



Kelp, C. (2017) Knowledge first virtue epistemology. In: Carter, J. A., Gordon, E. C. and Jarvis, B. W. (eds.) *Knowledge First: Approaches in Epistemology and Mind*. Oxford University Press: Oxford, pp. 223-245. ISBN 9780198716310

This is the author's final accepted version.

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/140971/>

Deposited on: 15 May 2017

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

Knowledge First Virtue Epistemology

Christoph Kelp*

Introduction

It was once widely agreed that knowledge is justified true belief. Enter Gettier [1963] who adduced two cases that showed, to the satisfaction of nearly all researchers in the field, that one can have a justified true belief and yet fail to know. Since then, fixing the problem with the ‘JTB account’ that Gettier identified has been one of the main points on the epistemological agenda. Unfortunately, this task has proved to be rather difficult. While the literature features a whole host of proposals for such fixes, there is an equally impressive number of Gettier-style cases that show that these fixes meet with the same fate as the original JTB account. In fact, it is fair to say that 50 years of research have not produced an account of knowledge that is even widely agreed among researchers in the field.

That said, there are a number of live proposals in the literature. One of the most popular ones is virtue epistemology, which aims to offer accounts of knowledge and/or justified belief in terms of epistemic abilities.¹ On the other hand, some researchers take the continued failure of attempts to fix the JTB account to suggest that the JTB account is irredeemably flawed and have scrapped this point from their agendas. Instead, many of them have adopted a ‘knowledge first’ approach to epistemology, which aims to explain various core phenomena in epistemology in terms of knowledge.² This paper connects these two developments in epistemology. Its core aims are (i) to develop *knowledge first virtue epistemological* accounts of knowledge and justified belief and (ii) to show that these accounts compare favourably with their traditionalist cousins.³

In order to achieve this aim, I will first introduce virtue epistemology (VE) in its traditional guise and sketch an argument that VE does not succeed in offering

*Centre for Logic and Analytic Philosophy, KU Leuven, Kardinaal Mercierplein 2, BE-3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

¹ Champions of virtue epistemology include John Greco [e.g. 2010], Alan Millar [e.g. 2010], Duncan Pritchard [e.g. 2012], Wayne Riggs [e.g. 2002], Ernest Sosa [e.g. 2015], John Turri [e.g. 2011] and Linda Zagzebski [e.g. 1996].

² Champions of knowledge first epistemology include Alexander Bird [2007], Clayton Littlejohn [e.g. 2013], Alan Millar [e.g. 2010], Jonathan Sutton [e.g. 2007] and, most famously, Timothy Williamson [e.g. 2000].

³ See [Kelp and Ghijzen 2015, Kelp 2015c] for arguments that knowledge first virtue epistemology compares favourably with alternative knowledge first accounts of justified belief.

a fully satisfactory solution to the Gettier problem (§1). The main part of the paper will be concerned with developing the alternative knowledge first version of virtue epistemology. §2 introduces knowledge first epistemology and describes the general project. In §3 I will give a general account of the normativity of a certain species of performances, which I will then (§4) apply to beliefs as epistemic performances, thus arriving at my knowledge first virtue epistemological accounts of knowledge and justified belief. §5 argues that these accounts can offer a more satisfactory solution to the Gettier problem than its traditionalist rivals.

1 Virtue Epistemology

1.1 VE's Accounts of Justified Belief and Knowledge

The kind of VE I am interested in has most prominently been championed by Ernest Sosa [e.g. 2015] and John Greco [e.g. 2010]. It starts with a general account of performance normativity according to which performances with a goal can be assessed along the following three dimensions: success, competence and aptness. The idea is that a performance is successful if and only if it reaches its goal, competent if and only if it is produced by the exercise of an ability to reach the goal, and apt if and only if it is successful because competent.

It may be worth noting that this account of performance normativity leaves at least a number of key notions undefined. It does not tell us what it takes to *possess an ability*, what it takes to *exercise it* and what it takes for a success to be *because of* the exercise of an ability. Different champions of VE have offered different accounts of these key notions⁴, thus arriving at different versions of VE.

While there are disagreements between champions of VE, there are also a number of points that are widely agreed on by virtue epistemologists. First, they all take belief to be a kind of performance with a goal.⁵ As a result, the normative framework for performances can be applied directly to belief and yields an account of the normativity of belief. Second, virtue epistemologists have tended to take belief as type of epistemic performance that has truth as its aim. This gives us:

Successful Belief. A belief is successful if and only if true.

⁴ For instance, Sosa and Turri unpack the because relation in terms of the manifestation of an ability, whereas Greco and Pritchard opt for an explanatory salience alternative, and Riggs and Zagzebski seem attracted to a primitive account. Millar takes the notion of exercise of an ability to be a success notion, in the sense that one cannot exercise an ability to ϕ unless one ϕ s. In contrast, most other virtue epistemologists allow for unsuccessful exercises of abilities. There are also differences in just how abilities depend on conditions. Millar and Greco take abilities to be very strongly dependent on conditions in the sense that in sufficiently unfavourable conditions the ability is unavailable to the agent altogether. In contrast, most other virtue epistemologists opt for a slightly weaker version according to which the abilities are relative to conditions but abilities remain available to agents in even in fairly unfavourable conditions.

⁵ For a recent criticism of this idea see [Chrisman 2012], for responses see [Sosa 2015, Broncano-Berrocal 2015].

Competent Belief. A belief is competent if and only if it is produced by an exercise of an ability to form true beliefs.

Apt Belief. A belief is apt if and only if it is true because competent.

Next, champions of VE identify knowledge and justified belief with different normative standings of beliefs as performances:

VE-JB. One justifiably believes that p if and only if competently believes that p .

VE-K. One knows that p if and only if one aptly believes that p .

VE carries a significant amount of promise. Besides offering an appealing account of the normativity of belief as well as nicely motivated accounts of knowledge and justified belief, the view promises to solve a number of venerable epistemological problems, including the value problem⁶ and, most importantly for present purposes, the notorious Gettier problem.

1.2 The Gettier Problem

To get a grip on the Gettier problem, consider first a run-of-the-mill Gettier case:

Stopped Clock. Having come down the stairs, Mr. White looks at the grandfather clock in the hallway, sees that it reads 8:22 and on that basis comes to believe that it is 8:22. The clock has an outstanding track-record of functioning properly and Mr. White has no reason to think that it is currently not accurate. His belief is true. It is in fact 8:22. Unbeknownst to Mr. White, however, the clock has stopped exactly twelve hours ago.

Intuitively, Mr. White's belief that it is 8:22 is both justified and true but does not qualify as knowledge. Cases like this one thus serve to show that justified true belief is not sufficient for knowledge. They also raise the question as to what condition must be satisfied in addition to justified true belief in order to get knowledge. The Gettier problem is the problem of providing a satisfactory answer to this question.

Here is an appealing diagnosis of why agents in Gettier cases fall short of knowledge: it is pure luck that they end up with a true belief [e.g. Pritchard 2005]. More specifically, their beliefs are afflicted by a stroke of bad epistemic luck that puts them at great risk of acquiring a false belief. However, thanks to an additional stroke of good epistemic luck, the agent ends up with a true belief after all [Zagzebski 1994]. In *Stopped Clock*, the stroke of bad epistemic luck consists in the fact that the clock Mr. White takes a reading from is stopped. The stroke of good epistemic luck is that the clock stopped exactly twelve hours earlier, thus displaying the time correctly.

