses a knowledge first accessibilism about justified belief. To see how Millar's view works, it will be useful to first have a look at his account of knowledge. According to Millar, knowledge requires the exercise of an ability to know. For instance, perceptual knowledge requires the exercise of an ability to tell that something is thus-and-so from the way it appears. Conversely, when one exercises an ability to know that something is thus and so, one does come to know that it is thus-and-so.⁴ Moreover, one comes to know that it is thus and so in a specific way. For instance, when you exercise your ability to tell that the animal before you is a zebra from the way it looks, you *see that* it is a zebra, where seeing that something is thus-and-so is a way of knowing that it is thus-and-so.

With these points in play, let's move on to Millar's account of justified belief. While Millar does not offer an explicit definition of justified belief, there is a clear enough account implicit in the text. In what follows, I will try to briefly reconstruct this account.

According to Millar, "justified belief involves being in a position to justify one's belief" and that justifying one's belief is "understood in terms of possession of a reason to believe." [2010: 113, n.15] These remarks suggest a broadly accessibilist account of justification according to which justification requires possession of reasons that one must be in a position to access to ensure that one will be able to justify one's belief should the occasion to do so arise.⁵

What makes Millar's accessibilism distinctively knowledge first epistemological is his account of the reasons that are required for justified belief. Here is Millar:

[I]nstead of explaining the knowledge as, so to speak, built up from justified belief, we treat the knowledge as what enables one to be justified in believing. Knowing that the animal is a zebra through seeing that it is enables me to be justified in believing that it is.

[Millar 2010: 139]

⁴ This is because, according to Millar, the notion of exercise of ability is a success notion. One cannot exercise an ability to ϕ unless one ϕ s.

⁵ Prominent champions of accessibilism include Roderick Chisholm [1977], Laurence Bonjour [1985], William Alston [1989] and Robert Audi [1993]. The link between the ability to justify belief and accessibilism is defended e.g. in [Alston 1989: 236].

The thought here is that it is precisely the ways of knowing that furnish the reasons required for justified belief. When you exercise your ability to tell that the animal before you is a zebra from the way it looks, you come to know that it is a zebra in a specific way: you see that it is. The fact that you see that it is a zebra in turn constitutes a reason to believe that it is. What's more, it is the kind of reason furnished by ways of knowing that, according to Millar, is the kind of reason required for justified belief.

The question remains whether Millar's knowledge first accessibilism qualifies as strong in the above sense. To answer it, notice first that Millar takes the reasons required for justified belief that p to be ways of knowing that p . If so, it is easy to see that one's belief that p will be justified only if one knows that p . Millar's account qualifies as a strong KFE account of justified belief.

1.3 Gettier Strikes Back

One might think that once the traditional analytical project is abandoned in favour of a knowledge first approach to epistemology, we can finally put the Gettier problem to one side. Of course, this is true in the sense that we will no longer need to search for a condition on knowledge that deals with Gettier cases, or, at the very least, not one that can be specified without invoking the concept of knowledge. Unfortunately, however, the Gettier problem is not that easy to get rid of. There is a variation of the problem that arises for knowledge first epistemological accounts of justified belief. To see this, let's take a look at the following well-known case:

Fake Barns. You are driving through the countryside and take a look out of the window of your car. You see what appears to be a barn in the field and form a perceptual belief that you are looking at a barn. Unbeknownst to you, you are looking at one of the few real barns in an area peppered with barn façades that are so cleverly constructed as to be indistinguishable from real barns from your position on the road [Goldman 1976].

The familiar point about Gettier cases is that their protagonists lack knowledge. In the present case, you do not know that you are looking at a barn. But now notice that it is equally plausible that the beliefs of agents in Gettier cases *are justified*. This should not come as much of a surprise. After all, Gettier cases were intended as counterexamples to the justified true belief account of knowledge. If it weren't plausible that the relevant agents' beliefs are justified, Gettier cases wouldn't have made for convincing counterexamples to this view in the first place. Unlike the intuition of absence of knowledge, the intuition that justified belief is present in Gettier cases is of interest at least to KFE accounts of justified belief. It is an intuition such accounts need to accommodate.

It is easy to see that strong KFE accounts of justified belief encounter this version of the Gettier problem. After all, intuitively, Gettier cases are cases in which agents have justified beliefs that fall short of knowledge. Since, according to strong KFE accounts of justified belief, justified belief entails knowledge, and

since champions of KFE accept that agents in Gettier cases lack knowledge, they will have to accept the counterintuitive result that their beliefs are not justified.

In fact, the range of problematic cases for strong KFE accounts of justified belief extends significantly beyond Gettier cases. To see this, just consider the non-lucky counterparts of Gettiered agents who end up with false beliefs. (In what follows, I will also refer to these cases as ‘Counterpart cases’.) For instance, consider a variation of *Fake Barns* in which you look out of the window of your car a couple of minutes earlier and end up looking at a fake barn. When, in this case, you form the belief that you are looking at a barn, your belief will be false. Intuitively, however, it is also justified. Since knowledge is factive, however, your belief falls short of knowledge. As a result, champions of strong KFE accounts of justified belief will have to accept the counterintuitive result here too.⁶

1.4 The Blamelessness Response

Champions of strong KFE accounts are of course aware that their accounts run into this problem. In fact, they accept that agents in Gettier and Counterpart cases do not have justified beliefs. What is going on, they claim, is that the intuition of justification is not trustworthy in these cases. The beliefs of agents in Gettier and Counterpart cases are blameless but not justified. Those who have the intuition of justification mistake blamelessness for justification [Littlejohn 2015b, Millar 2010: 102-3, Sutton 2005, Sutton 2007: 33-34, Williamson 2010, 2014].

I believe that this response remains unsatisfactory for two reasons. First, it collapses a normative difference worth marking in epistemology. Second, the intuition of justification is not the only reason to think that agents in Gettier and Counterpart cases have justified beliefs.

To see the first point, consider the following two cases:

Insanity. You have gone insane. As a result, you form your beliefs in all sorts of crazy manners. When hearing the wind blow you think your long lost love is speaking to you, when the sky is red at sunset, you think that doom is impending, and so on.

Benighted Isolation. You are part of an isolated and benighted community the members of which share a common belief that thunderstorms indicate that their twenty-eared

⁶ In a similar vein, one might consider pressing the new evil demon problem [Lehrer and Cohen 1983, Cohen 1984] against strong KFE accounts of justified belief. This problem exploits the intuition that agents in sceptical scenarios have justified beliefs to argue against a variety of externalist accounts of justification. It is easy to see that strong KFE accounts of justified belief face the new evil demon problem. At the same time, it is also easy to see that the problem Gettier and Counterpart cases pose for strong KFE accounts of justified belief is of a kind with the new evil demon problem. Both venture to show that strong KFE accounts are too strong. The reason I favour Gettier and Counterpart cases to make this point is twofold. First, it avoids the charge that the difficulty is generated by problematic internalist intuitions. After all, the intuition that agents in Gettier and Counterpart cases have justified belief is shared by many externalists. Second, Gettier and Counterpart cases are less extravagant cases than the sceptical cases at issue in the new evil demon problem. In fact, they are a part of everyday life. If so, the intuition cannot be discredited solely on the grounds that the cases are too extravagant to generate trustworthy intuitions in the first place.

deity is about to scratch its largest left ear. Just now you are witnessing a thunderstorm and come to believe that the deity is about to scratch an ear.

