



Kelp, C. (2019) How to be a reliabilist. *Philosophy and Phenomenological Research*, 98(2), pp. 346-374.

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.

This is the peer reviewed version of the following article:

Kelp, C. (2019) How to be a reliabilist. *Philosophy and Phenomenological Research*, 98(2), pp. 346-374.
(doi:[10.1111/phpr.12438](https://doi.org/10.1111/phpr.12438))

This article may be used for non-commercial purposes in accordance with [Wiley Terms and Conditions for Self-Archiving](#).

<http://eprints.gla.ac.uk/140979/>

Deposited on: 16 May 2017

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

How to Be a Reliabilist

Christoph Kelp*

Abstract

In this paper, I aim to develop a novel virtue reliabilist account of justified belief, which incorporates insights from both process reliabilism and extant versions of virtue reliabilism. Like extant virtue reliabilist accounts of justified belief, the proposed view takes it that justified belief is a kind of competent performance and that competent performances require reliable agent abilities. However, unlike extant versions of virtue reliabilism, the view takes abilities to essentially involve reliable processes. In this way, the proposed view takes a leaf from process reliabilism. Finally, I will provide reason to believe that the view compares favourably with both extant versions of virtue reliabilism and process reliabilism. In particular, I will show that in taking abilities to essentially involve reliable processes, the view has an edge over extant versions of virtue reliabilism. Moreover, I will argue that the proposed view can either solve or defuse a number of classical problems of process reliabilism, including the new evil demon problem, the problem of clairvoyant cases and the generality problem.

Introduction

One of the most prominent accounts of justified belief in recent literature is reliabilism. Among reliabilist accounts, at least two species can be distinguished. The most widely discussed reliabilist theory is process reliabilism, which was first stated in Alvin Goldman's seminal 1979 paper. According to process reliabilism what matters to whether one believes justifiably is, roughly, whether one's belief was produced by a process that is reliable in the sense that it tends to produce beliefs with a favourable truth to falsity ratio.¹ An alternative to process reliabilism is virtue reliabilism, which first surfaced in the contemporary debate in [Sosa 1980]. It agrees with process reliabilism that there is a reliability condition on justified belief. However, what matters for justified belief is not so much whether a belief was produced by a reliable process, but, again roughly, whether it was produced

*KU Leuven, Email: christoph.kelp@hiw.kuleuven.be.

¹ Process reliabilism has been defended extensively by Alvin Goldman [e.g. 1979, 1986, 2012]. Other champions include Hilary Kornblith [e.g. 2002, 2004, 2009], Jack Lyons [e.g. 2009] and Sandy Goldberg [e.g. 2010].

by a reliable agent ability, where the notion of ability is analysed in terms of an agent disposition.²

For the purposes of this paper I will assume without further argument that some reliabilist account of justified belief is correct. My central aim will be to develop a novel reliabilist account of justified belief, which improves on extant versions of both process and virtue reliabilism by incorporating insights from both sides. It sides with extant versions of virtue reliabilism in that it takes justified belief to require reliable agent abilities. However, unlike extant versions of virtue reliabilism, it does not analyse agent abilities in terms of agent dispositions. Rather, it takes abilities to essentially involve processes. In this way, the account also incorporates a process reliabilist element.³

In order to achieve this, I will first outline process reliabilism as well as a number of prominent problems the view encounters (section 1). In section 2, I will present virtue reliabilism and highlight some advantages the view has vis-à-vis process reliabilism. However, I will also develop a problem for virtue reliabilism and argue that this problem can be solved by introducing processes into the account of abilities. Section 3 develops a novel general account of (a certain kind of) competent performance. More specifically, I analyse competent performances in terms of abilities and their exercises, where abilities, in turn, are taken to essentially involve reliable processes of a sort. In section 4, I apply the general account of competent performance to the case of belief and derive a novel reliabilist account of justified belief. Since competent performances in general involve reliable processes, so do justified beliefs according to the proposed account. Given that the resulting reliabilist account employs process reliabilist ideas, I will ask whether the view falls prey to the classical problems for process reliabilism. I will argue the answer to this question is no. Finally, section 5 considers some objections to this view and offers responses.

1 Process Reliabilism

1.1 The View

Process reliabilism offers an account of justified belief. The core idea of process reliabilism is that justified beliefs are beliefs that are produced by processes that tend⁴ to produce beliefs with a favourable truth to falsity ratio. In other words,

² Sosa has published extensively on the virtue reliabilism [see e.g. Sosa 1991, 2011, 2015 as well as his contribution to *BonJour & Sosa* 2003] and continues to be its most prominent advocate. Other prominent champions of virtue reliabilism include John Greco [e.g. 1999, 2000, 2010] and Linda Zagzebski [e.g. 1996]. I have also defended a version of the view e.g. in [Kelp 2016, Forthcoming].

³ Note that I do not aim to offer an account that effects a happy marriage between virtue and process reliabilism where both are taken to be equal partners. Rather, I take the virtue reliabilist component to do the bulk of the work. Accordingly, the view may be thought of as a version of virtue reliabilism with a process reliabilist twist.

⁴ Goldman deliberately leaves open whether to give ‘tendency’ an frequentist or a propensity interpretation. His reason for this is that our ordinary conception of justification is vague on this

justified beliefs are beliefs produced by processes that reliably produce true beliefs (henceforth also ‘reliable processes’).

With the core idea in play, let’s turn to some important qualifications. Note that not all belief-forming processes are created equal. Some belief-forming processes, most notably perception, do not have beliefs among their inputs. For instance, my perceptual belief that I am sitting at my desk has no beliefs among its inputs. It is not as if I believe that I am sitting at my desk based on a prior belief that I have an experience as of sitting at my desk. Rather, the only inputs to the process are non-doxastic. They may comprise an experience as of a desk, certain retinal stimuli or perhaps something entirely different. Other belief-forming processes do have beliefs among their inputs. The most prominent example here is inference. When you form a belief that q by inference from a belief that p and a belief that if p , then q , the inferential process that outputs your belief that q has your belief that p and your belief that if p then q among its inputs. Following Goldman, I will call the former kinds of process ‘belief-independent’ and the latter ‘belief-dependent’.