⁶ Very roughly, the value problem is the problem of explaining just how knowledge is more valuable than mere true belief or perhaps belief that falls short of knowledge. For more on the value problem see [Kvanvig 2003, Pritchard 2007]. For virtue epistemological approaches to the value problem see e.g. [Riggs 2002, Greco 2010, Sosa 2015].

This diagnosis of Gettier cases motivates a constraint on solutions to the Gettier problem. Any condition that, in conjunction with justified true belief, is said to be sufficient for knowledge must rule out the kind of luck at issue in Gettier cases.

With these points in play, we can now see why VE may, at first glance, appear to be particularly well positioned to solve the Gettier problem. According to VE, justified true belief is competent and successful belief, whereas knowledge is apt belief. But now recall that, in addition to competence and success, apt belief requires that the belief to be successful *because* competent. The reason why VE carries so much promise to solve the Gettier problem is that there is independent reason to believe that apt belief, i.e. belief that is successful because competent, is incompatible with the kind of luck at issue in Gettier cases.

How so? The answer is that there is independent reason to think that apt performance in general is incompatible with the kind of luck at issue in Gettier cases. To see this consider the following non-epistemic analogue of a Gettier case:

Diverted Arrow. Mr. Blonde, a highly competent archer, fires a competent shot at the target before him. On its way to the target the arrow is brought off its trajectory by a gust of wind. However, a moment later, a second gust of wind brings the arrow back on target with the result that it hits the bulls-eye.

In this case, Mr. Blonde's shot is both successful and competent. However, it is intuitively clear that it is not successful because competent and thus not apt.

Notice also just how appealing the parallel diagnosis is: the reason why Mr. Blonde's shot is not apt is that it is a matter pure luck that it hits the target. Moreover, we also find the familiar combination of a stroke of good luck and a stroke of bad luck. In *Diverted Arrow*, the stroke of bad luck consists in the arrow's being brought off target by a gust of wind. Mr. Blonde's shot is running a serious risk of missing the target. However, this risk is prevented from materialising by a stroke of good luck, which here takes the form of a second gust of wind that brings the arrow back on target.

It comes to light that there are non-epistemic analogues of Gettier cases in which the agents' performances are successful and competent but not apt. What's more, these non-epistemic analogues exhibit exactly the combination of two strokes of luck that are widely agreed to be the driving force in Gettier cases. Given that this is so, it may now seem that VE has what it takes to offer a successful solution to the Gettier problem. After all, champions of VE may now claim that Gettier cases are simply an instantiation of a more general phenomenon. Since it is independently plausible that such cases in general involve performances that are successful and competent but not apt, there is reason to believe that the same will hold true for its epistemic instantiations, belief. And since VE identifies knowledge with apt belief, there is independent reason to believe that beliefs in Gettier cases fall short of knowledge.

1.3 Fake Barn Cases

Unfortunately, there is a fly in the ointment. There is a second kind of Gettier case, besides the one exemplified by *Stopped Clock*. The most famous example was first stated by Goldman in 1976⁷:

Fake Barns. Mr. Pink is driving through the countryside and is currently looking out of the window of his car. He sees what appears to be a barn in the field and forms a perceptual belief that there is a barn in the field. Unbeknownst to Mr. Pink, he is 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 Mr. Pink's position on the road.

Just as in *Stopped Clock* before, the agent's (here: Mr. Pink's) belief is justified and true but falls short of knowledge. What's more, the same diagnosis as in *Stopped Clock* is plausible here: it is a matter of pure luck that Mr. Pink arrives at a true belief. In particular, Mr. Pink suffers from a stroke of bad luck in that he is in a part of the world in which fake barns prevail. This stroke of bad luck is cancelled out by a stroke of good luck: he is looking at one of the few real barns in the area.

The difficulty with *Fake Barns* is that VE cannot plausibly explain Mr. Pink's lack of knowledge in the same way as in *Stopped Clock*. The reason for this is that there are non-epistemic analogues of *Fake Barns* in which the agents' performances are very plausibly apt. Consider the following case by Pritchard [2008]:

Sabotaged Range. Mr. Orange, a skilled archer, is taking a shot a target at his shooting range. Unbeknownst to him, nearly all targets at the range are sabotaged to prevent any shot fired at them. The target he is shooting at is the only non-sabotaged target at the range.

It is intuitively clear that Mr. Orange's shot is not only successful and competent but also successful because competent and hence apt. If so, champions of VE cannot treat cases like *Fake Barns* as just another instantiation of a general phenomenon. VE's attractive account of standard Gettier cases does not extend to fake barn cases.

In fact, things are worse than this. It seems that the vast majority of non-epistemic analogues of *Fake Barns* are cases in which the agent attains aptness. Besides *Sabotaged Range*, think of a chef who successfully prepares a tasty omelette but happened to take the only salt shaker in which the salt wasn't replaced by sugar; or think of an artist who successfully produces a beautiful monochrome after having taken the only can in which the colour wasn't replaced by acid. In all of these cases, the agents' performances are intuitively apt. If so, the pressure on champions of VE to acknowledge that the Mr. Pink's belief in *Fake Barns* is apt also.

Fake barn cases constitute a significant obstacle for VE. Champions of VE must explain the absence of knowledge in such cases. As a first observation, note

⁷ It may be worth noting that Goldman attributes this example to Carl Ginet. While Ginet does describe the case in print, to the best of my knowledge he doesn't do so before his 1988 paper.

that most virtue epistemologists aim to offer a separate account for fake barn cases, i.e. one that diverges from the account they give for standard Gettier cases such as *Stopped Clock*. I'd like to point out that already the fact that fake barn cases are given a separate treatment constitutes a potential cost for VE. It will do so if there is an alternative account that offers a unified explanation of both standard Gettier cases and fake barn cases. After all, such an account will have the theoretical advantage of being simpler.⁸

What's worse, there is reason to believe that extant virtue epistemological accounts of *Fake Barns* remain ultimately unsuccessful. To see this, note first that these accounts venture to deal with fake barn cases by introducing a modal condition into the account of knowledge in one way or another.⁹ The problem now is that these modal conditions are too strong to accommodate intuitions of presence of knowledge in other cases. My own favourite argument to this effect involves epistemic Frankfurt cases, which are inspired by Frankfurt's famous 1969 cases.¹⁰

⁸ One might worry that this argument begs the question against VE. After all, champions of VE often argue that fake barn cases are structurally importantly different from standard Gettier cases. If so, by the lights of champions of VE, it may seem to make perfect sense to give them a separate treatment. I have two points to make by way of response. First, the question remains as to whether this structural difference cuts at the epistemological joints. If the reason these cases are taken to be different is simply that VE's treatment of Gettier cases doesn't extend to fake barn cases, the above complaint can hardly be said to beg the question against the view. (Compare: a champion of a 'no false lemma' (NFL) account claims that there are important structural differences between inferential and non-inferential standard Gettier cases and that, as a result, non-inferential Gettier cases afford separate treatment. If the reason for this claim is simply that the NFL account doesn't extend to non-inferential Gettier cases, this will do little to defuse a complaint of the kind made above.) Second, and more importantly, I don't take the above considerations to constitute an argument against VE, at least not by themselves. Rather, the thought is that if, in addition, there is an alternative view that can offer a uniform account and if all else is equal between the alternative account and VE, then there is some reason to think (i) that the alternative view is preferable and (ii) that the structural difference in question does not cut at the epistemological joints. (Compare: there is an alternative to the NFL account (perhaps a version of VE) such that it offers a uniform treatment of inferential and non-inferential Gettier cases and all else is equal between it and the NFL account. If so, (i) the alternative account is preferable to the NFL account and (ii) there is reason to think that the difference between inferential and non-inferential Gettier cases does not cut at the epistemological joints.) While I will develop just such an account in what follows, I will do so only in §4. As a result, the above argument is presently still awaiting completion.