In both of these cases you form a blameless belief. This is confirmed by the following two widely accepted theses concerning blamelessness: (B1) one is blameless for ϕ -ing if it is out of one's control that one ϕ s; (B2) one is blameless for ϕ -ing if one ϕ s in the light of good reason to believe that ϕ -ing is permissible [e.g. Haji 1998, Zimmermann 1997].⁷ *Insanity* is an instance of (B1). Here it is out of your control that you believe that doom is impending. You have gone insane. *Benighted Isolation* is instance of (B2). In this case, your belief about impending ear-scratching is formed in the light of good reason to believe that it is permissible. After all, you reasonably believe that thunderstorms indicate ear-scratching and that a thunderstorm has occurred. If so you have good reason to believe that it is permissible for you believe as you do. Moreover, since you base your belief on reasonable beliefs that provide you with good reason to believe that it is permissible, you believe in the light of this reason.

Crucially, there is a difference between agents in Gettier and Counterpart cases on the one hand, and agents in cases like *Insanity* and *Benighted Isolation* on the other. Agents in the Gettier and Counterpart cases form their beliefs in perfectly fine ways, ways that normally lead them toward epistemic goods such as true belief or knowledge. In contrast, agents in cases like *Insanity* and *Benighted Isolation* form their beliefs in highly problematic ways, ways that will normally not bring them on the path toward epistemic good. In this way, agents in the former cases are in a much stronger epistemic position than agents in the latter cases. Their beliefs have a connection with knowledge and truth that the beliefs of agents in cases like *Insanity* and *Benighted Isolation* lack altogether. To see this, compare, for example *Fake Barns* and *Benighted Isolation*. In *Fake Barns*, you are simply unlucky not to acquire knowledge on this occasion, whereas, in *Benighted Isolation* you fail to acquire knowledge because you are part of a community that is on the wrong epistemic track entirely.

This difference is a difference worth marking in epistemology. One reason for this is that marking this difference enables us to select the right course of action when working toward improvements of the agents' future epistemic performances. For agents in cases like *Benighted Isolation* and *Insanity* we will need to get agents to change the ways in which they form their beliefs. For agents in Gettier and Counterpart cases, in contrast, this is not necessary. Accordingly, when aiming for future improvements, a different strategy would seem more appropriate. For instance, we might consider engineering a more hospitable epistemic environment.

In fact, the difference is a difference worth marking as a *distinctively normative* difference in epistemology. It makes sense to evaluate the beliefs of agents in

⁷ See [Kelp 2015a] for a full account of blameless action that incorporates both (B1) and (B2). Could it be the case that one ϕ s in the light of good reason to believe that ϕ -ing is permissible whilst one should have known that ϕ -ing is impermissible? Perhaps. However, for present purposes, this issue is of little importance. Accordingly, I will simply assume that one cannot ϕ s in the light of good reason to believe that ϕ -ing is permissible when one should have known that it isn't.

Gettier and Counterpart cases positively and the beliefs of agents in cases like *Insanity* and *Benighted Isolation* negatively. To see that it makes sense to evaluate the beliefs of agents in Gettier and Counterpart cases positively, note that so doing will reinforce their ways of proceeding as epistemic agents, which is a good thing because agents will start to reap epistemic goods again as soon as they are back in epistemically more hospitable territories [cf. Bird 2007: 106]. In contrast, it makes sense to evaluate the beliefs of agents in cases like *Insanity* and *Benighted Isolation* negatively because so doing will discourage agents from continuing in their ways of proceeding as epistemic agents, which is also good thing given that no epistemic goods are to be gained by their way of proceeding. (Of course, it also makes sense grant these agents an excuse for proceeding in the way they do, but this point is by-the-by here.)

With these points in play, it can be argued that Gettier and Counterpart cases pose a problem for strong KFE accounts of justified belief independently of the intuition that agents in these cases have justified belief. To see this notice first that it is hard to see how champions of strong KFE accounts of justified belief can accommodate the aforementioned normative difference between agents in Gettier and Counterpart cases and agents in cases like *Insanity* and *Benighted Isolation*, at least on the present strategy. If the only thing that the beliefs of agents in Gettier and Counterpart cases have going for themselves is that they are blameless, then they are on par with the beliefs of agents in cases like *Insanity* and *Benighted Isolation*. The normative difference between agents in Gettier and Counterpart cases on the one hand and agents in cases like *Insanity* and *Benighted Isolation* is going to be lost. Since this difference is a difference worth marking, this means strong KFE accounts of justified belief collapse a difference worth marking.⁸

Let's move on, then, to the question of whether there is reason to think that beliefs of agents in Gettier and Counterpart case are justified, other than the contested intuition. I believe that the answer to this question is 'yes'. The beliefs of agents in Gettier and Counterpart cases have a couple of properties that are widely considered to be hallmark properties of justified belief, to wit, the property of enjoying a strong connection to truth and knowledge (P1), and the property of having positive normative status (P2). That the beliefs of agents in Gettier and Counterpart have P1

⁸ [Littlejohn 2015b] draws a distinction among blameless agents between agents who get an *excuse* and agents who get an *exemption*. In the case of exemptions the rational capacities of the blameless agent are absent or compromised. In contrast, agents who are blameless but have intact rational capacities get excuses [Littlejohn 2015b: 10]. Could champions of strong KFE accounts of justified belief venture to accommodate the difference between agents in Gettier and Counterpart cases on the one hand and agents in cases like *Insanity* and *Benighted Isolation* on the other by means of the distinction between exemptions and excuses? It is easy to see that the answer to this question is 'no'. While agents in cases like *Insanity* come out as getting an exemption on Littlejohn's view, agents in cases like *Benighted Isolation* will get an excuse. After all, their rational capacities are in perfect working order. Even with the proposed distinction in play, champions of strong KFE accounts of justified belief will be unable to accommodate the difference between agents in Gettier and Counterpart cases on the one hand and agents in cases like *Benighted Isolation* on the other.

was argued above. And while I haven't strictly speaking shown that these beliefs also have P2, I believe that I have come pretty close. After all, I have effectively argued that it makes sense to accord positive normative status to these beliefs. But since it plausibly makes sense to accord positive normative status to these beliefs only if they do indeed have positive normative status, it is also plausible that the beliefs of agents in Gettier and Counterpart cases have P2. The fact that the beliefs of agents in Gettier and Counterpart cases have properties that are widely considered to be hallmark properties of justified belief confirms the hypothesis that these beliefs are indeed justified.⁹

The fact that the blamelessness response collapses a normative difference worth marking in epistemology suggests that this response remains unsatisfactory. This suggestion is reinforced by the fact that we have more reason to believe that the beliefs of agents in Gettier and Counterpart cases are justified than just the contested intuition. In this way, there is not only evidence that the blamelessness response is unsuccessful, but also some positive reason to think that the beliefs of agents in Gettier cases are justified after all. If so, strong KFE accounts do face the Gettier problem.

2 Weak Knowledge First Accounts of Justified Belief

Perhaps, then, knowledge first epistemologists should start looking for a weaker account of justified belief, one that allows for justification in Gettier and Counterpart cases. One promising proposal on this front has been made by Alexander Bird. Let's take a look at it.

2.1 Bird's Knowledge First Mentalism

In his 2007 paper, Bird offers the following knowledge first version of a mentalist approach to justified belief:

KNOWLEDGE FIRST MENTALISM (KFM). If in world w_1 S has mental states M and then forms a [belief], that [belief] is justified if and only if there is some world w_2 where, with the same mental states M , S forms a corresponding [belief] and that [belief qualifies as] knowledge.