While it makes sense to require belief-independent processes to be reliable in the sense that in order to deliver justification they must produce beliefs with a favourable truth to falsity ratio, the same does not hold for belief-dependent processes. To see this notice that belief-dependent processes need not be expected to produce beliefs with a favourable truth to falsity ratio when the input beliefs happen to be false. Moreover, they need not be expected to ensure that the input beliefs be true. We will do well, then, to weaken the reliability condition on belief-dependent processes. Rather than requiring that they produce beliefs with a favourable truth to falsity ratio unconditionally, these processes need only be conditionally reliable in the sense that they must produce belief with a favourable truth to falsity ratio given that the input beliefs are true.⁵

Finally since conditionally reliable processes can only transmit justification but

front and that, as a result, its appropriate to leave the theory vague also [Goldman 1979: 11]. For the purposes of this paper, I will follow Goldman on this front. For the record, I think there is reason to prefer something along the lines of a propensity interpretation of ‘tendency’ over the frequentist alternative.

⁵ It may be worth noting that the usefulness of the distinction between belief-dependent and belief-independent processes does not go undisputed. For instance, [Lyons 2009] argues that the key distinction for process reliabilists is the distinction between basic and non-basic beliefs. While it might be thought that the distinction between belief-dependent and belief-independent processes unpacks just this distinction in process reliabilist terms, there is excellent reason to think that this cannot be the case after all. For instance, the resulting account will categorise introspective beliefs about other beliefs as non-basic, whereas they should come out as basic. Another way in which the distinction may be problematic is if the processes that produce perceptual beliefs turn out to have perceptual experiences with propositional contents as input. In this case, an argument parallel to the above argument should provide reason to think that the reliability condition for perceptual processes should be weakened also. Since perceptual experiences are not beliefs, Goldman’s distinction is problematic on yet another count. (Thanks to an anonymous referee for pointing this out.) While these are important problems for process reliabilism, I will not take issue with them within the confines of this paper. Rather, I will rest content with simply following Goldman’s exposition here.

not generate it, belief-dependent processes will produce justified beliefs only if the input beliefs are themselves justified. This gives us the process reliabilist account of prima facie justified belief in Goldman's classic formulation:

Process Reliabilist Justification (PRJ). *S*'s belief that *p* at *t* is prima facie justified if and only if, at *t*, (i) it "results ('immediately') from a belief-independent process that is (unconditionally) reliable" or (ii) it "results ('immediately') from a belief-dependent process that is (at least) conditionally reliable, and . . . the beliefs (if any) on which this process operates in producing *S*'s belief that *p* at *t* are themselves justified." [Goldman 1979: 13-14]

Even this cannot be the whole story. After all, compatibly with a given belief's being formed reliably, the justification of the belief can be defeated. Goldman proposes a process reliabilist-friendly account of defeat along the following lines:

Process Reliabilist Defeat (PRD). *S*'s belief in *p* at *t* is defeated if and only if "there is an alternative reliable or conditionally reliable process available to *S* which, had it been used by *S* in addition to the process actually used, would have resulted in *S*'s not believing *p* at *t*."⁶ [Goldman 1979: 20]

Whether a belief is ultima facie justified depends on whether it is prima facie justified and not defeated in the sense specified in *PRD*.

Goldman thus offers an account of justified belief covering basic and non-basic justification as well as prima facie and ultima facie justification. I'd like to emphasise that, for the purposes of this paper, I will restrict my discussion to prima facie and basic justification. That is to say, I will bracket the phenomena of beliefs formed by belief-dependent processes and defeat. In addition, I would like to focus on justification of first-order beliefs only. Accordingly, I will henceforth take 'justified belief' to mean 'prima facie basic justified first-order belief', and I will work with the following process reliabilist account of justification:

Process Reliabilism (PR). *S*'s belief that *p* at *t* is justified if and only if, at *t*, it results ('immediately') from a belief-independent process that is (unconditionally) reliable.

With PR thus in play, let's look at some of the problems the view is said to encounter.

⁶ There are various problems with PRD. For instance, process reliabilists will have provide an account of what it takes for an alternative process to be available in the relevant sense. And not every interpretation of 'available' will work for process reliabilists. (Thanks to an anonymous referee for pointing this out. For one interpretation of 'available' that will spell trouble for process reliabilism see [Kvanvig 2007]; see [Beddor 2015] for further critical discussion of PRD.) While these are also significant issues for process reliabilism, I will again not stop to discuss them here. Rather, I will again rest content with simply following Goldman's exposition of these issues.

1.2 Problems

PROBLEM 1: THE NEW EVIL DEMON PROBLEM

Consider cases of radical deception. Let's suppose, for instance, that you have become the victim of a mad scientist who has turned you into a brain in a vat (henceforth just 'BIV') that is being fed deceptive experiences as if everything were normal. It would seem that, in this case, the beliefs you form based on your experiences continue to be justified. For instance, when you appear to be waking up in your bed in the morning, and appear to see that your alarm clock reads 7:45, your corresponding beliefs that you are lying in your bed and that it is 7:45 are justified. Or, at any rate, they are no less justified than they were before you were turned into a BIV. The problem for PR is that, since you are subject to radical deception, the processes that produce your beliefs are highly unreliable. You are a disembodied BIV. You are not lying and surely not in your bed. Your beliefs about the time (let's suppose) are wildly mistaken also. As a result, the processes that produce the corresponding beliefs tend to produce beliefs with a highly unfavourable truth to falsity ratio. PR predicts, mistakenly, that the beliefs you form are not justified here.⁷

PROBLEM 2: CLAIRVOYANCE CASES

PR meets with a similar fate when it comes to cases like the following. Suppose, as a result of exposure to radiation, you start forming beliefs via a 'clairvoyance' belief forming process that reliably produces true beliefs about distant events. You do not know that you form beliefs in this way. In fact, you have no evidence that there exists a process of this kind or that it should be so much as possible for it to exist. On the other hand, you also do not have evidence that such a process does not exist/is not possible. From your point of view, you spontaneously form beliefs about distant events. Just now, whilst being on vacation in a faraway country, the clairvoyance process has produced a belief that your house is on fire. This belief is not justified. At the same time, the process that produced your belief is highly

⁷ This problem, which is also known as the new evil demon problem, was first stated in [Lehrer & Cohen 1983] and [Cohen 1984] as a problem for reliabilist accounts of justified belief. [Wedgwood 2002] argues that it generalises to all externalist accounts of justified belief.