⁹ For instance, Pritchard [2012] adds a safety condition on knowledge to deal with fake barn cases. Greco [2010] uses a modal condition on abilities (which in essence is a safety condition) to deal with fake barn cases. Sosa [Sosa 2015] uses a distinction between animal and reflective knowledge and argues for a modal condition (again, in essence a safety condition) on reflective knowledge.

¹⁰ Note, however, that Frankfurt cases are not essential to the argument. Any case in which the agent knows but doesn't satisfy the proposed modal conditions on knowledge will do. Since these conditions all amount to something like a safety condition (see fn.9), cases of unsafe knowledge are a good candidates for doing the work that I take Frankfurt cases to do. Crucially, the literature features a wide variety of such cases [e.g. Baumann 2008, Bogardus 2013, Comesaña 2005, Neta and Rohrbaugh 2004]. As a result, the argument may be successful even if Frankfurt cases eventually fail to do the trick (e.g. because they don't generate the required intuition of knowledge in readers). Finally, suppose that there is no one case that works against all of the proposed modal

Here is one example:

Frankfurt Clock. Mr. Blue's arch-nemesis, a powerful demon, has an interest that Mr. Blue forms a belief that it's 8:22 by looking at the grandfather clock in the hallway when he comes down the stairs. In order to achieve this, Mr. Blue's arch-nemesis is prepared to set the clock to 8:22 when Mr. Blue comes down the stairs. However, Mr. Blue's arch-nemesis is also lazy. He will act only if Mr. Blue does not come down the stairs at 8:22 of his own accord. Suppose, as it so happens, Mr. Blue does come down the stairs at 8:22. Mr. Blue's arch-nemesis remains inactive. Mr. Blue forms a belief that it's 8:22. It is 8:22. The grandfather clock is working reliably as always [Kelp 2009, 2015a].

Intuitively, Mr. Blue knows that it is 8:22 in this case. After all, we may assume, he has the ability to read the clock and forms a belief via an exercise of this ability. Moreover, the clock is actually functioning properly and the reading is accurate.

Recall that I pointed out that extant virtue epistemological accounts of knowledge introduce a modal condition to explain the absence of knowledge in fake barn cases. The problem they are now facing is that agents in epistemic Frankfurt cases, such as Mr. Blue above, will not satisfy any modal condition on knowledge that serves to explain the absence of knowledge in fake barn cases. The reason for this is that fake barn cases and epistemic Frankfurt cases have exactly the same modal profile. Any modal condition on knowledge that successfully predicts absence of knowledge in fake barn cases will also predict absence of knowledge in Frankfurt cases. Correlatively, any modal condition on knowledge that allows for knowledge in Frankfurt cases will not be able to explain the absence of knowledge in fake barn cases. In any case, accounts of knowledge that venture to accommodate the intuition of absence of knowledge in fake barn cases by means of a modal condition are in trouble. That's why extant virtue epistemological accounts of fake barn cases remain unsatisfactory.¹¹

2 Knowledge First Epistemology

While, at first glance, VE appeared to offer promising accounts of justified belief and knowledge, it turns out that VE continues to struggle with the Gettier problem. In what follows, I would like to develop a version of VE that avoids this difficulty. It differs from standard versions of the view in that it takes beliefs to be performance that aim not at truth but at *knowledge*. Accordingly, it takes the epistemic abilities required for knowledge and justified belief to be abilities to know.

It is easy to see that this version of VE will not yield a so-called 'reductive analysis' of knowledge. That is to say, it does not provide an analysis of knowledge in terms of a non-circular set of individually necessary and jointly sufficient

conditions. Even so, the argument may be successful provided that as there is at least one case that can be put to work against each of them.

¹¹ [Kelp 2015a] provides a more in-depth presentation of this argument. There I also argue in some detail the most prominent virtue theoretic accounts of knowledge in the literature, including Greco's and Sosa's, all succumb to this problem.

conditions for knowledge. After all, if the epistemic abilities required for knowledge are abilities to know, any account of knowledge in terms of epistemic abilities so understood will violate the non-circularity constraint.

I am not the first to abandon the ambition of offering a reductive analysis of knowledge. Others have taken the persistent failure of proposals in recent literature as a reason to do so as well. The perhaps most prominent alternative to the ‘traditional analytical project’ is knowledge first epistemology. Its founder, Timothy Williamson, characterises the view in the following passage:

“Knowledge first” is a slogan for epistemology that takes the distinction between knowledge and ignorance as the starting point from which to explain other cognitive matters [W1]. It reverses the direction dominant in much twentieth-century epistemology, which treated belief as explanatorily prior to knowledge, attempting to analyze knowledge as belief that meets further conditions, such as truth and justification [W2]. By contrast, a knowledge first epistemologist might treat believing something as treating it as if one knew it [W3].

[Williamson 2010: 208]

This characterisation of knowledge first epistemology contains three claims, which I have labeled W1, W2 and W3, respectively. Williamson’s characterisation already indicates just how ambitious his project is. In fact, it exceeds the boundaries of epistemology, at least as it is traditionally understood. By way of evidence, notice that W1 takes knowledge to be the starting point from which to explain other *cognitive* matters. Moreover, W3 anticipates one of Williamson’s famous theses, to wit, that knowledge is a mental state in its own right and that belief is to be analysed in terms of knowledge. Since accounts of mental states are traditionally taken to fall within the domain of the philosophy of mind and perhaps philosophical psychology, Williamson anticipates that he will go beyond the boundaries of epistemology as it is traditionally understood.

I’d like to flag that I will not follow Williamson in his pursuit of this ambitious project. The reason for this is that I am interested developing a viable *virtue epistemology*, one that is better than its traditionalist cousins. Since, for this purpose, the question as to whether knowledge is a state of mind in its own right is by-the-by, Williamson’s thesis concerning the philosophy of mind can safely be put to one side.¹²

¹² At the same time, knowledge first virtue epistemology is compatible with Williamson’s thesis that knowledge is a state of mind (henceforth also ‘KSM’). In fact, I recently discovered that Lisa Miracchi [2014] has developed a knowledge first virtue epistemology in which KSM takes centre stage. Like Miracchi’s, the account I develop below is compatible with KSM. However, my account neither entails KSM, nor does KSM play a key role in it. I take this to be a benefit of my account. After all, KSM remains a highly controversial thesis in the philosophy of mind. For that reason, it will be of some interest if it is possible to reap at least some of the epistemological benefits of knowledge first epistemology, without taking on the heavyweight commitments in the philosophy of mind that KSM brings with it. I’d also like to emphasise that I have developed my own version of knowledge first virtue epistemology independently of Miracchi’s.

At this stage one might naturally ask the following question: In what sense exactly is my account going to be knowledge first epistemological? The answer is that I do adopt the core epistemological claim of Williamson's characterisation of knowledge first epistemology, that is, the following version of W2:

Knowledge First Epistemology (KF). Knowledge first epistemology 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. In other words, rather than treating justified belief as explanatorily prior to knowledge, knowledge is taken to have explanatory priority over justified belief.