[Bird 2007: 84]

In order to get a good grasp of what exactly KFM amounts to, we also need to be clear on (a) which mental states we need to hold fixed when applying KFM and (b) what makes for a corresponding belief. The crucial point concerning (a) that needs to be kept in mind is that Bird accepts that there are factive mental states and that knowledge is a mental state in its own right that fits this description [Bird 2007: 82]. Thus, what we need to hold fixed in our applications of KFM is not only

⁹ It is easy to see that the beliefs of agents in cases like *Insanity* and *Benighted Isolation* do not have all of these hallmark properties of justified belief. In particular, they do not have P1 and P2. If so, we have reason to think that the beliefs of agents in these cases are not justified.

the relevant agent's non-factive mental states (states like experiences, apparent memories, intuitions and beliefs) but also the agent's factive mental states. In particular, we also need to hold the agent's knowledge fixed. Moreover, while we will always have to hold fixed the agent's mental states at the start of the belief forming process, sometimes we will even have to hold fixed larger stretches of the agent's mental history. Since Bird himself thinks that the relevant mental states can in every case be extended "to include a person's total mental history" [2007: 86], for the purposes of this paper, I will interpret KFM in this way. Concerning (b), Bird claims that a belief at a possible world corresponds to an actual belief if, first, there is "at most a minor variation in content" and, second, the belief is formed in the same way and that is to say it is "produced by the same mental dispositions and capacities" [2007: 87].

2.2 Transmission of Justification

There is reason to believe that KFM can improve on strong KFE accounts of justified belief in that it successfully predicts the presence of justified belief in cases like *Fake Barns* and its counterpart. After all, while in these cases your belief that you are looking at a barn falls short of knowledge there are possible worlds at which you have the same mental history as you do in the actual world and yet come to know that you are looking at a barn. One such world is a world at which there are no fake barns in the environment.

While KFM thus makes some progress on strong KFE accounts of justified belief, unfortunately, the view runs into problems with a different kind of Gettier case. Consider:

Inference. You are driving through the countryside and take a look out of the window of your car. You see what appears to be a barn in the field and form a perceptual belief that you are looking at a barn. Your perceptual belief is false. The structure you are looking at is a fake. A while later, you remember that in your logic course you learned that disjunction introduction is a valid rule of inference. By a competent application of this rule to your belief that you were looking at a barn earlier on you form the belief that either you were looking at a barn or the the number of stars in the universe is even. Your belief is true. The number of stars in the universe is indeed even.

Inference is another classical Gettier case. In fact, it's a variation of one of Gettier's original cases. Unsurprisingly, we find the characteristic pattern of intuitions here too. The belief you arrive at is justified and true but falls short of knowledge. Can KFM accommodate the intuition of justification here?

To answer this question we need to ask whether anyone with the same mental history could have acquired knowledge instead. Crucially, recall that knowledge is taken to be a mental state in its own right. This means that, when answering this question, we need to hold fixed that your premise belief falls short of knowledge. After all, if it did qualify as knowledge, your mental history would have been different. But now notice that you cannot come to know the conclusion of your

inference unless you know the premise. If so, at any world at which your premise belief that falls short of knowledge, you do not end up knowing the conclusion belief. Moreover, this is no matter what the content of the conclusion belief may be.¹⁰ As a result, there is no possible world at which you have the same mental history as in *Inference* at which you arrive at the same belief or at a corresponding belief that qualifies as knowledge. By, KFM, then your belief in *Inference* is not justified. KFM, too, faces the Gettier problem.

Again, the issue here is not just an intuitive one. Again, there is an underlying theoretical problem. To see this consider first the following plausible transmission principle for justified belief:

TRANSMISSION-JB. If one competently deductively reasons from p to q ; if one thereupon comes to believe that q ; and if one justifiably believes that p throughout, then one's belief that q is justified.¹¹

The underlying theoretical problem for KFM here is that KFM is incompatible with TRANSMISSION-JB. Given that competent deductive reasoning cannot generate knowledge from non-knowledge, it is easy to show that, on KFM, TRANSMISSION-JB fails in all cases in which the premise belief is justified but falls short of knowledge.¹² Here is how. Let S be any subject, b_1 be any justified premise belief that falls short of knowledge and b_2 any conclusion belief arrived at by competent deduction from b_1 . By KFM the worlds that are relevant to whether b_2 is justified in the actual world is a subset, Σ , of the worlds at which b_1 falls short of knowledge. Since competent deductive reasoning cannot generate knowledge

¹⁰ One might worry here that I am effectively assuming the truth of 'Counter-Transmission', the principle that competent deductive reasoning can generate knowledge only if the premise beliefs are known. This principle has recently come under attack [Warfield 2005, Klein 2008, Luzzi 2010]. Even if Counter-Transmission is false as a general principle, there will be many true instances of it. Moreover, all that I need for my argument at this point is that the relevant instances of Counter-Transmission are true. However, this is independently plausible. After all, it is uncontroversial that you are simply not in a position to know that the number of stars in the universe is even. If so, you can know that either you were looking at a barn or the number of stars in the universe is even via an inference from the proposition that you were looking at a barn only if you know that you were looking at a barn. And the same holds, *mutatis mutandis*, in cases in which you form "corresponding" conclusion beliefs, e.g. if you form a conclusion belief that you were looking at a barn or the number of stars in the universe is odd/prime/perfect, etc. via an inference from the proposition that you were looking at a barn.

¹¹ There is a live debate on whether TRANSMISSION-JB affords substantive restriction. See e.g. [Pryor 2000, 2004] and [Davis 2004, 2009] for defences of unrestricted transmission principles, and [Wright 2002, 2003] and [White 2006] for arguments that transmission principles afford restriction. For the purposes of this paper, I would like to bracket this issue. I will work with TRANSMISSION-JB as it is, without restrictions. To compensate for the presumptuousness of an unrestricted transmission principle, I will here consider TRANSMISSION-JB only insofar as it applies to uncontroversial cases like *Inference*.

¹² Again, this argument presupposes Counter-Transmission. Foes of this principle may note that my argument will still serve to show that TRANSMISSION-JB fails in all cases in which (i) the justified premise belief that falls short of knowledge and (ii) the relevant instances of Counter-Transmission hold.

from non-knowledge, b_2 falls short of knowledge at all members of Σ . Hence, by KFM, b_2 , while arrived at via competently deductive reasoning from a justified premise belief, will itself be unjustified. On KFM, TRANSMISSION-JB fails for any justified belief that falls short of knowledge. *Inference* is but one case in point. No surprise, then, that KFM makes the wrong prediction in this case.¹³

3 Knowledge First Virtue Epistemology

In what follows, I would like to develop an alternative knowledge first epistemological account of justified belief. Unlike the accounts discussed in the previous sections, my account is virtue epistemological. In particular, it follows a popular trend among recent virtue epistemologists in it takes the normativity of belief to drop out of an attractive general framework for the normativity of performances with an aim [Sosa 2007, 2011, 2015, Greco 2010, 2012]. According to this framework, performances with an aim can be evaluated along the following three dimensions: (i) is the performance successful, i.e. does it attain its aim?; (ii) is the performance competent, i.e., roughly, is it produced by an exercise of an ability to attain its aim?; (iii) is the performance apt, i.e., roughly, is it successful because competent? Moreover, my account also follows recent virtue epistemology in identifying knowledge with apt belief and justified belief with competent belief.