But couldn't PR solve the new evil demon problem by taking the belief-forming processes to start at the periphery of the brain? After all, even though you will be forming mostly false beliefs when you become the unlucky victim of a mad scientist, the vast majority of people won't. Since processes that start at the periphery of the brain will be shared between victims and non-victims, the processes that produce beliefs even in victims continue to be reliable. Note, however, that it is a contingent matter of fact that the vast majority of people who use the same processes form true beliefs reliably. It might be otherwise. Nearly everyone may be in your predicament. As a result, it's at least possible that your beliefs are formed via unreliable processes. Since PR is a necessarily true if true at all, this will be enough for those who want to charge PR with new evil demon problem. Thanks to an anonymous referee for pressing me on this point.

reliable and so PR predicts, again mistakenly, that your belief is justified.^{8,9}

PROBLEM 3: THE GENERALITY PROBLEM

The core idea of PR is that a belief is justified if and only if it is produced by a reliable process. To be more precise, the core idea here is that a belief is justified if and only if it is produced by a token of a reliable process *type*. Individual beliefs are produced by token processes. Token processes are particulars, which, among other things, are not repeatable. As a result, it's not clear that the notion of reliability even applies to token processes which is why champions of PR typically take justified belief to require production via a reliable process type.¹⁰

To see the difficulty for PR, notice that every token process instantiates indefinitely many types. The process that produced my belief that I am sitting at my desk, for instance, instantiates all of the following types: cognitive process, belief-forming process, perceptual belief-forming process, visual belief-forming process, to name but a few. Each process type has a different degree of reliability. Of course, unless it is specified which process types exactly are the ones at issue in PR, PR simply does not make determinate predictions about whether or not individual beliefs are justified. For instance, before PR can make a determinate prediction on whether my belief that I am sitting at my desk is justified, we need a specification of which of the various process types that the token process instantiates is the one at issue in PR. In order to work as an account of justified belief at all, PR must be supplemented with an account that specifies the relevant type of

⁸ Clairvoyant cases were first adduced by Bonjour [1980, 1985] to argue against reliabilist accounts of justified belief. The related case of Mr. Truetemp is due to Lehrer [1990].

⁹ It may be worth noting that the new evil demon problem and the problem posed by clairvoyance cases are especially worrisome for PR in tandem. The new evil demon problem provides reason to believe that process reliability is not necessary for justified belief, while clairvoyant cases suggest that it is not sufficient either. If process reliability is neither necessary nor sufficient for justified belief, it looks as though PR is heading down the wrong epistemological track entirely.

¹⁰ [Conee & Feldman 1998: 2]. It may be worth noting that if 'tendency' is given a propensity (rather than frequentist) interpretation, then it will be possible for token processes to be reliable and unreliable. Note, however, that 'token versions' of process reliabilism will face analogues of all the classical problems for 'type versions'. To see this, note first that tokens of the wildest belief forming processes can be reliable, e.g. when one has a helper in the wings who has the power to make the belief true and is determined (in a modally robust manner) to do so. Suppose, for instance, that I believe that I will be the world's richest man based on wishful thinking. My belief is not justified. However, with a powerful and committed helper in the wings who is ready to (modally robustly) see to it that I become the world's richest man, the token process that produced this belief may well be reliable. On the other hand, even tokens of normally excellent belief forming processes can be unreliable, e.g. when one is up against an opponent who has the power to make the belief come out false and is determined (in a modally robust manner) to do so. Suppose, for instance, that I believe that I am looking at an apple based on normally highly reliable visual perceptual processes. My belief is justified. However, when up against a powerful and committed opponent, who on this particular occasion (modally robustly) made me look at an indistinguishable fake, the token process that produced this belief is unreliable. Finally, as [Comesaña 2006] convincingly argues, token versions of process reliabilism also encounter a version of the generality problem. So, not much will be gained by abandoning the classical type version of process reliabilism in any case.

processes at issue in PR. The problem of specifying the relevant type of process has become known as the generality problem for PR.¹¹

PROBLEM 4: WORLD-BOUND RELIABILITY

The last problem that I will discuss here is somewhat less prominent. As will become clear in due course, the main reason I introduce it is for dialectical purposes. To see how it arises notice that, according to PR, process reliability is evaluated relative to worlds. That is to say, whether a given process type is reliable depends on the ratio of true to false beliefs it attains at a set of worlds.¹²

But now consider a process type, *T*, that takes a certain look of a certain mushroom as input and outputs a classification of the mushroom as edible. Suppose *T* is a highly unreliable process type because, at the set of worlds at which it is evaluated for reliability, both edible chanterelles and poisonous jack-o'-lantern mushrooms have this look and throughout this set of worlds nearly everywhere where there are chanterelles, there are also jack-o'-lanterns. But now suppose, compatibly with this, that throughout this set of worlds there also exists a remote island with only a handful inhabitants where there are no jack-o'-lanterns. None of the inhabitants has ever left the island and they never will. Suppose you are an inhabitant of this island and form a belief of the mushroom before you that it is edible via a token of *T*. Your belief is intuitively justified. Since *T* is highly unreliable due to its tendency to produce too many false beliefs at the set of worlds at which it is evaluated for reliability, however, PR predicts, incorrectly, that your belief is unjustified.

Consider also the converse situation in which *T* is highly reliable partly because, throughout the set of worlds at which it is evaluated for reliability, nearly only chanterelles have the look in question. Jack-o'-lanterns are nowhere to be found, with the exception of the remote island you inhabit where they grow as abundantly as chanterelles. When you form the belief of the mushroom before you that it is edible via *T*, your belief is intuitively not justified. Since *T* is highly reliable as it tends to produce true beliefs at the set of worlds at which it is evaluated for reliability, PR makes yet another incorrect prediction, this time of the presence of justified belief.¹³

¹¹ The generality problem was already noted in Goldman's original 1979 paper. It was developed into a serious challenge for process reliabilism by Conee and Feldman [1998].