I submit that KF retains enough of the spirit of Williamson's project to qualify as a knowledge first view. Since it reverses the direction of analysis, it is incompatible with the traditional analytical project. Since it takes knowledge to enjoy explanatory priority over justified belief, there is an important respect in which knowledge comes first. For that reason, a virtue epistemology that satisfies KF can plausibly be regarded as a knowledge first virtue epistemology (KFVE). It is exactly this view that I will develop in more detail in the next sections.¹³

3 Simple Goal-Directed Practices

3.1 A Basic Framework

Let's start with a framework for simple goal-directed practices (SGPs). For a practice to be *goal-directed* is for it to have a success condition, a condition under which the practice's goal is attained. One very simple kind of goal-directed practice involves two types of particular, targets and moves, and a designated relation. The success condition of this kind of practice can be defined as obtaining if and only if a move stands in the designated relation to the target. In (a very simple version of) target archery, call it 'ARCH', the target is a disc with a set surface area, moves are shots taken from a set distance, and the designated relation is the hit relation. A success in ARCH is a shot that hits the target.

Practitioners of SGPs are move-producers. They may *attain success* in a given SGP. They do so if and only if they produce a move that stands in the designated relation to the target. Practitioners of ARCH are shot-producers. A practitioner of ARCH attains a success in ARCH if and only if he produces a shot that hits the target.

¹³ Or, to be more precise, I will develop a knowledge first virtue epistemological account of *basic first-order* knowledge and justified belief. That is to say, I want to set aside the issues of second-order knowledge/justified belief as well as of non-basic knowledge/justified belief, i.e. knowledge/justified belief which depends for its status as knowledge/justified belief on the epistemic status of other beliefs. My main reason for this is to keep the complexity of the paper manageable.

3.2 Abilities

Practitioners of SGPs may have the *ability* to attain success in a given SGP. You, the reigning world champion in ARCH, have the ability to hit the target, while I, a blind man, don't. Let's take a closer look at these abilities (henceforth 'SGP abilities').

SGP abilities involve *ways of move production*. Your ability to hit the target involves a way of shooting. More than one way of move production may be available to any one agent. For instance, you may shoot with your right hand or with your left hand. Of course, not all ways of move production will qualify as SGP abilities. While you may have the ability to hit the target in ARCH when shooting with your right hand, you may not have this ability when shooting with your left.

SGP abilities are *relative to conditions*. You may have the ability to hit the target in ARCH when shooting with your left when sufficiently concentrated, sober, not being shoved while releasing the arrow, when shooting in normal winds, when there are no jokesters destroying the target at the last second and so on. The conditions relative to which one may have an SGP ability may differ across different ways of move production. While you may have the ability to hit the target in ARCH when shooting with your left hand but only when you are completely sober, you may have this ability when shooting with your right even after a couple of beers.

SGP abilities involve *dispositions to attain success*. Dispositions are associated with (i) trigger (\mathcal{T}) and (ii) manifestation (\mathcal{M}) conditions. In the case of SGP abilities, the manifestations are, of course, attainments of success. In addition, the trigger conditions are uses of ways of move production. Furthermore, dispositions are (iii) relative to conditions (\mathcal{C}). In the case of SGP abilities these conditions are the ones to which the relevant ability is relative. To possess an ability to attain success in a given SGP, S , relative to conditions C , a practitioner must have a way of move production, W , such that he is disposed to attain S 's success (\mathcal{M}) when using W (\mathcal{T}) in C (\mathcal{C}). For instance, to have the ability to hit the target in ARCH relative to C you must have a way of shooting such that you are disposed to hit the target when using it in C . Finally, (iv) dispositions are widely regarded to be associated with counterfactual conditionals of the form were \mathcal{T} to obtain in \mathcal{C} , the possessor of the disposition would exhibit \mathcal{M} . This means that practitioners of SGPs who have an ability to attain success in a given SGP, S , relative to conditions C will be such that were they to use the way of move production at issue in this ability (\mathcal{T}) in C (\mathcal{C}), they would attain S 's success (\mathcal{M}). Those who have the ability to hit the target in ARCH relative to C will be such that were they to use the way of shooting at issue in this ability in C , they would hit the target.¹⁴

¹⁴ Properties (i), (ii) and (iv) above are among what Jennifer McKittrick [2003: 157] calls "markers of dispositionality", that is properties of dispositions that are widely agreed upon among contributors to the debate on dispositions. For what I take to be convincing motivations of property (iii) see e.g. [Mumford 1998, Sosa 2015]. Finally, it might be better to construe dispositions as being associated with a *probability of manifestation conditional on triggering in suitable conditions*, rather than with the counterfactual conditionals at issue in (iv) above. One advantage that this kind

A way of move production that qualifies as an ability to attain success in one SGP may or may not so qualify for another SGP. Consider another SGP, ARCH'. ARCH' is just like ARCH except that the target has a different shape than the target in ARCH. Instead of being disc-shaped, it is star-shaped. It may be that, in a certain set of conditions, *C*, a way of move production—shooting with your left hand, say—disposes you to hit the target in ARCH but not ARCH'. Your left-handed way of move production qualifies as an ability to hit the target in *C* relative to ARCH but not ARCH'. You are not that good an archer when shooting with your left hand. At the same time, it may also happen that a(nother) way of move production qualifies as an ability to hit the target in *C* for both ARCH and ARCH'. It may be that, in *C*, shooting with your right hand disposes you to hit the target not only in ARCH but also in ARCH'. Your right-handed way of move production qualifies as an ability to hit the target in *C* relative to ARCH and ARCH'. You are that good an archer when shooting with your right hand.

For that reason a way of move production may also qualify as an ability to attain success *across a range of SGPs*. A way of move production qualifies as an ability to attain success across a given range, *R*, of SGPs only if it disposes one to attain success when used for all SGPs in *R*. Your right-handed way of move production may constitute an ability to hit the target across ARCH, ARCH', ARCH'' (square-shaped), etc. Your left-handed way of move production may constitute an ability to hit the target across ARCH, ARCH'' but not ARCH', etc.

With these points in play, I would like to propose the following account of SGP abilities:

SGP Ability. One has an ability to attain success in a range, *R*, of SGPs and relative to conditions, *C*, if and only if one has a way of move production, *W*, such that, for any $S \in R$, using *W* in *C* disposes one to attain success in *S*.¹⁵

3.3 Exercises of Abilities

I want to suggest that exercises of SGP abilities are uses of ways of move production involved in SGP abilities. Or, more precisely,

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

of account of dispositions has for the application to abilities is that it can easily make sense of the fallibility of abilities, even in favourable conditions. After all, the relevant probability of manifestation may but need not be 1, even when the triggering conditions obtain in favourable conditions. See e.g. [Healey 1991, Suarez 2007] for more on probabilistic approaches to dispositions.

¹⁵ See [Kelp 2015b] for a more detailed account on SGP Abilities (as well as SGP Exercise and Competent SGP Move below). Note also that the full account of SGP abilities additionally features what I call a groundedness condition on abilities, which is a generalisation of Millikan's [2000: ch. 4] idea that genuine abilities are distinguished from mere behavioural dispositions by being relative to conditions in which they were acquired via learning or natural selection. That said, the grounding condition does no substantive theoretical work in this paper. Accordingly, for present purposes at least, it can safely be set aside.