¹³ It may be worth briefly considering an argument Bird offers to the effect that KFM can accommodate transmission principles like TRANSMISSION-JB:

Knowledge is spread by such inferences because they preserve a reliable connection between true belief and corresponding fact. It is because of the possibility of spreading knowledge by inference that justification too can typically be transmitted in this way (even in cases where knowledge is only possible and not actual). Consider a straightforward case. Let S have a justified belief that p . Let S infer from the proposition that p the proposition that q , by a simple deduction. We need to show that S 's belief that q is also justified. Since S is justified in believing that p , by [KFM] there is a possible world w (which might or might not be the actual world) in which S has some corresponding knowledge. In this case let us assume that it is knowledge that p . Let it be that in some world w^* near to w (which might be w itself) in which S also knows that p , S infers that q by the aforementioned simple deductive inference. Since knowledge is transmitted by such inferences, in w^* S thereby comes to know that q . Since w^* is a possible world in which S comes to know that q by the same process by which S comes to believe that q in the actual world, then by [KFM] S 's belief that q in the actual world is justified.

[Bird 2007: 91]

This argument fails. To see why, let's focus on the last step. Bird claims that since w^* is a possible world in which S comes to know that q by the same process by which S comes to believe that q in the actual world, then, by KFM, S 's belief that q is justified in the actual world. In fact, this is not quite right. The fact that, at w^* , S comes to know that q via the same process by which S comes to believe that q in the actual world is not sufficient for S 's belief in the actual world to qualify as knowledge, at least not by KFM. It will only be sufficient given the additional assumption that S and his counterpart at w^* share the same mental history up until the point of belief formation. It is easy to see that once this additional condition is satisfied the above argument serves to show that KFM is as a matter of fact incompatible with TRANSMISSION-JB.

On the account I want to propose, then, a belief is justified if and only if it is competent. I also already said that this means, roughly, that it must have been produced by an exercise of ability to attain its aim. Of course, this immediately raises the following questions. What is the aim of belief? What is an ability? And what does it take to exercise an ability? In order to arrive at an account that is sufficiently well developed to be compared to the views discussed above, these questions need to be answered. In other words, the account needs to be made more precise. This is the task I will tackle in what follows. Since I want the account to drop out of an account that works for performances with an aim (or at least a certain type thereof) in general, I will start by developing an account of competent performances for performances in a certain kind of practice, which I will call a simple goal-directed practices (SGPs).

3.1 Simple Goal-Directed Practices

THE FRAMEWORK

A simple goal-directed practice is *goal-directed*. This means that it has a success condition, a condition under which the practice's goal is attained. A simple goal-directed practice is simple in the following sense. It features two types of particular, *moves* and *targets*, and a *designated relation*. A success in an SGP can be defined as a move that stands in the designated relation to the target. One example of an SGP is the following version of target archery, call it 'ARCH', in which moves are shots taken from a certain distance, targets are discs of a certain circumference and the designated relation is the hit relation. A success in ARCH is a shot that hits the target.

With the basic framework in play, I would like to offer the following accounts of (i) abilities to attain success in a given SGP, (ii) exercises thereof and (iii) competent moves in an SGP.

SGP ABILITIES

Roughly, SGP abilities are ways of move production that dispose its possessor to attain success in SGPs, at least provided he is in suitable shape and situational conditions.¹⁴

¹⁴ Since, on this account, abilities crucially involve dispositions to succeed, the account is similar to the accounts of ability offered by Sosa [2010, 2015] and Greco [2010, 2012]. Despite this similarity, my account differs from Sosa's and Greco's in a number of relevant respects. For starters, on my account abilities are relative to ways of move production. In contrast, neither Sosa nor Greco relativise abilities to ways of performing. Rather, what matters is that the agent attains success across nearby worlds at a high rate [Greco 2010: 77] or that the agent would succeed were she to try [Sosa 2015: 96]. There is reason to believe that, as a result, my account is preferable to both Sosa's and Greco's. To see why, consider an agent who has two ways of performing such as an archer who can shoot with his left or with his right. Suppose one and only one way of producing performances qualifies as an ability: the archer is right-handed, say. Even so, it might be that the agent is disposed to produce performances in the way that doesn't qualify as an ability. Perhaps our archer vowed never to shoot with his right again and has now taken to shoot with his left. In

For instance, to have the ability to hit the target in ARCH is to have a way of shooting that disposes one to hit the target in ARCH, at least when in suitable shape (e.g. awake, sufficiently concentrated, not too drunk, etc.) and situational conditions (e.g. normal winds, enough light, etc.).

SGP abilities need not just be one-trick ponies, as it were. Rather, they may span ranges of SGPs. For instance, a way of shooting that disposes one to hit the target in ARCH, in which the target is a disc with a certain circumference, may also dispose one to hit the target in ARCH', in which the target is a square with the same surface area, or in ARCH'', in which the target is a disc with an ever so slightly smaller circumference. In that case, one's ability to hit the target spans a range of SGPs, including ARCH, ARCH' and ARCH''.

Here is a more precise account of SGP abilities:

SGP ABILITY. One has an ability to attain success in a range, R , of SGPs and relative to conditions of shape, SH , and situational conditions SI if and only if one has a way of move production, W , such that, for any $S \in R$, using W in SH and SI disposes one to attain success in S .^{15,16}

this case, the agent will have the ability even though she will not be such as to attain success across nearby worlds at a high rate, nor will she be such as to succeed were she to try. Having taken the vow, the archer will not hit the target across nearby worlds with a high rate of success, nor would she hit the target were she to try. The reason for this is of course that were she to try, she would shoot with her left and fail to hit the target. Sosa and Greco's account of abilities are too coarse-grained. It is easy to see that the present account, which relativises abilities to ways of move production, avoids this problem. For more on how my account of abilities compares with standard virtue epistemological ones (such as Sosa's and Greco's) see [Kelp 2015c].

¹⁵ Note that there is a difference between genuine abilities and mere dispositions to attain some success or other. SGP ABILITIES does not appear to have the resources to accommodate this difference. I agree. In fact, I discuss this issue in more detail in [Kelp 2015c], where I follow [Millikan 2000: ch. 4] in arguing that genuine abilities must satisfy an etiological condition. Very roughly, the thought is that to qualify as an SGP ability, a way of move production must have been shaped through interaction with the environment, by a learning process or by natural selection, and, as a result must have become "tuned" [Millikan 2000: 63] to the environment. (Note that, as a result, my account differs in yet another respect from the accounts of Sosa and Greco (see fn.14) who do not countenance such an etiological condition on abilities.) I also argue that this etiological condition does some epistemological work. In particular, it serves to deal with clairvoyant-style cases [Kelp 2015c] and the lottery paradox [Kelp 2015e]. That said, the distinction between genuine abilities and mere dispositions does no substantive theoretical work in this paper. To see this, notice that any account of abilities will need some way of drawing the distinction between genuine abilities and mere dispositions. At the same time, none of the arguments I adduce presuppose any particular way of doing this. Moreover, in all the relevant cases discussed below, it is independently plausible that the agents' ways of belief formation qualify as genuine abilities rather than mere dispositions. For the purposes of this paper, then, the distinction between genuine abilities and mere dispositions can safely be set aside and the etiological condition can safely be assumed to be satisfied.

¹⁶ The dispositions involved in SGP abilities can be either *surefire* or *probabilistic* dispositions [Healey 1991, Suarez 2007]. Surefire dispositions are dispositions such that the (conditional) probability of manifestation given the stimulus condition and suitable conditions is 1. Probabilistic dispositions are dispositions such that the probability of manifestation given the stimulus conditions and suitable conditions is smaller than 1. An SGP ability features a surefire disposition just in case the probability of success given an exercise of the ability in suitable SH and SI is 1, otherwise it features a probabilistic disposition.

EXERCISES OF SGP ABILITIES

Exercises of SGP abilities are then defined in terms of uses of the ways of move production at issue in the ability:

SGP EXERCISE. One exercises an ability, *A*, to attain success for a range, *R*, of SGPs and relative to *SH* and *SI*, if and only if one has *A* and produces a move via the way of move production at issue in *A*.