¹² This is clear especially from various responses to the new evil demon problem. Goldman himself considers 'actual world' and 'normal worlds' versions of PR according to which whether a belief is justified according to PR depends on whether the relevant process type is reliable at the actual world or at normal worlds, i.e. "worlds that are consistent with our general beliefs about the actual world" [Goldman 1986: 107]. The idea is that the new evil demon problem can be avoided since the relevant process type is reliable at, respectively, the actual world and normal worlds, at both of which the relevant processes are reliable. Both of these solutions presuppose that the domain of evaluation are worlds of some denomination.

¹³ Isn't this problem structurally analogous to the fake barns problem? And in view of recent x-phi results, doesn't this mean that proper epistemological theorising should not rather explain why intuitions clash here? Thanks to an anonymous referee for raising this issue.

2 Virtue Reliabilism

2.1 Performance Normativity

Recent statements of virtue reliabilism (VR) start from a general account of performance normativity. The core idea is that performances with a goal can be assessed along the following three dimensions:

Success. Does the performance reach its goal? In other words, is it successful?

Competence. Does the agent perform competently? In other words, is the performance produced by an ability to attain the performance's success?

Aptness. Is the performance successful because competent? In other words, does the right kind of (explanatory) relation obtain between competence and success.

In addition, VR assumes that beliefs are a type of performance with a goal. More specifically, they are a type of epistemic performance. This means that the general account of performance normativity can be applied to the special case of belief, delivering an account of the normativity of belief.

A key question then concerns the nature of the goal of belief understood as performance. According to standard versions of VR, the goal of belief is truth.¹⁴ Given that the goal of belief is truth, we get the following:

First, I am just not sure just what lesson we epistemologists should take the x-phi results to teach us. The reason for this is that it is not clear to me why we should prioritise laypeople's judgements on these cases over expert judgements. On the contrary, I am attracted to side with those who think that expert judgements should take precedence over laypeople's [e.g. Williamson 2007, 2011].

Second, while I'd agree that an account of *fake barn cases* (or a full epistemological theory for that matter) will do well to explain the clash of intuitions in fake barn cases, it's not clear to me that the same goes for an account of *knowledge*. Rather, it seems to me that what an account of knowledge needs to do is accommodate the epistemic facts, whatever they may be. For instance, if agents in fake barn cases lack knowledge, that's what the account of knowledge should predict. And the same goes, *mutatis mutandis*, for accounts of justified belief. Since I am interested in giving an account of justified belief, it's not clear that proper theorising should really explain the clash of intuitions here.

Third, even if (i) the cases are structurally analogous, (ii) the x-phi results do suggest a clash of intuitions about fake barn cases and (iii) as a result we should explain the clash of intuitions about fake barn cases, it's not clear this means that the same goes, *mutatis mutandis*, for this case. After all, this case is a case about justified belief rather than knowledge. Since an unclear case of knowledge can still be a clear case of justified belief, we can't derive conclusions about cases of justified belief from otherwise structurally analogous cases of knowledge.

But once we consider the question whether the belief in the fake barn case is justified, it seems clear that the answer is that it is. If so, doesn't the analogy provide us with reason to think that the same goes for the above case? No (and this is my last point). It's not clear that the cases are analogous. In the fake barn case, the environment in which the agent normally finds himself—his natural habitat, as it were—does not include any fakes, whereas in the above case it does. That's a potentially important structural difference between the two cases.

¹⁴ Note that alternatives to standard VR are conceivable and have, as a matter of fact, been defended in the literature. [Sosa 2015] argues that judgemental belief is of central epistemological interest and that judgemental belief essentially aims at a form of aptness. If this turns out to

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

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.

Crucially, VR identifies knowledge and justified belief with different normative properties of beliefs as performances. In particular, according to VR:

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

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

Note that the resulting virtue reliabilism has a number of highly attractive features. First, it offers an attractive account of the normativity of belief. Second, it can offer a theoretical motivation for its accounts of justified belief and knowledge. After all, the accounts are backed by a general framework for the normativity of performances. Knowledge and justified belief are instances of familiar normative categories, to wit, that of apt and competent performance. Third, a number of plausible theses about knowledge and justified belief and their relations drop right out of the account. Since aptness entails both success and competence, knowledge entails both true belief and justified belief. In contrast, since success and competence do not entail aptness, justified true belief does not entail knowledge. Moreover, since success does not entail competence and vice versa, justified belief and true belief turn out to be logically independent. We can have true but unjustified beliefs and justified false beliefs.

With these remarks about VR in general in place, I'd like to stress that, in what follows, I will focus mainly on VR's account of justified belief as competent belief. In particular, I will now sketch the virtue reliabilist account of competent performance from which the key notion of competent belief is derived.

2.2 Competent Performances

THE GENERAL CASE

According to the general account of performance normativity champions of VR have appealed to, a performance with the aim of attaining a certain success is competent if and only if it is produced by an ability to attain the relevant success. This immediately raises the questions as to what abilities are.

be incompatible with the idea that truth is the goal of belief, this means that Sosa departs from standard VR in this respect. While Sosa may not be a clear case of departure from standard VR, my own view is. I favour a knowledge first version of virtue reliabilism according to which the goal is knowledge, which I defend in more detail e.g. in [Kelp 2016, Forthcoming]. For the purposes of this paper, nothing hinges on whether we opt for a traditional or a knowledge first version of virtue reliabilism. For that reason, I will follow orthodoxy and work with a traditionalist version of the view.

It is widely agreed among champions of VR that abilities are *agent dispositions*.¹⁵ More specifically, they are agent dispositions to produce successful performances. For instance, in order to have the ability to hit the target in target archery, you must have the disposition to hit the target.

Since abilities are dispositions, a closer look at dispositions may help illuminate the nature of abilities. First, it is plausible that dispositions are relative to conditions. Take the disposition of water to boil when heated to 100°C. While water has this disposition in certain conditions, i.e. at sea level, it does not have this disposition in other conditions. For instance, it doesn't have it at altitudes below sea level.

Second, dispositions have trigger conditions (*T*) and manifestation conditions (*M*). In the case of water's disposition to boil when heated to 100°C, the trigger condition is heating to 100°C and the manifestation condition is boiling.