It is important to note that placing the agent in conditions relative to which he does not have an SGP ability can have different effects on an agent's performances. Some such conditions will result in preventing an agent from using his way of move production. For instance, being too drunk, distracted, nervous, shoved while releasing the arrow and so on will prevent you from using the way of shooting that qualifies as an SGP ability relative to some (albeit different) conditions. I will henceforth refer to conditions that, when not satisfied, prevent the agent from using his way of move production as conditions of shape (*SH*). According to *SGP Exercise*, then, exercising an SGP ability requires that *SH* be satisfied.

In contrast, other such conditions do not prevent the agent from using his way of move production when not satisfied and so allow him to exercise his ability anyway. Suppose, for instance, that you fire a shot that would have hit the target had it not been for a jokester who destroys the target when the arrow is about to hit it. Even though your shot misses the target, you do get to produce a move via the way of shooting that qualifies as an SGP ability relative to some (albeit different) conditions. I will henceforth refer to conditions that, when not satisfied, do not prevent the agent from using his way of move production as situational conditions (*SI*). According to *SGP Exercise* exercising an SGP ability does not require that *SI* be satisfied.

3.4 Competent Moves

Competent moves in an SGP require the exercise of an SGP ability. When producing a shot in ARCH, your shot will be competent only if it is produced by an ability to hit the target.

However, a competent move requires more than the exercise of an SGP ability. To see this, let's return to the case in which you are the reigning world champion of ARCH. Suppose that you are currently engaging in ARCH_X in which the target changes its position discontinuously, randomly, and rapidly. Let's assume, as is plausible anyway, that you do not have the ability to hit the target in ARCH_X. You have no way of shooting that disposes you to produce successful moves in ARCH_X, no matter what conditions we may place you in. Suppose you take a shot using a way of move production that disposes you to hit the target in a range *R* of SGPs and relative to conditions *C*. Here you exercise your SGP ability to hit targets in range *R* and relative to *C*. However, that does not make your shot competent. The ability you exercise is the wrong ability for the SGP you are engaging in. (An even clearer example may be the following: Suppose you have a way *W* of producing layups in basketball that qualifies as an ability to score relative to some *C*. Currently you are standing at the three-point line and have two seconds to score a basket to win the game. Suppose you produce a shot via *W*, which, of course, doesn't even get close to the basket. On the present view, you exercise an ability to score. However, your shot is not competent. Again, the reason for this is that you are exercising the wrong ability to score here.) For a move to be competent, then, it must be a move in an SGP, *S*, that is within the range *R* for which your way of move production qualifies as an ability. Contrast the situation described

above with one in which you engage in an SGP that, we may assume, is within the range, R , of your SGP ability, but in which a jokester prevents the shot from being successful. Here you not only exercise an ability to hit the target, your shot is also competent.

There is thus reason to believe that, in order to produce a competent move in a given SGP, the SGP must be within the range (of SGPs) of the SGP ability exercised. The above considerations thus motivate the following account of competent moves:

Competent SGP Moves. A move in a given SGP, S , is competent if and only if it is produced by an exercise of an SGP ability to attain success in a range, R , of SGPs and relative to conditions, C , such that $S \in R$.

This gives us a detailed account of what competence amounts to in SGP moves. Let's now look at aptness.

3.5 Aptness

A SGP move is apt if and only if it is successful because competent. We already have precise accounts of success and competence. What is still needed is an account of because condition on aptness.

I will follow Sosa [2015] in adopting a competence manifestation account of this condition. That is to say, an SGP move is apt if and only if its success manifests its competence. The success of an SGP move manifests its competence if and only if the competent move is successful in the SI of the ability exercised. For instance, a shot you take in ARCH is apt if and only if it is competent, i.e. produced by the exercise of an ability to hit the target such that ARCH is in the range of this ability, successful, i.e. the shot hits the target, and the ability's SI are satisfied. If the shot is not competent, unsuccessful, or the SI are not in place, it will fall short of aptness.

I would thus like to suggest the following account of aptness:

SGP Aptness. A move in a given SGP is *apt* if and only if it is (i) successful, (ii) competent and (iii) the SI of the ability exercised are satisfied.

Before moving on, I'd like to make one last point, which will be of some importance in what follows. Note that aptness, competence and success can come apart in various ways. In particular, a successful move in an SGP need not be competent. And moves that are both successful and competent can still fall short of aptness. Both of these point are nicely illustrated by the case of ARCH. Consider a case in which you take a shot in ARCH. Suppose you are completely drunk when taking the shot. As a result, your shot is not competent. But now suppose that, by an incredible stroke of luck, your shot hits the target anyway. Your shot is successful without being competent. Consider, next, a case in which, unbeknownst to you, there is a magnet at the shooting range that brings any shot fired at it off target. You take a competent shot at that shooting range which is brought off target

by the magnet. Fortunately for you, a helper with a wind machine brings it back on target with the result that shot is successful after all. Your move is both successful and competent and yet falls short of aptness.

While it is true that there are SGPs such that aptness and success come apart in this way, things are different for other SGPs. Consider the practice of joining consenting adults in lawful marriage (henceforth LM). LM can be understood as an SGP. Targets here are couples of consenting adults, moves are tokens of ceremonies of a certain type, call it *c*, and the designated relation is effecting the lawful marriage of the couple. A success in LM is a token of *c* that effects the lawful marriage of the target couple. (The definitions of LM ability, exercise thereof, competent move and apt move are as expected.)

Crucially, there is reason to believe that one joins partners in lawful marriage only if one competently conducts a token of *c*. To see this, suppose you, the master of ceremony, fail to competently conduct a token of *c*. This may be for two reasons. You may fail to conduct a token of *c* altogether, say because you are too drunk and so do not manage to say enough of the relevant words. Or else the SGP may not be in the range of *c*, say because one of the partners is under age and so a different kind of ceremony is needed. In either case, the couple will not end up lawfully married. Since joining the couple in lawful marriage is the success condition of LM, this means that the success of LM will not be attained. If so, there is reason to believe that a move in LM will be successful only if competent.

What's more, there is also reason to believe that one joins partners in lawful marriage only if the *SI* of the ability exercised are satisfied. To see this, suppose that the situational conditions are not satisfied. For instance, one key situational condition for the ability to join partners in lawful marriage is that any caveat lodged against the marriage has been discharged. So suppose that this *SI* is not in place. There remains an undischarged caveat against the marriage. If so, again, the couple will not end up lawfully married, even if you have competently conducted a token of *c*. Since joining the couple in lawful marriage is the success condition of LM, this means that the success of LM will not be attained. As a result, there is also reason to believe that a move in LM will be successful only if the *SI* of the ability exercised are satisfied.

These considerations motivate the following condition on success in LM:

LM Success. A move in LM is successful only if it is competent and the *SI* of the ability exercised are satisfied.¹⁶

LM Success is an informative condition on success in LM. After all, it does not hold for successes in other SGPs such as ARCH and so captures relevant information about success in LM. At the same time, LM Success is unfit to constitute part of a reductive analysis of success in LM/joining consenting couples in lawful marriage. After all, it features the notion of competence, which, in turn, is unpacked

¹⁶ Other plausible candidates for types of performances such that success entails a competence-in-*SI* condition are certain types of intentional action, such as winking, waving, applauding and reading as well as Zagzebski's [1996] acts of virtue.

in terms of an ability to attain success in LM/join consenting couples in lawful marriage.