What's important here is that unsuitable *SI* do not prevent one from using one's way of move production and hence from exercising one's ability, whilst unsuitable *SH* do. For instance, a skilled practitioner of ARCH who produces a shot whilst being too drunk does not exercise his ability to hit the target in ARCH. In contrast, when his shot is blown of target by a gust of wind, he may still have exercised his ability.¹⁷

COMPETENT SGP MOVES

Recall that I said that a competent performance is, roughly, a performance that is produced by the exercise of an ability to attain the performance's aim. With SGP ABILITY and SGP EXERCISE in play, then, we may think that we have all we need in order to offer a precise account of competent moves in an SGP. On reflection, however, there is reason to believe that competent moves require more than just the exercise of ability. To see this, note that you may not have the ability to hit the target in certain variations of ARCH, such as variations in which the target is moving quickly, randomly and discontinuously. Even so, you may have the ability to hit the target in ARCH and you may exercise it. However, when you do, it would seem that you have not produced a competent move in the version of ARCH you are engaging in.

These considerations suggest that we need an additional condition on competent moves in SGPs. Here is my proposal: For a move in an SGP to be competent it must (i) be produced by an exercise of an SGP ability and (ii) the SGP must be in the range of the ability exercised. Or, again, more precisely:

¹⁷ It is worth noting that, on the present account, the notion of exercise of ability is not a success notion. In this way, it differs from Millar's (see fn.4). In fact, this difference is one of two key differences that ultimately allow my account to solve the Gettier problem. The other difference is, of course, that my account replaces Millar's accessibilist approach to justified belief by a virtue epistemological one, which identifies justified belief with competent belief. Since the notion of exercise of ability is not a success notion, my account allows for competent but unsuccessful moves and, as a result, for competent but false beliefs. (This will become clear in due course.) Since my account is virtue epistemological, this means that there can be justified but false beliefs. And this is of course essential to standing any chance of accommodating the presence of justified belief in Gettier and Counterpart cases. For that reason I take it that the fact that, on my account, the notion of exercise of ability is not a success notion constitutes a key improvement on Millar's alternative. For a more detailed comparison of my account of abilities and their exercises and Millar's see [Kelp and Ghijzen 2016].

COMPETENT SGP MOVE. A move in a given SGP, S , is competent if and only if it is produced by an exercise of an ability to attain success within a certain range R of SGPs and relative to SH and SI such that $S \in R$.

3.2 Belief

INQUIRY AS AN SGP

With the general framework for SGPs and the account of competent move in an SGP in play, I would now move on to the application to belief. In particular, I would like to suggest that inquiry into specific whether questions can be viewed as an SGP, or better: as a set of SGPs (one for each question). The idea here is that the targets of inquiry are *correct answers*. For instance, the target of an inquiry into whether p is the proposition that p , if p is true and the proposition that not- p , if p is false. Moves in inquiry are *beliefs*. For instance, believing p constitutes a move in an inquiry into whether p , as does believing not- p . The designated relation in inquiry is the *knowledge relation*.¹⁸ A success in inquiry, then, is a belief that qualifies as knowledge (henceforth also ‘a knowledgeable belief’ for short). For instance, my belief that p stands in the designated relation to the target of an inquiry into whether p if and only if it qualifies as knowledge that p .

ABILITIES TO KNOW, ETC.

The corresponding accounts of abilities to know, their exercises and competent moves in inquiry are then straightforward:

ABILITY TO KNOW. One has an ability to know propositions in a range, R , and relative to SH and SI , if and only if one has a way of belief formation, W , such that, for any $p \in R$, using W in SH and SI disposes one to form knowledgeable beliefs that p .

EXERCISES OF ABILITIES TO KNOW. One exercises an ability, A , to know propositions in range R and relative to conditions SH and SI if and only if one has A and forms a belief via the way of belief formation at issue in A .

COMPETENT BELIEF. One competently believes that p if and only if one’s belief that p is formed by an exercise of an ability to know propositions in range R and relative to conditions SH and SI such that $p \in R$.

¹⁸ Note that taking knowledge to be the designated relation is what makes the account knowledge first epistemological. The framework can also be used to develop a traditionalist virtue epistemology. All we need to do is take the designated relation to be the correspondence relation between belief and true answer, where a belief corresponds to a true answer if and only if its content is identical to the true answer. In [Kelp 2015c] I argue that the resulting account of justified belief allows us to make progress on both traditional process reliabilist and alternative traditionalist virtue epistemological accounts.

JUSTIFIED BELIEF

Recall that I follow virtue epistemology in identifying justified belief with competent belief. This means that we get the following precise account of justified belief:

KFVE. One justifiably believes that p just in case one believes that p via an exercise of an ability to know propositions within R and relative to SH and SI such that $p \in R$.

KFVE is clearly a virtue epistemological account of justified belief. That said, it is also knowledge first epistemological. After all, it identifies justified beliefs with competent belief and analyses competent belief in terms of abilities to know. Knowledge thus enjoys explanatory priority over justified belief.

One nice feature of the account is that it serves to accommodate all of the hallmark properties of justified belief outlined in §1.4 above. KFVE features a kind of reliability condition on justified belief. As a result, it can accommodate justified belief's hallmark property of being connected to knowledge and truth. Moreover, KFVE identifies justified belief with a kind of competent performance. Since, according to the general account of performance normativity, competent performances in general enjoy positive normative standing, the same goes for justified belief.

KFVE offers an appealing knowledge first epistemological account of justified belief. In what follows, I will return to the problems that rival accounts in the knowledge first camp encountered. I will argue that KFVE steers clear of these problems, thus making the account look even more attractive.

4 KFVE in Action

Strong KFE accounts of justified belief ran into difficulties already with Gettier (and Counterpart) cases involving non-inferential knowledge. Bird, who offers a weak KFE account, could avoid Gettier cases involving non-inferential knowledge but ran into trouble with inferential Gettier cases. In what follows, I will look at how KFVE fares with respect to both kinds of Gettier (and Counterpart) case.

4.1 Non-Inferential Gettier Cases

To begin with, recall why strong KFE accounts of justified belief face the Gettier problem. According to such accounts, justified belief is sufficient for knowledge. Gettier cases provide reason to think that this is too strong, i.e. that one can have a justified belief that falls short of knowledge.

In contrast, we need not expect justified belief to entail knowledge on KFVE. Here is one way to see this. According to the general account of performance normativity associated with virtue epistemology, while apt performance requires competent performance, competent performance does not require apt performance. Given that my favoured form of virtue epistemology identifies justified belief with competent belief and knowledge with apt belief, we need not expect that, on this

form of virtue epistemology, justified belief entails knowledge. This means that KFVE is at least in the running for improving on strong KFE accounts of justified belief.