Third, dispositions correspond to trigger-manifestation conditionals. To say that something, *x*, has the disposition to *M* when *T* is to say that were *T* to obtain in *C*, *x* would (likely enough) exhibit *M*. For instance, to say that water has the disposition to boil when heated to 100°C at sea level is to say that were water heated to 100°C at sea level, it would (likely enough) boil.

If abilities are dispositions and dispositions have these properties, we may expect that abilities have these properties as well. And this is exactly what we find.

First, abilities are relative to conditions. Your ability to hit the target in target archery is relative to conditions. You do not have it when drunk beyond comprehension, when strong winds are blowing, etc. Why not? Because you don't have the disposition to hit the target in those conditions.

Second, abilities have trigger and manifestation conditions. According to Sosa, for instance, the trigger conditions for abilities are tryings and the manifestation conditions successes [Sosa 2015: 96]. In the case of your ability to hit the target, the trigger conditions is trying to hit the target and the manifestation condition is hitting the target.

Third, abilities correspond to trigger-manifestation conditionals. Your ability to hit the target in conditions *C* (sufficiently sober, not too strong winds etc.) corresponds to the conditional: if you were to try to hit the target in *C*, you would (likely enough) succeed.¹⁶

It may be worth noting that the trigger-manifestation conditionals effectively impose a reliability condition on abilities. Abilities are, by nature, reliably connected to the relevant successes in the sense that is specified by the trigger-manifestation conditional.

¹⁵ See e.g. [Sosa 2010, Greco 2010, Kelp 2011, Pritchard 2012, ?].

¹⁶ My presentation of VR's account of abilities follows [Sosa 2015]. Other champions of VR have offered slightly different accounts. Greco, for instance, holds that to have an ability to attain *S* in certain conditions, *C*, one must be such that one attains *S* with a high rate success across nearby possible worlds at which *C* obtain [Greco 2010: 77]. However, these differences are of little consequence for the purposes of this paper.

THE EPISTEMIC CASE

Let's apply this account to the case of belief, still understood as a type of performance that aims at truth. The result that we get is that a belief is justified if and only if it is produced by an exercise of an ability to believe truths, that is, a disposition on the part of the agent to form true beliefs.

Of course, the dispositions at issue in these abilities also share the properties of dispositions. That is to say, they are relative to conditions and have the relevant trigger and manifestation conditions. They also correspond to trigger manifestation-conditionals and, as a result, feature a reliability condition.

What comes to light is that, according to VR, a reliability condition on justified belief drops right out of VR's account of justified belief. As a result, VR can indeed be seen as a reliabilist account of justified belief. What's more, note that VR's reliability condition on justified belief is but an instance of a general and independently plausible reliability condition on competent performance. As a result, VR not only incorporates a core reliabilist idea, it also rationalises it.

THE PROBLEM OF WORLD-BOUND RELIABILITY REVISITED

Note that, with this account of abilities in hand, champions of VR can already make progress on one of the problems that beset PR, to wit, the problem of world-bound reliability. To see how VR can solve this problem note that that what matters to justification according to VR is not the reliability of a process type at a world, but the reliability of an agent in a particular set of conditions. In the case in which you are one of a few inhabitants of a remote island on which there are no jack-o'-lanterns, your belief that the mushroom you are holding is edible will be justified according to VR. After all, you have a highly reliable disposition to form true beliefs about edible mushrooms in your conditions. Were you to form a belief that a mushroom with a certain look is edible in your conditions, your belief would very likely be true. In the converse case in which jack-o'-lanterns exist nowhere except on the island you inhabit, your belief is unjustified. After all, you do not have a reliable disposition to produce true beliefs about edible mushrooms in your conditions. It is not the case that were you to form a belief that a mushroom with a certain look is edible in your conditions, your belief would very likely be true. Agent and condition relativity appear to give VR an important advantage here. We thus have some reason to believe that VR is on the right track.¹⁷

¹⁷ Consider also the following problem for PR, which an anonymous referee pointed out to me. Suppose that a small set of powerful agents are seeing to it that, say, the perceptual belief-forming processes of vast majority of agents are rendered highly unreliable. Since the epistemic oppressors use the same perceptual belief-forming processes as their victims, this means that they are harming themselves epistemically: they make it so that their own processes are not reliable and so will not deliver justified belief and knowledge. However, that seems to be the wrong result. It's easy to see that agent and condition relativity will solve this problem as well. In this way, cases like this one also serve to motivate VR vis-à-vis PR.

2.3 A Problem: Abilities without Agent Dispositions

While the general account of abilities that virtue reliabilists have offered carries a fair degree of promise, it also faces an important problem. To see how it arises, notice that an agent may have different ways of producing performances at his disposal. To keep things simple, let's say he has two such ways. For instance, you may produce shots in target archery with your left hand or with your right hand. Now, it may well be that a given agent who has two ways of producing performances of a certain kind at his disposal is disposed to perform successfully when producing performances of that kind in a certain set of conditions in one way but not the other. Call the former 'the good way' and the latter 'the bad way'. You may be disposed to hit the target when shooting with your right hand, but not when shooting with your left hand. Suppose, finally, that the agent is disposed to produce performances of said kind in the bad way. His dominant way of producing performances is the bad way. You took a sacred vow never to shoot with your right again. However, you still love practising archery. That's why you have taken to shooting with your left hand.

In this situation, there is a clear sense in which the agent is not disposed to perform successfully. According to the standard virtue reliabilist account of abilities, the agent does not have the relevant ability. For instance, you do not have the ability to hit the target in ARCH. After all, since you sacredly vowed never to shoot with your right again, if you were to produce a shot, you would produce a shot with your left. Since there are no conditions relative to which producing a shot with your left makes success highly likely, the trigger-manifestation conditional that corresponds to the ability to hit the target is false of you no matter what conditions you may be in. As a result, you do not have the ability to hit the target.

Crucially, this appears to be the wrong result. The mere fact that an agent is disposed to produce performances of a certain kind in a bad way does not entail that the agent no longer has the ability to produce successful performances of that kind, at least not so long as he still has the good way at his disposal. A vow never to shoot with your right again does not make your ability to hit the target when shooting with your right disappear. Given that this is so, something is amiss with the standard virtue reliabilist account of abilities.