Now, let Δ be any SGP such that a move in Δ is successful only if it is competent and the *SI* of the ability exercised are satisfied. Since, trivially, a move in Δ will be successful only if it is successful, it follows that a move in Δ is successful only if successful, competent and the *SI* of the ability exercised are satisfied. However, by the relevant instance of *SGP Aptness*, we get the result that a move in Δ is successful if and only if apt. Since, by LM Success, LM is an SGP that features the relevant condition on success, we get:

LM Aptness. A move in LM is successful if and only if it is apt.

It comes to light, then, that while there are SGPs in which aptness and success come apart in the sense that one can attain success whilst falling short of aptness, in others aptness not only entails success, but success also entails aptness.

4 Knowledge and Justified Belief

I will now apply this account of SGPs, SGP abilities, exercises thereof and competent SGP moves to the epistemic case. To begin with, I would like to suggest that a relevant fragment of epistemic activity—*viz.* inquiry into specific whether questions (henceforth simply ‘inquiry’)—can be understood as an SGP. Or, to be more precise, it can be understood as a collection of SGPs, one for each question.

More specifically, my suggestion is that the targets of inquiry are *true answers*. For instance, the target of an inquiry into whether p is the true member of the set including the proposition that p and the proposition that not- p . Moves in inquiry are *beliefs*.¹⁷ 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*.^{18,19} A success in inquiry, then, is a belief that qualifies as knowledge (henceforth also ‘knowledgeable belief’ for short). In other words,

Successful Belief. A belief is successful if and only if it qualifies as knowledge.

¹⁷ A complete account would also countenance suspension of judgement as a type of move. Notice, however, that suspension of judgement is a second-order attitude. Given my aim of providing an account of first-order knowledge and justified belief (fn.13), the issue of suspension of judgement can safely be set aside.

¹⁸ Given that in inquiry targets are true answers and moves are beliefs, doesn’t this commit me to saying that true beliefs constitute successes in inquiry? No. To see this, imagine a variation of Eight-ball in which pockets must be designated and balls must be potted intentionally—flukes that find the designated pocket via a set of cushions do not count. Here, there is a clear sense in which the pocket is the target, but not any move that finds the pocket qualifies as a success.

¹⁹ I argue in some detail that knowledge rather than weaker epistemic standings (such as justified and/or true belief) is the goal of inquiry into particular whether questions in [Kelp 2014c]. For an argument that understanding why is not the goal of inquiry into such questions see [Kelp 2014b]. Compatibly with this, understanding may be the goal of inquiry into various phenomena and subject matters. The reason for this is that understanding is a form of knowledge (see [Kelp 2015d] for an argument).

What's more, we can apply the accounts of SGP abilities, their exercises, competent belief and apt belief from the previous section to the epistemic case. This gives us first the following accounts of abilities to know and their exercises:

Ability to Know. One has an ability to know propositions in a range, R , and relative to conditions, C , if and only if one has a way of belief formation, W , such that, for any $p \in R$, using W in C 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 C 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 C such that $p \in R$.

Apt Belief. A belief is apt if and only if it is (i) successful, (ii) competent and (iii) the *SI* of the ability to know exercised are satisfied.

Recall that traditional virtue epistemology identifies justified belief and knowledge with different normative standings of beliefs as performances. In particular, they identify justified belief with competent belief and knowledge with apt belief. I would like to follow traditional virtue epistemologists on this front. That is to say, I would like to offer the following accounts of justified belief and knowledge:

KFVE-JB. One justifiably believes that p if and only if one competently believes that p .

KFVE-K. One knows that p if and only if one aptly believes that p .

Let me say a few words to explain these accounts of justified belief and knowledge. They are superficially indistinguishable from their traditionalist rivals. At the same time, once they are properly unpacked, there are substantial differences between the two.

Starting with KFVE-JB, notice that, on my account, competent belief is analysed in terms of abilities to know. In true knowledge first style, then, KFVE-JB reverses the traditional direction of analysis and analyses justified belief in terms of knowledge, where knowledge has explanatory priority over justified belief.

What's more, whilst the accounts of abilities to know and their exercises do not figure directly in KFVE-JB they are of crucial importance to it. After all, they contribute to making the core notions in terms of which justified belief is defined more precise.

Let's move on to KFVE-K, then. First, it is of crucial importance to keep in mind that, unlike its traditionalist cousin, KFVE-K does not (and actually could not) offer a reductive analysis of knowledge. The easiest way to appreciate this

point is by taking note of the fact that apt belief requires successful belief and successful belief is knowledgeable belief. According to KFVE-K, then, one key condition a belief must satisfy in order to qualify as knowledge is that it qualifies as knowledge. KFVE-K is blatantly circular and thus unfit for the purposes of reductive analysis.

Second, notice that it follows from Successful Belief and KFVE-K that a belief is successful if and only if it is apt. After all, by Successful Belief, a belief is successful if and only if it qualifies as knowledge and, by KFVE-K, a belief qualifies as knowledge if and only if it is apt. It follows that a belief is successful if and only if it is apt.

As we have seen in the previous section, in and of itself, this is no cause for concern. There are other cases of SGPs in which a move in the relevant SGP is successful if and only if apt. The practice of joining consenting adults in lawful marriage (LM) is a case in point.

What's more, on reflection, it is independently plausible that in the case of inquiry, moves are successful if and only if apt. In fact, an argument parallel to the one that made this point for LM will do the trick here too. Here goes.

To begin with, there is reason to believe that knowledge requires competent belief. To see this, suppose you, an epistemic agent, form a belief that you are facing a barn, say, that falls short of being competent. This may be for two reasons. You may fail to exercise an ability to know, say because you form your belief on the basis of a coin toss. Or else the proposition may not be in the range of your ability, say because it has been produced by an ability to recognise colours. In either case, you end up not knowing that you are facing a barn. Since knowledge is success in inquiry, this means that success in inquiry will not be attained. If so, there is reason to believe that a move in inquiry will be successful only if competent.

What's more, there is also reason to believe that knowledge requires that the *SI* of the ability exercised are satisfied. To see this, suppose that the situational conditions are not satisfied: you are in fake barn county, are currently looking at a fake barn etc. If so, again, you end up not knowing that you are facing a barn. Since knowledge is the success condition of inquiry, this means that success in inquiry will not be attained. As a result, there is also reason to believe that a move in inquiry will be successful only if the *SI* of the ability exercised are satisfied.

Just as in the case of LM before, these considerations motivate the following condition on success in inquiry:

Inquiry Success. One's move in inquiry is successful only if it is competent and the *SI* of ability exercised are satisfied.

But now recall that we saw in the last section that, for any SGP such that this condition holds, it follows that a move in that SGP is successful if and only if apt. Since we have seen that there is independent reason to think that Inquiry Success is true, this means that it is independently plausible that a belief is successful if and only if it is apt.

Finally, it may be worth noting that, just as in the case of LM Success, Inquiry Success is an informative condition on success in inquiry and yet is not suited to constitute part of a reductive analysis of success in inquiry/knowledge. It is informative because the parallel condition does not hold for successes in other SGPs such as ARCH and so captures relevant information about success in inquiry. It is unfit to constitute part of a reductive analysis of success in inquiry/knowledge because it features the notion of competence, which, in turn, is unpacked in terms of an ability to attain success in inquiry/knowledge.