Of course, not being disqualified from getting Gettier and Counterpart cases right from the very start is one thing; actually getting them right is quite another. Let's return, then, to our toy example of the kind of Gettier/Counterpart case that causes trouble for strong KFE accounts of justified belief, to wit *Fake Barns*. Recall that here you are in Fake Barn County, see what appears to be a barn and thereupon come to believe that you are facing a barn. In this case, you acquire your belief that you are looking at a barn in a way, W , that does not qualify as such an ability relative to your actual situational conditions, SI_A . After all, in these situational conditions, using your way of belief formation does not dispose you to acquire knowledgeable beliefs about the presence of barns. That said, in more favourable situational conditions (SI_F), using W does dispose you to form knowledgeable beliefs about the presence of barns. For instance, it does so in certain situational conditions such that everything that looks like a barn actually is a barn. By ABILITY TO KNOW, W qualifies as an ability to know propositions about the presence of barns relative to SI_F . Since you form your belief via W , by EXERCISE OF ABILITY TO KNOW, you form your belief via the exercise of an ability to know propositions about the presence of barns relative to SI_F . Moreover, the target proposition is within the range of this ability as it concerns the presence of a barn. Since you form your belief via the exercise of an ability to know such that the target proposition is within the range of this ability, by COMPETENT BELIEF, your belief is competent and so, by KFVE, justified. All we need to do to turn the case into a Gettier case is stipulate that your belief is true as you happen to be looking at one of the few real barns. Alternative, we can turn the case into a Counterpart case by assuming that you structure you are looking at is a fake. Since KFVE's verdict of justification stands no matter how we fill in the details of the case, it comes to light that KFVE makes the right predictions in both the Gettier and the Counterpart version of the case.¹⁹

4.2 Inferential Gettier Cases

Let's move on to the kinds of inferential Gettier cases that Bird's weak KFE account of justified belief, KFM, falls prey to. Recall that the reason why KFM couldn't get inferential Gettier cases right is that, on KFM, TRANSMISSION-JB fails for any justified premise belief that falls short of knowledge. This, in turn, is because, KFM accepts both of the following two claims: (i) one's belief that q is justified only if some agent who shares one's mental history can come to know that q , whilst (ii) knowledge is a mental state in its own right. Anyone who accepts

¹⁹ It is easy to see the account will work equally well for standard non-inferential Gettier/Counterpart cases such as the so-called stopped clock case in which you acquire an intuitively justified belief about the time by taking an accurate and competent reading from a stopped clock you know to have worked highly reliable and you have no reason to think isn't doing so now. See [Kelp 2015d] for further details.

both of (i) and (ii) will have the problematic result. After all, given (ii), no agent who knows that p shares one's mental history. Moreover, since no one who doesn't know that p could come to know that q via competent deductive reasoning from p , no one who shares one's mental history can come to know that q via competent deductive reasoning from p and so, by (i) no one who shares one's mental history and believes that q via such reasoning will have a justified belief. In contrast, KFVE is wedded to neither (i) nor (ii).²⁰ Once again, this means that KFVE is at least in the running for doing better here.

Now it might be thought that so long as we have no reason to think that KFVE is incompatible with TRANSMISSION-JB, then we are entitled to assume that KFVE is compatible with it. If so, champions of KFVE are surely free to help themselves to TRANSMISSION-JB, at least until further notice. But if champions of KFVE may help themselves to TRANSMISSION-JB, then there is every reason to believe that they can handle inferential Gettier and Counterpart cases without much difficulty. After all, what's key to these cases is that the conclusion belief is arrived at via competent deductive reasoning from a justified but unknown premise belief. For instance, in *Inference* you come to believe a true conclusion—that you were looking at a barn or that the number of stars in the universe is even—from a justified but false premise belief—that you were looking at a barn. The status of the premise belief is not really what's problematic in these cases. Moreover, we have already seen how, according to KFVE, premise beliefs can be justified even if they are not known. For instance, the above treatment of *Fake Barns* shows how this works for the premise belief in *Inference*. If champions of KFVE are entitled to help themselves to TRANSMISSION-JB, then, they can easily secure the result that the conclusion beliefs of agents in inferential Gettier cases are justified also.

But perhaps not everybody is so lenient in granting that champions of KFVE are free to help themselves to TRANSMISSION-JB, just because we have no reason to think the two are incompatible. Some might doubt that the mere fact that there is no reason to think the two to be incompatible suffices for an entitlement to assume their compatibility. Others might worry that even if champions of KFVE are entitled to assume that TRANSMISSION-JB is compatible with their view, this doesn't mean that they are also entitled to help themselves to it. Such sceptics will maintain that before they are willing to grant that champions of KFVE are entitled to help themselves to TRANSMISSION-JB, they'd want to see some positive reason that this principle is indeed available to champions of KFVE. In what follows, I will address the worries of these sceptics. More specifically, I will argue that, modulo certain assumptions, KFVE entails TRANSMISSION-JB. Since the assumptions are clearly available to champions of KFVE, so is TRANSMISSION-

²⁰ Interestingly, KFVE, even in conjunction with TRANSMISSION-JB, is compatible with either of them. It is compatible with (ii) given that justification does not supervene on one's mental states and so one can come by a justified conclusion belief even when, holding fixed one's mental states, one could not have come to know the conclusion of one's deduction. It is compatible with (i) provided that knowledge is not a mental state in its own right and so one can be in the same total mental state even when one's premise belief is false.

JB.

4.3 The Transmission of Justification

To begin with, recall the relevant principle:

TRANSMISSION-JB. If one competently deductively reasons from p to q [J1]; if one thereupon comes to believe that q [J2]; and if one justifiably believes that p throughout [J3], then one's belief that q is justified.

Consider next the corresponding transmission principle for knowledge, i.e.:

TRANSMISSION-K. If one competently deductively reasons from p to q [K1]; if one thereupon comes to believe that q [K2]; if one knows that p throughout [K3]; and if one's deductive reasoning was successful [K4], then one knows that q .²¹

TRANSMISSION-K is available to champions of KFVE. After all, TRANSMISSION-K does not explicitly feature the concept of justified belief. Moreover, on KFVE, knowledge enjoys explanatory priority over justified belief and so TRANSMISSION-K does not presuppose this concept either. In this way, on KFVE, TRANSMISSION-K does not depend for its truth on facts about justified belief. If so, it should be uncontroversial that TRANSMISSION-K is available to champions of KFVE.²²

²¹ For defences of an unrestricted transmission principle for knowledge see e.g. [Williamson 2000] and [2004], for denials see e.g. [Dretske 1970, 2005] and [Nozick 1981].

²² Here is my response to those sceptical even of this claim. First, according to my account of knowledge [Kelp 2015d], the following is a necessary condition for knowledge:

- (1) One knows that p only if one competently believes that p in SI .

Moreover, when the relevant ability to know features a surefire disposition (fn.16), the converse holds also, i.e.:

- (2) One knows that p if and only if one competently believes that p in SI .

Second, consider the following bridge principles between TRANSMISSION-K and (2):

- (3) [K1] and [K2] hold if and only if one believes that q via the exercise of an ability to know propositions by deductive reasoning, $A_{Deduction}$, where q is within the range of $A_{Deduction}$.
- (4) [K3] and [K4] identify SI required by $A_{Deduction}$.

Third, consider these assumptions:

- (5) The disposition at issue in $A_{Deduction}$ is a surefire disposition.
- (6) There are no other SI required by $A_{Deduction}$ besides the ones mentioned in (4), i.e. [K3] and [K4].