I want to suggest that the problem here is that the kind of agent dispositions that the above general account of abilities countenances are too coarse-grained. What we need is something more fine-grained. One obvious proposal is to relativise abilities to ways of producing performances. Once we do so, we can allow that some ways of producing performances constitute abilities to succeed, whilst others don't. As a result, agents may have abilities to attain a certain kind of success, whilst not being disposed to exercise them. In fact, they may be disposed to exercise ways of producing performances that do not qualify as abilities. In that case, agents will not be disposed to attain the relevant kind of success, even though they do have the ability to do so. And, of course, this is exactly the situation you find yourself in after you have taken the vow never to shoot with your right again.

There is thus reason to relativise abilities to ways of producing performances. What does that mean for VR's account of justified belief as competent belief? Given that competent belief requires the exercise of an ability to form true beliefs and given that abilities are relative to ways of producing performances, this means that the abilities required for justified belief are relative to ways of belief formation. More specifically, to have an epistemic ability of the kind required by justified belief one must have a way of belief formation that disposes one to form true beliefs, at least when in suitable conditions.

But now notice that by relativising abilities to ways of producing performances, VR effectively incorporates a key idea from PR, to wit, that whether a belief is justified depends on properties of the process that produced it. Of course, this is not to say that VR collapses into PR. After all, abilities still remain agent and condition relative as well. In consequence, VR can still secure the advantage vis-à-vis PR outlined above, to wit, that VR avoids the problem of world-bound reliability. Moreover, VR still serves to rationalise PR. After all, the relevant reliability condition on justified belief still drops out of a general account of competent performances. However, once we have a version of VR that adopts a key idea from PR, the question naturally arises as to whether the view falls prey to the original problems PR encountered.

3 Competent Performances

In what follows, I will prepare the ground for addressing this question. More specifically, I will develop a detailed account of competent performance that takes on board the suggestion that the abilities required for competent performances are relative to ways of producing performances. I will then apply this general account to the case of belief and derive a detailed virtue reliabilist account of justified belief. Finally, I will argue that this account can either solve or defuse the problems of process reliabilism I introduced in section 1.

3.1 Simple Goal-Directed Practices

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.

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').

I agree with the general account of abilities offered by extant versions of VR that abilities, including SGP abilities, involve dispositions and that, as a result, they are relative to conditions, have trigger and manifestation conditions, and correspond to trigger-manifestation conditionals.

Crucially, I would like to suggest that SGP abilities involve *ways of move production*. Your ability to hit the target involves a way of shooting. As I already indicated, 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.

Since SGP abilities are relative to ways of move production, so are the dispositions at issue in them. And, of course, the same applies to the conditions to which the dispositions are relative. Accordingly, my suggestion is that a practitioner's way of move production, *W*, qualifies as an ability to attain success in a given SGP, *S*, relative to conditions *C* only if he is disposed to attain *S*'s success when using *W* in *C*. A given way of shooting you may have will qualify as an ability to hit the target in ARCH relative to *C* only if you are disposed to hit the target when using it in *C*.

Not every way of move production that disposes one to attain success in a certain SGP and in certain conditions qualifies as an SGP ability. Suppose I am an extravagant archer. The only occasions on which I do is when I see a sculpture by my favourite sculptor. When I do take a shot, I fire it right up in the air. As it so happens, the only remaining sculpture is located at a shooting range that is manipulated by an army of clandestine helpers who will see to it that all and only shots fired right up in the air hit the target. Currently I am at the shooting range. I have noticed the sculpture and have fired arrows straight up into the air. I don't stay in order to verify whether my shot actually hit the target. Why should I? There is no reason for this. However, my shot hits the target. In fact, in this case I am strongly disposed to hit the target when using my way of shooting in the conditions that obtain at the range where I take these shots.

Does my way of shooting qualify as an SGP ability, if only relative to those conditions? I take it to be intuitively clear that the answer to this question is no. Following Ruth Millikan [2000: 61] I would like to suggest that there is a distinction between mere dispositions and genuine abilities. Not all dispositions to attain success qualify as abilities. My disposition to hit the target in the above case

is a case in point. In consequence, the above picture of SGP abilities as ways of move production that dispose agents to attain success in certain conditions needs refinement.

What makes the difference between a genuine ability and a mere disposition? Taking another leaf from Millikan [2000: ch. 4], I want to say that etiology matters. Here is Millikan's proposal:

In general, the conditions under which any ability will manifest itself are the conditions under which it was historically designed as an ability. These are conditions in which it was learned, or conditions in which it was naturally selected for. They are conditions necessary to completing the mechanisms by which past successes were reached by the systems or programs responsible for the abilities.

[Millikan 2000: 61]

Here is a natural way of connecting Millikan's proposal with the picture of SGP abilities sketched above: to qualify as an SGP ability, a way of move production must have led one to successful SGP moves in the conditions in which it was acquired by learning/selected for. If so, since my extravagant way of shooting arrows up in the air did not lead me to hits in the conditions in which I acquired it, it does not qualify as an ability.

Unfortunately, however, this way of connecting Millikan's proposal with the above picture of SGP abilities encounters a problem. As Millikan also acknowledges, abilities can be acquired by coming to know that a way of producing moves works, i.e. by coming to know that a way of move production leads to success in certain conditions [Millikan 2000: 64]. The problem is that when one acquires an ability in this way, one need not acquire a new way of move production. Rather, one may simply learn something about an existing way of move production. For instance, suppose that, in the above case, I am told that, at the relevant shooting range, I will hit the target if and only if I fire my shots straight up in the air. I now possess the ability to hit the target at that shooting range. At the same time, I have not acquired a new way of move production. Rather, I learned something about an existing way of move production. The problem is that the way of move production at issue in my SGP ability never led me to success in the conditions in which *it* was acquired by learning.

Here is a way of addressing this problem. In certain cases, agents acquire an SGP ability by acquiring a new way of move production. In the learning process, the way of move production gets its shape. The process involves interaction with the environment during which the way of move production becomes "tuned" [Millikan 2000: 63] to producing successful SGP moves in the conditions that obtain during the learning process. This is what happened during your training period, after you had first taken up arching. You practised shooting, adjusting your way of performing in the light of past successes, tips from your coach and so on. You underwent a learning process that shaped your way of shooting with the result that using this way of shooting now disposes you to produce hits in the conditions of

learning. In this way, the learning process tuned your way of shooting to certain conditions. As a result, you now have the ability to hit targets in ARCH in those conditions. Let us say that, in this kind of case, the agent's way of move production is 'metaphysically grounded'.