5 The Gettier Problem

With KFVE in play, I would now like to take a closer look at how the view deals with the Gettier problem. One might think that once the traditional analytical project is abandoned in favour of a knowledge first approach to epistemology, the Gettier problem does not arise in the first place. While this is true in a sense, there is a variation of the problem that arises for knowledge first epistemology. To see this, let's have another look at our two Gettier cases:

Stopped Clock. Having come down the stairs, Mr. White looks at the grandfather clock in the hallway, sees that it reads 8:22 and on that basis comes to believe that it is 8:22. The clock has an outstanding track-record of functioning properly and Mr. White has no reason to think that it is currently not accurate. His belief is true. It is in fact 8:22. Unbeknownst to Mr. White, however, the clock has stopped exactly twelve hours ago.

Fake Barns. Mr. Pink is driving through the countryside and is looking out of the window of his car. He sees what appears to be a barn in the field and forms a perceptual belief that there is a barn in the field. Unbeknownst to Mr. Pink, he is 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 Mr. Pink's position on the road.

Recall that, intuitively, Mr. White's and Mr. Pink's beliefs do not qualify as knowledge. Crucially, it is intuitively equally plausible that their beliefs *are justified*. Knowledge first epistemological accounts of justified belief will have to accommodate the intuition that beliefs of agents in Gettier cases are *justified*. While the original version of the Gettier problem does not arise within knowledge first epistemology, a variation of it does.

What's more, even though knowledge first epistemological accounts do not offer a reductive analysis of knowledge, they may offer informative necessary conditions on knowledge. If so, the question remains whether the necessary conditions countenanced are satisfied in Gettier cases. If not, knowledge first accounts may even be able to predict the absence of knowledge here.

In what follows, I will look at how KFVE deals with each kind of Gettier case. In particular, I will argue that KFVE successfully predicts both the presence of

justification and the absence of knowledge in both kinds of case. I will start once more with standard Gettier cases.

5.1 Standard Gettier Cases

First, in *Stopped Clock*, Mr. White's belief that it is 8:22 is intuitively justified. Can KFVE accommodate this intuition? To answer this question, notice first that Mr. White acquires his belief that it is 8:22 in a way that qualifies as an ability to know propositions about the time. After all, forming his belief via this way disposes Mr. White to acquire knowledgeable beliefs about the time across a range of conditions. For instance, it does so in cases just like *Stopped Clock* with the exception that the clock he is taking a reading from is functioning properly. Given that Mr. White acquires his belief that it is 8:22 in a way that qualifies as an ability to know propositions about the time, his belief is the product of an exercise of an ability to know propositions about the time. Since the belief Mr. White forms is a belief about the time, the target proposition is within the range of the ability to know that produced his belief. By *Competent Belief*, Mr. White's belief that it is 8:22 is competent and, by KFVE-JB, justified.²⁰

Second, Mr. White's belief intuitively falls short of knowledge. Recall that KFVE does provide an informative virtue condition on knowledge. As a result, we may ask whether this condition is satisfied here. If it isn't, then KFVE has the additional benefit of being able to explain the intuition of absence of knowledge, besides the intuition of presence of justified belief in this kind of case.

Let's return to *Stopped Clock* once more, then. Notice that the clock Mr. White is taking a reading from is stopped. As a result, he is not in conditions *C* relative to which he has the ability to acquire knowledge about the time. After all, Mr. White's way of belief formation does not dispose him to form knowledgeable beliefs about the time in conditions in which the clock he is taking a reading from is stopped. At the same time, the conditions Mr. White finds himself in do not prevent him from forming his belief via the way that constitutes an ability to know relative to some (but different) conditions. This means that Mr. White is in unsuitable *SI* for his ability to acquire knowledge about the time. By *Apt Belief*, his belief that it is 8:22 is not apt and, by KFVE-K, it falls short of knowledge. In this way, KFVE predicts not only presence of justified belief, but also the absence of knowledge in *Stopped Clock*. There is thus reason to believe that KFVE can successfully handle standard Gettier cases.

²⁰ It may be worth noting that, as a result, KFVE can also allow for justified false beliefs. If this isn't immediately obvious, consider a variation of *Stopped Clock* in which Mr. White acquires his belief about the time a minute earlier or later. In that case, his belief that it is 8:22 will be false. At the same time, according to KFVE-JB it will still be justified. In fact, as I argue elsewhere [Kelp 2015b] in more detail the virtue epistemology at issue in KFVE can solve or defuse a number of long-standing problems for broadly reliabilist accounts of justified belief, including the new evil demon problem, the problem of clairvoyance cases and the generality problem.

5.2 Fake Barn Cases

Arguably, standard Gettier cases have never posed much of a difficulty for VE. The hard nut to crack are fake barn cases. Let's return to *Fake Barns* then. Here, too, Mr. Pink's belief that he is facing a barn is intuitively justified but falls short of knowledge. Fortunately, it is easy enough to see that KFVE can explain the presence of justified belief and the absence of knowledge here in much the same way as in standard Gettier cases.

First, concerning the presence of justified belief, the way in which Mr. Pink forms his belief qualifies as an ability to recognise barns, for the same reasons as Mr. White's way of belief formation in *Stopped Clock* qualifies as an ability to know propositions about the time. If so, Mr. Pink's belief that he is facing a barn is the product of an exercise of an ability to recognise barns. Since the belief Mr. Pink forms is a belief about the presence of a barn, the target proposition is within the range of the ability to know that produced his belief. By *Competent Belief*, Mr. Pink's belief that he is facing a barn is competent and, by KFVE-JB, justified.

Second, concerning the absence of knowledge, the fact that Mr. Pink is in a part of the country in which fake barns predominate means that he is not in conditions *C* relative to which he has the ability to recognise barns. At the same time, the conditions Mr. Pink finds himself in do not prevent him from forming his belief via the way that constitutes an ability to know relative to some (but different) conditions. This means that Mr. Pink is in unsuitable *SI* for his ability to recognise barns. By *Apt Belief*, Mr. Pink's belief that he is facing a barn is not apt. By KFVE-K, it falls short of knowledge.

Of course, the real problem that fake barn cases pose for VE is not just to accommodate the intuition of absence of knowledge in fake barn cases. Traditional versions of VE can achieve this much. Rather, the challenge is to offer a satisfactory account of fake barn cases, without at the same time ending up making incorrect predictions in Frankfurt cases. Let's see whether KFVE can rise to this challenge as well.

Recall our toy Frankfurt case:

Frankfurt Clock. Mr. Blue's arch-nemesis, a powerful demon, has an interest that Mr. Blue forms a belief that it's 8:22 by looking at the grandfather clock in the hallway when he comes down the stairs. In order to achieve this, Mr. Blue's arch-nemesis is prepared to set the clock to 8:22 when Mr. Blue comes down the stairs. However, Mr. Blue's arch-nemesis is also lazy. He will act only if Mr. Blue does not come down the stairs at 8:22 of his own accord. Suppose, as it so happens, Mr. Blue does come down the stairs at 8:22. Mr. Blue's arch-nemesis remains inactive. Mr. Blue forms a belief that it's 8:22. It is 8:22. The grandfather clock is working reliably as always.