(2) – (6) jointly entail (and thus validate) TRANSMISSION-K. To see this notice first that if the disposition at issue in some SGP ability is of the surefire kind, then the probability of success given an exercise of it is 1 and so the success is guaranteed, at least given that the SGP is within the ability's range and the SH and SI required by the ability are in place. As per (5), the disposition at issue in $A_{Deduction}$ is a surefire disposition. For that reason, the exercise of $A_{Deduction}$ guarantees knowledge that q , at least given that q is within the range of $A_{Deduction}$ and $A_{Deduction}$'s SH and SI are in place. Now, by (3), q is within $A_{Deduction}$'s range. Moreover, also by (3), $A_{Deduction}$'s SH are satisfied. After all, an exercise of an SGP ability requires that the SH are satisfied and so the fact

Note that in order to acquire inferential knowledge that q in accordance with TRANSMISSION-K, one must [K1] have arrived at one's belief that q via [K2] *competent deductive reasoning*. In this way, TRANSMISSION-K effectively features a competence condition. This raises the question as to what it takes to arrive at a belief that q via competent deductive reasoning. KFVE's account of competent belief serves to provide an answer to this question: one arrives at a belief that q via competent deductive reasoning if and only if one arrives at it via the exercise of an ability to know propositions by deduction such that q is in the range of this ability. Given that this is so, TRANSMISSION-K can be plausibly connected with KFVE's account of competent belief in the following way:

BRIDGE 1. [K1] and [K2] hold if and only if one believes that q via the exercise of an ability to know propositions by deductive reasoning, $A_{Deduction}$, where q is within the range of $A_{Deduction}$.²³

Finally, a claim that is not only available to but very much in the spirit of KFVE is that TRANSMISSION-JB is derivative from TRANSMISSION-K in the sense that the ability one needs to exercise in order to acquire a justified conclusion belief is the very ability to know propositions by deductive reasoning at issue in BRIDGE 1, i.e. $A_{Deduction}$. In particular, champions of KFVE are free to embrace the following bridge principle between KFVE, BRIDGE 1 and TRANSMISSION-JB:

BRIDGE 2. [J1] and [J2] hold if and only if one believes that q via the exercise of $A_{Deduction}$, where q is within the range of $A_{Deduction}$.

Note that BRIDGE 2 and KFVE entail that [J1] and [J2] hold if and only if one justifiably believes that q . It follows from this claim that if [J1], [J2] and [J3] hold, then one justifiably believes that q . And that's, of course, just TRANSMISSION-JB. In conjunction with TRANSMISSION-K, BRIDGE 1 and BRIDGE 2, KFVE entails TRANSMISSION-JB. Since all of TRANSMISSION-K, BRIDGE 1 and BRIDGE 2 are clearly available to champions of KFVE, so is TRANSMISSION-JB.

that, as per (3), $A_{Deduction}$ is exercised entails that $A_{Deduction}$'s *SH* are satisfied. Finally, given (4) and (6), $A_{Deduction}$'s *SI* must be in place. After all, by (4) and (6), [K3] and [K4] specify *the only SI* required by $A_{Deduction}$. It comes to light that, given (2), (3), (4), (5) and (6), believing that q via the exercise of an ability to know propositions by deductive reasoning, $A_{Deduction}$, where q is within the range of $A_{Deduction}$, is sufficient for knowledge that q . Since (2) – (6) are clearly available to champions of KFVE and since the jointly validate TRANSMISSION-K, TRANSMISSION-K is available to champions of KFVE also.

²³ Doesn't this mean that the resulting version of TRANSMISSION-K depends on the proposed account of justified belief? No. To see why not, just note that the proposed account of competent belief enjoys explanatory priority over the proposed account of justified belief. After all, according to KFVE, justified belief is analysed in terms of competent belief, not the other way around. Given that this is so, the fact that BRIDGE 1 employs the account of competent belief does not mean that it the resulting version of TRANSMISSION-K depends on the proposed account of justified belief.

It may also be worth noting that, on my account of knowledge, there is a competence condition on knowledge of exactly the form at issue in the right-hand side of BRIDGE 1 (see fn.22 above). As I argue in more detail in [Kelp 2015d], this competence condition is independently motivated. If so, BRIDGE 1 is not just ad hoc either. Rather, it provide an attractive way of connecting TRANSMISSION-K with an independently motivated condition on knowledge.

Unfortunately, however, there remains one fly in the ointment. The claim entailed by BRIDGE 2 and KFVE, that [J1] and [J2] hold if and only if justifiably believes that q , is too strong. To see this, suppose that one satisfies [J1] and [J2], i.e. one competently deductively reasons from p to q and thereupon forms a belief that q . By the above claim, it follows that one's belief is justified. But now suppose, furthermore, one's belief that p is entirely unjustified. In that case, one's belief that q is certainly not justified either.²⁴ If so, the claim entailed by BRIDGE 2 and KFVE is too strong.

Fortunately, this difficulty can be overcome. One attractive way of doing this is by restricting KFVE to epistemically basic beliefs, i.e. beliefs that, for their justification, do not depend on the epistemic status of other beliefs one holds. In line with a modest foundationalist requirement on the structure of justification champions of KFVE may add that, in the case of epistemically non-basic beliefs, the status of justification will be granted if and only if, in addition to the conditions specified in KFVE, the following proviso is satisfied: the beliefs on which they depend are also justified. It is easy to see that once the modest foundationalist requirement is in place, BRIDGE 2 and KFVE do not entail that one's belief that q is justified if and only if [J1] and [J2] hold. After all, in cases in which one comes to believe that q by deductive reasoning from p , the epistemic status of the former belief depends on the latter's and so the modest foundationalist requirement applies. Rather, once we have added the modest foundationalist requirement to TRANSMISSION-K, BRIDGE1, BRIDGE 2 and KFVE, the result that we get is that one's belief that q is justified if and only if [J1], [J2], and [J3] hold.

It comes to light that KFVE can be shown to entail (a plausible version of) TRANSMISSION-JB, at least on the assumption that TRANSMISSION-K, BRIDGE1, BRIDGE 2 and the modest foundationalist requirement are true. All of these assumptions are clearly available to champions of KFVE. In fact, I believe that they are also plausible. If so, there is every reason to believe that champions of KFVE can plausibly accept TRANSMISSION-JB. The worries of the above sceptics in this regard can be laid to rest.

Conclusion

It comes to light that KFVE improves on rival knowledge first accounts of justified belief in both the strong and the weak camp. Strong competitors struggle with Gettier cases. They cannot accommodate the intuition of justified belief in these cases. Moreover, their attempt at discrediting this intuition by holding that it arises from a confusion between justification and blamelessness also failed. On this line,

²⁴ Again, I am bracketing cases of failure of Counter-Transmission here (and in my response to this kind of worry below). Even if it turns out that Counter-Transmission does not hold with full generality, it remains overwhelmingly plausible that a restricted version of Counter-Transmission will have to be true. Given that this is so, we can leave the fine-tuning of the response I am about to give to the kind of problem at hand until we have a suitably restricted version of Counter-Transmission.

they are bound to collapse an important normative difference between agents in Gettier cases and agents in cases like *Insanity* and *Benighted Isolation*. Bird's weak alternative can improve on the shortcomings of its strong cousins in that it manages to successfully deal with non-inferential Gettier cases. However, in the end, Bird's view also doesn't manage to solve the Gettier problem. He stumbles over inferential Gettier cases. The reason for this is that his view is incompatible with a plausible transmission principle for justified belief. KFVE, in contrast, steers clear of the problems on both sides. It passes the right verdicts in both inferential and non-inferential Gettier cases. KFVE can also uphold the normative distinction between agents in Gettier cases and agents in cases like *Insanity* and *Benighted Isolation*, and the plausible transmission principle for justified belief is arguably available to its champions.²⁵

References

- Alston, W. 1989. *Epistemic Justification. Essays in the Theory of Knowledge*. Cornell University Press, Ithaca/NY.
- Audi, R. 1993. *The Structure of Justification*. Cambridge University Press, New York.
- Bird, A. 2007. Justified judging. *Philosophy and Phenomenological Research*, 74:81–110.
- BonJour, L. 1985. *The Structure of Empirical Justification*. Harvard University Press, Cambridge/MA.
- Chisholm, R. 1977. *Theory of Knowledge*. Prentice Hall, Englewood Cliffs, 2 edition.
- Cohen, S. 1984. Justification and truth. *Philosophical Studies*, 46:279–96.
- Davis, M. 2004. Epistemic entitlement, warrant transmission and easy knowledge. *Proceedings of the Aristotelian Society*, 78(1):213–245.