In contrast, in the case in which I am told that shooting arrows up in the air will lead to success at the relevant shooting range, I do not acquire a new way of shooting. Rather, I learn something about an existing way of shooting. That said, there is a clear sense in which my new knowledge has the same effect here as your training did. It also tunes my way of move production to producing successful SGP moves in certain conditions. Crucially, the tuning is of a rather different kind. In contrast with the metaphysical kind of tuning we found in your case, here the tuning is epistemic in nature. Accordingly, let us say that, in this kind of case, the agent's way of move production is 'epistemically grounded'.

I want to suggest that in order to qualify as an SGP ability, a way of move production must have been tuned by learning or natural selection to some conditions, thereby grounding the way of move production, be it metaphysically or epistemically. This grounding condition on abilities enables my account to distinguish between genuine abilities and mere behavioural dispositions.¹⁸

I also want to allow that grounded ways of move production can qualify as SGP abilities for SGPs and conditions to which they have not been tuned. Suppose ARCH is practised only in strongly controlled conditions: indoors, under very specific artificial lighting conditions, whilst completely sober, etc. (henceforth *C*). Suppose you, the reigning world-champion of ARCH, are trained in *C* with the result that your way of shooting is tuned to *C*. Even so, it is hard to deny that you may have the ability to hit targets in ARCH in different conditions (e.g. outdoors, under different lighting conditions, after a beer = *C'*), at least so long as your way of shooting continues to dispose you to hit targets in ARCH in those conditions. Similarly, even if your ability is tuned specifically to ARCH, there can be no question that you may have the ability to hit targets in certain other SGPs (e.g. in ARCH' which is just like ARCH except that the target is a square rather than a disc), at least so long as your way of shooting continues to dispose you to hit targets in those SGPs. SGP abilities are thus relative to a *range of SGPs and conditions* that may differ from the SGP and conditions to which the underlying way of move production had been tuned. What matters is that the way of move production continues to dispose the agent to produce successful moves in the relevant SGPs and/or conditions.

Finally, the SGPs and conditions to which an ability is tuned and the SGPs and

¹⁸ It may be worth pointing out that, contrary to what Millikan suggests, these conditions need not be the condition *in* which the ability was acquired by learning. To see this, suppose I, the extravagant archer, am not at the shooting range when I am told that shooting arrows up in the air will produce successful shots there. I have now acquired the ability to hit the target at that range. Evidently, in this case, the way of move production does not dispose me to produce successful moves in the conditions *in* which it was acquired (i.e. my present conditions). However, my way of move production disposes me to produce successful moves in the conditions *for* which it was acquired (i.e. the relevant shooting range).

conditions relative to which an agent may have an SGP ability may vary from one way of move production to another. One of your ways of move production—shooting with your left hand (W_1)—may be tuned to ARCH', slight intoxication and natural light, while another—shooting with your right (W_2)—is tuned to ARCH, sobriety and artificial light. It may also be that W_2 disposes you to produce successful moves not only in ARCH but also in ARCH', not only in artificial light, but also in natural light but only when entirely sober. In contrast, W_1 may work only for ARCH' and only in natural light, no matter whether slightly intoxicated or entirely sober. Further extensions and other combinations are of course possible.

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

SGP Ability, General. One has an ability to attain success in a range, R_S , of SGPs and relative to a range, R_C , of conditions if and only if one has a grounded way of move production, W , such that, for any $S \in R_S$, there is some $C \in R_C$ such that using W in C disposes one to attain success in S , and for any $C \in R_C$, there is some $S \in R_S$ such that using W in C disposes one to attain success in S .

While this offers a fully general account of SGP abilities, it has the disadvantage of being rather complex. At the same time, for present purposes, I rarely need the account in its full generality. For that reason, I will be working with the following slightly simplified version of the account:

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 grounded way of move production, W , such that, for any $S \in R$, using W in C disposes one to attain success in S .

3.3 Exercises of Abilities

What does it take to exercise an ability? While virtue reliabilists have developed detailed accounts of ability, they have rarely stopped to take up this question. I would like to supply this lack. In particular, 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 .

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 does not hit 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 grounded 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 grounded 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 the 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. For a move to be competent, 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.¹⁹ 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.

¹⁹ An even clearer example may be the following: Suppose you have a grounded way *W* of producing layups in basketball that qualifies as an ability to score relative to some *C*. Currently you are standing at the midcourt 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. By the relevant instance of *SGP Exercise*, you exercise an ability to score. However, your shot is not competent. The shot you are taking is not within the range of the ability you exercise.

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 $S \in R$.

4 A Different Virtue Reliabilism

4.1 The View

With the general account of competent moves in SGPs in play, I will now move on to the application to the case of belief. 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*.²⁰ 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 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.²¹ For instance, a belief that p stands in the designated relation to the target of an inquiry into whether p if and only if its content, i.e. the proposition that p , is identical to the true answer, i.e. the true member of the set including proposition that p and the proposition that not- p . It is easy to see that this gives us the standard virtue reliabilist account of successful belief according to which a belief is successful if and only if true.

With the account of successful belief in play, we can now apply the above accounts of SGP abilities, their exercises and competent moves to the case of belief. This gives us:

Epistemic Ability. One has an ability to form true beliefs about propositions in a range, R , and relative to conditions, C , if and only if one has a grounded way of belief formation, W , such that, for any $p \in R$, using W in C disposes one to form true beliefs that p .

²⁰ 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, the issue of suspension of judgement can safely be set aside.

²¹ Note that the view can also easily be given a knowledge first spin. All we need to do is maintain that the designated relation is the knowledge relation.

Exercises of Epistemic Abilities. One exercises an ability, A , to form true beliefs about 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 form true beliefs about propositions in range R and relative to conditions C such that $p \in R$.