Intuitively, Mr. Blue knows that it is 8:22. What does KFVE predict about this case? And does KFVE's treatment of *Fake Barns* entail that Mr. Blue does not know that it's 8:22? Let's start with the second question. Fortunately, the answer here is no. In a nutshell, the reason for this is that KFVE does not offer a reductive

analysis of knowledge and hence is free to treat the difference between *Fake Barns* and *Frankfurt Clock* as a fundamentally epistemic difference, that is, roughly, a difference in knowledge. More specifically, while both Mr. Pink and Mr. Blue form competent beliefs, only Mr. Blue is in suitable *SI* for the ability to know he exercises and so satisfies KFVE's necessary condition on knowledge. Why is that? The answer is that only Mr. Blue is in *SI* such that his way of belief formation disposes him to form the relevant knowledgeable beliefs. In this way, KFVE can make a difference between Mr. Pink and Mr. Blue, albeit one that is fundamentally epistemic. As a result, KFVE's treatment of *Fake Barns* does not lead to a corresponding incorrect prediction in *Frankfurt Clock*.²¹

This leaves the question as to what KFVE predicts about *Frankfurt Clock*: presence or absence of knowledge? The answer to this question is neither. After all, KFVE does not offer a sufficient condition for knowledge. What's more, it is in the spirit of knowledge first epistemology to consider knowledge as basic in the sense that we need not seek to offer an account of knowledge that explains the presence of knowledge in particular cases. But then we need not expect KFVE to make predictions about the presence of knowledge in *Frankfurt Clock* to begin with. Be that as it may, the important point for present purposes is that KFVE's treatment of fake barn cases does not lead the view into trouble with Frankfurt cases.

Conclusion

In this paper, I have developed a novel virtue epistemological account of knowledge and justified belief, one that gives VE a knowledge-first twist. In order to achieve this, I first outlined a framework for simple goal-directed practices, i.e. practices that can be characterised in terms of targets, moves and a designated relation, and defined the notions of ability, exercise of ability, competent and apt moves for such practices. I then suggested that inquiry into whether questions can be viewed as a collection of simple goal-direct practices in which targets are true answers, moves are beliefs and the designated relation is the knowledge relation. I derived the corresponding accounts of abilities to know, their exercises, competent and apt belief. With these accounts in play, I went on to embrace the familiar virtue epistemological accounts of knowledge as apt belief and justified belief as competent belief. The crucial difference between my account and its traditionalist

²¹ Note that while KFVE can allow that Mr. Blue knows while Mr. Pink doesn't, it is not committed to these verdicts. For instance, KFVE can also accommodate the intuitions of those who think that Mr. Pink knows, *viz.* by allowing that the *SI* in *Fake Barns* are suitable (as Mr. Pink's way of belief formation does dispose him to form the relevant knowledgeable beliefs in those *SI*). Similarly, KFVE can accommodate the intuitions of who think that Mr. Blue doesn't know, *viz.* by allowing that the *SI* in *Frankfurt Clock* are unsuitable (as Mr. Blue's way of belief formation does not dispose him to form the relevant knowledgeable beliefs in those *SI*). Of course, if either one of these intuitions turns out correct, my argument from Frankfurt cases will not serve to establish an advantage for KFVE over traditional VE after all. For what it's worth, however, my own intuitions are that Mr. Blue knows, whereas Mr. Pink doesn't, which is why I think the argument does go through.

cousins is that abilities are unpacked as abilities to know. In this way, the account qualifies as a knowledge first version of virtue epistemology. Finally, I showed that KFVE has an edge over its traditionalist cousins. Not only can it handle fake barn cases without further difficulty, it can also offer a uniform account of all Gettier cases, including the fake barn and the standard variety. In view of all this, I submit, knowledge first virtue epistemology is a promising view that deserves to be taken seriously.²²

References

- Baumann, P. 2008. Is knowledge safe? *American Philosophical Quarterly*, 45:19–30.
- Bird, A. 2007. Justified judging. *Philosophy and Phenomenological Research*, 74:81–110.
- Bogardus, T. 2013. Knowledge under threat. *Philosophy and Phenomenological Research*, 86:289–313.
- Broncano-Berrocal, F. 2015. In defence of performance epistemology. *Unpublished Manuscript*.
- Chrisman, M. 2012. The normative evaluation of belief and the aspectual classification of belief and knowledge attributions. *Journal of Philosophy*, 109:588–612.
- Comesaña, J. 2005. Unsafe knowledge. *Synthese*, 146:395–404.
- Frankfurt, H. 1969. Alternate possibilities and moral responsibility. *The Journal of Philosophy*, 66:829–39.
- Gettier, E. 1963. Is justified true belief knowledge? *Analysis*, 23:121–23.

²² *Acknowledgements.* Thanks to the audiences of the following conferences for helpful feedback on 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). Thanks also to the audiences of the following conferences for helpful feedback on other papers on knowledge first virtue epistemology, from which this paper has benefitted: *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 Conor McHugh, Daniel Whiting, Chris Tucker, Adam Carter, Emma Gordon, Benjamin Jarvis, Jan Heylen, Fernando Broncano-Berrocal, Ernest Sosa, anonymous referees of various journals and OUP as well as the Leuven Epistemology Group for detailed 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).

- Ginet, C. 1988. The fourth condition. In Austin, D., editor, *Philosophical Analysis. A Defense by Example*. Springer, Dordrecht.
- Goldman, A. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy*, 73:771–91.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press, Cambridge.
- 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, C. 2014b. Knowledge, understanding and virtue. In Fairweather, A., editor, *Virtue Scientia. Virtue Epistemology and Philosophy of Science*. Springer, Dordrecht.
- Kelp, C. 2014c. Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51:227–32.
- Kelp, C. 2015a. Epistemic Frankfurt cases revisited (forthcoming). *American Philosophical Quarterly*.
- Kelp, C. 2015b. How to be a reliabilist. *Unpublished Manuscript*.
- Kelp, C. 2015c. Justified belief: knowledge first-style. *Unpublished Manuscript*.
- Kelp, C. 2015d. Understanding phenomena. *Synthese*.
- Kelp, C. and Ghijzen, H. 2015. Perceptual justification: factive reasons and fallible virtues (forthcoming). In Mi, C., Slote, M., and Sosa, E., editors, *Moral and Intellectual Virtues in Western and Chinese Philosophy*. Routledge, London.
- Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press, Cambridge.
- Littlejohn, C. 2013. The Russellian retreat. *Proceedings of the Aristotelian Society*, 113:293–320.
- McKittrick, J. 2003. A case for extrinsic dispositions. *Australasian Journal of Philosophy*, 81:155–174.
- 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.

- Miracchi, L. 2014. Competence to know (forthcoming). *Philosophical Studies*.
- Mumford, S. 1998. *Dispositions*. Oxford University Press, Oxford.
- Neta, R. and Rohrbaugh, G. 2004. Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, 85:396–406.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press, Oxford.
- Pritchard, D. 2007. Recent work on epistemic value. *American Philosophical Quarterly*, 44:85–110.
- Pritchard, D. 2008. Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, 58:437–47.
- Pritchard, D. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy*, 109:247–79.
- Pritchard, D., Millar, A., and Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford University Press, Oxford.
- Riggs, W. 2002. Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64:79–96.
- 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. 2007. *Without Justification*. MIT Press, Cambridge/MA.
- Turri, J. 2011. Manifest failure: the Gettier problem solved. *Philosopher's Imprint*, 11:1–11.
- 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.
- Zagzebski, L. 1994. The inescapability of Gettier problems. *The Philosophical Quarterly*, 44:65–73.
- Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press, Cambridge.