²⁵ *Acknowledgements*. Thanks to the audiences of the following conferences for helpful feedback on the material of this paper: 2012 EEN Meeting, Universities of Bologna and Modena (2012); CCPEA, Academia Sinica, Taipei (2012); 2013 Bled Philosophical Conferences, Bled (2013); *Epistemic Justification and Reasons*, University of Luxembourg (2013); *Neue Perspektiven der Epistemischen Rechtfertigung*, University of Dresden (2013), *Saving Safety*, Bonn (2013), *Yonsei Philosophy Summer Conference*, Yonsei University (2014); *Normative Epistemic Reasons*, University of Luxembourg (2014), *The Virtue Turn*, University of Taipei (2014); 2015 Bled Philosophical Conferences, Bled (2015). Thanks also to the Leuven Epistemology Group and two anonymous referees of this journal for their feedback on the paper. Special thanks to Harmen Ghijsen and Mona Simionescu for commenting on various versions of the paper. This work was funded by grants from KU Leuven's Special Research Fund (BOF) and Research Foundation Flanders (FWO).

- Davis, M. 2009. Two purposes of arguing and two epistemic projects. In Ravenscroft, I., editor, *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*. Oxford University Press, Oxford.
- Dretske, F. 1970. Epistemic operators. *Journal of Philosophy*, 67:1007–23.
- Dretske, F. 2005. The case against closure. In Steup, M. and Sosa, E., editors, *Contemporary Debates in Epistemology*. Blackwell, Oxford.
- Gettier, E. 1963. Is justified true belief knowledge? *Analysis*, 23:121–23.
- Goldman, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy*, 73:771–91.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press, Cambridge.
- Greco, J. 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research*, 85:1–26.
- Haji, I. 1998. *Moral Appraisability*. Oxford University Press, Oxford.
- Hawthorne, J. 2004. *Knowledge and Lotteries*. Oxford University Press, Oxford.
- Healey, R. 1991. *The Philosophy of Quantum Mechanics: An Interactive Interpretation*. Cambridge University Press, Cambridge.
- Kelp, C. 2009. Knowledge and safety. *Journal of Philosophical Research*, 34:21–31.
- Kelp 2011. In defence of virtue epistemology. *Synthese*, 179:409–33.
- Kelp 2013a. Extended cognition and robust virtue epistemology. *Erkenntnis*, 78:245–52.
- Kelp, C. 2013b. Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, 91:265–78.
- Kelp, C. 2014. Extended cognition and robust virtue epistemology: Response to Vaesen. *Erkenntnis*, 79:729–32.
- Kelp 2015a. Criticism and blame in action and assertion. *Manuscript*.
- Kelp, C. 2015b. Epistemic Frankfurt cases revisited (forthcoming). *American Philosophical Quarterly*.
- Kelp, C. 2015c. How to be a reliabilist. *Unpublished Manuscript*.
- Kelp, C. 2015d. Knowledge first virtue epistemology (forthcoming). In Carter, A., Gordon, E., and Jarvis, B., editors, *Knowledge First: Approaches in Epistemology and Mind*. Oxford University Press, Oxford.

- Kelp, C. 2015e. Lotteries and justification. *Manuscript*.
- Kelp, C. and Ghijsen, H. 2016. Perceptual justification: factive reasons and fallible virtues. In Mi, C., Slote, M., and Sosa, E., editors, *Moral and Intellectual Virtues in Western and Chinese Philosophy*. Routledge, London.
- Klein, P. 2008. Useful false beliefs. In Smith, Q., editor, *Epistemology: New Essays*. Oxford University Press, Oxford.
- Lackey, J. 2007. Why we don't deserve credit for everything we know. *Synthese*, 158:345–61.
- Lackey, J. 2009. Knowledge and credit. *Philosophical Studies*, 142:27–42.
- Lehrer, K. and Cohen, S. 1983. Justification, truth, and coherence. *Synthese*, 55:191–207.
- Littlejohn, C. 2015a. How and why knowledge is first (forthcoming). In Carter, A., Gordon, E., and Jarvis, B., editors, *Knowledge First: Approaches in Epistemology and Mind*. Oxford University Press, Oxford.
- Littlejohn, C. 2015b. A plea for epistemic excuses. In Dorsch, F. and Dutant, J., editors, *The New Evil Demon*. Oxford University Press, Oxford.
- Luzzi, F. 2010. Counter-closure. *Australasian Journal of Philosophy*, 88:673–83.
- Lycan, W. 2006. On the Gettier problem problem. In Hetherington, S., editor, *Epistemology Futures*. Clarendon Press, Oxford.
- Millar, A. 2010. Perceptual knowledge and recognitional abilities. In Pritchard, D., Millar, A., and Haddock, A., editors, *The Nature and Value of Knowledge*. Oxford University Press, Oxford.
- Millikan, R. 2000. *On Clear and Confused Ideas*. Cambridge University Press, Cambridge.
- Nozick, R. 1981. *Philosophical Explanations*. Oxford University Press, Oxford.
- Pritchard, D. 2008. Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, 58:437–47.
- Pritchard, D., Millar, A., and Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford University Press, Oxford.
- Pryor, J. 2000. The skeptic and the dogmatist. *Noûs*, (34):517–49.
- Pryor, J. 2004. What's wrong with moore's argument? *Philosophical Issues*, 14:349–78.

- Shope, R. 1983. *The Analysis of Knowing: A Decade of Research*. Princeton University Press, Princeton/NJ.
- Slaght, R. 1977. Is justified true belief knowledge?: a selective critical survey of recent work. *Philosophy Research Archives*, 3:1–135.
- Sosa, E. 2007. *A Virtue Epistemology. Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press, Oxford.
- Sosa, E. 2010. How competence matters in epistemology. *Philosophical Perspectives*, 24:465–475.
- Sosa, E. 2011. *Knowing Full-Well*. Princeton University Press, Princeton/NJ.
- Sosa, E. 2015. *Judgment and Agency*. Oxford University Press, Oxford.
- Suarez, M. 2007. Quantum propensities. *Studies in History and Philosophy of Modern Physics*, 38:418–38.
- Sutton, J. 2005. Stick to what you know. *Noûs*, 39:359–96.
- Sutton, J. 2007. *Without Justification*. MIT Press, Cambridge/MA.
- Warfield, T. 2005. Knowledge from falsehoods. *Philosophical Perspectives*, 19:405–16.
- White, R. 2006. Problems for dogmatism. *Philosophical Studies*, 131:525–57.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press, Oxford.
- Williamson, T. 2010. Knowledge first epistemology. In Bernecker, S. and Pritchard, D., editors, *The Routledge Companion to Epistemology*, pages 208–18. Routledge, New York.
- Williamson, T. Forthcoming/2014. Justifications, excuses, and sceptical scenarios. In Dutant, J. and Dorsch, F., editors, *The New Evil Demon*. Oxford University Press, Oxford.
- Wright, C. 2002. (Anti-)sceptics simple and subtle: G.E. Moore and John McDowell. *Philosophy and Phenomenological Research*, 65:330–48.
- Wright, C. 2003. Some reflections on the acquisition of warrant by inference. In Nuccetelli, S., editor, *New Essays on Semantic Externalism and Self-Knowledge*. MIT Press, Cambridge/Mass.
- Zimmermann, M. 1997. Moral responsibility and ignorance. *Ethics*, 107:410–26.