Recall that, according to extant virtue reliabilist accounts of justified belief, a belief is justified if and only if competent. I would like to adopt this account of justified belief as competent belief. In conjunction with *Competent Belief*, it entails the following detailed account of justified belief:

*VR-JB**. One justifiably believes that p if and only if one's belief that p is formed by an exercise of an ability to form true beliefs about propositions in range R and relative to conditions C such that $p \in R$.

It is important to keep in mind that, while VR-JB* does not explicitly feature a reliable process condition, it is implicitly present in it. After all, the notions of ability and their exercise at issue in the right hand side of VR-JB* are analysed in terms of *Exercise of Epistemic Ability* and *Epistemic Ability*. And, according to *Epistemic Ability*, epistemic abilities require reliable ways of belief formation. Since VR-JB* does feature a reliable process condition on justified belief, the question as to whether VR-JB* succumbs to the classical problems of PR remains. This question will be taken up in the remainder of this section.

4.2 The Problems for Process Reliabilism Revisited

Recall that PR encounters at least four problems: the new evil demon problem, the problem of clairvoyant cases, the generality problem and the problem of world-bound reliability. I have already shown how virtue reliabilism can avoid the problem of world-bound reliability. Since VR-JB* has little to add on this front, I will not return to this problem. Rather, I will restrict my focus on the remaining three problems, starting with the new evil demon problem.

THE NEW EVIL DEMON PROBLEM

Recall the new evil demon problem: you are abducted by a mad scientist and turned into a radically deceived BIV. As a result, many of your belief forming processes are now highly unreliable. Even so, many of the beliefs you go on to form are intuitively justified.

In a nutshell, the reason why VR-JB* can solve the new evil demon problem is (i) that reliability is relative to conditions of grounding and (ii) abilities can be exercised in unfavourable conditions. Let me explain how this pans out by means of an example: having recently been envatted, you appear to be taking a reading from a clock and thereupon come to believe that it is 7:45. Since you are radically

deceived, your belief is false. In fact, your process of forming beliefs about the time is now highly unreliable.

According to VR-JB*, what is going on here is that you acquired a way of forming beliefs about the time in much the same way as other normal (non-deceived) human beings, i.e. via interaction with the world that shaped your belief forming process with the result that using it disposes you to form true beliefs about the time in the conditions of learning. You have grounded way, *W*, of forming beliefs about the time that disposes you to form true beliefs about the time when using it in the conditions of learning, *C*. By *Epistemic Ability*, you have the ability to form true beliefs about the time in *C*. What happens when you are abducted and envatted is that you are being moved from *C* to different conditions *C'*. These conditions are of course highly unsuitable for your ability to form true beliefs about the time in the sense that using *W* in *C'* does not dispose you to form true beliefs about the time. Crucially, however, having been moved to *C'* does not prevent you from continuing to form beliefs via *W*. In particular, your belief that it is 7:45 was produced by using *W*. By *Exercise of Epistemic Ability*, you continue to exercise your ability to form true beliefs about the time. What's more, you form a belief that it is 7:45, which is a belief about the time and so falls within the range of the ability exercised. By VR-JB*, it follows that your belief is justified.²²

CLAIRVOYANT CASES

Let's move on to clairvoyant cases. Recall that, in our toy case, you take to forming beliefs via a 'clairvoyance' belief forming process that produces true beliefs about distant events with a high degree of reliability. Even so, the beliefs you go on to form via this clairvoyance process are intuitively not justified.

The key to VR-JB*'s account of clairvoyant cases is the grounding condition on abilities. In particular, I want to suggest that even though your clairvoyant process produces true beliefs with a high degree of reliability, it is not grounded. After all, it is not the case that you underwent a learning process involving interaction with the environment during which your clairvoyant process was tuned to

²² What about agents who are born BIVs? Such agents plausibly never have the chance to acquire the ability to recognise e.g. sunshine. However, since they cannot exercise abilities they do not have, they are not even in a position to acquire justified beliefs about the presence of sunshine. It might be objected that, as a result, the present solution to the New Evil Demon problem is less than fully satisfactory. While I agree that the solution to the New Evil Demon problem is incomplete, I don't think it is unsatisfactory. This is because it can be supplemented by a content externalist solution to the problem of scepticism that arises from cases of agents who are born BIVs. According to content externalism, agents who are born BIVs may even acquire *knowledge* of their enviroing world. It's just that their concepts and thoughts have different contents than ours [see e.g. Putnam 1981]. Notice also that when agents who have not been born BIVs have been envatted long enough, their thought contents may change. When this happens the character of their epistemic abilities also changes. They may once again be in a position to know propositions about their environment. However, just like agents who are born BIVs, the propositions these agents are now in a position to know are different than the ones inhabitants of normal environments routinely come to know.

producing true belief in the conditions that obtained during the learning process. This means that your clairvoyant process is not metaphysically grounded. Moreover, you also do not have any evidence that your clairvoyant process works and so it is not epistemically grounded either. Since your clairvoyant process is neither metaphysically nor epistemically grounded, it does not satisfy the grounding requirement on abilities. In consequence, it does not qualify as an ability to form true beliefs about any range of propositions. Since you cannot exercise an ability you do not have, you do not form a belief via an exercise of an ability to form true beliefs. By VR-JB*, your belief is not justified.²³

THE GENERALITY PROBLEM

Recall that the generality problem for PR arises from the fact that every token process instantiates indefinitely many types. Many of these process types have different degrees of reliability, some differ rather dramatically in the degree of reliability they achieve. Unless it is specified which process types exactly are the ones at issue in PR, PR simply does not make determinate predictions about whether or not individual beliefs are justified.

I'd like to flag that I do not mean to offer a head-on solution to the generality problem here, not in the least because a full solution would certainly outstrip the confines of this paper. Rather, my strategy will be to try and defuse it.

A large step in the direction of defusing the generality problem has already been made in [Comesaña 2006] and [Bishop 2010]. Comesaña and Bishop both argue, convincingly to my mind, that the generality problem is a problem that affects not just process reliabilism but all theories of justified belief. If so, the generality problem is a general epistemological problem, rather than a problem for process reliabilism in particular.

In addition, I will now argue that the generality problem is a problem in the general theory of performance normativity as it is a problem that arises for competent performances in general. This is good news especially for VR-JB*. After all, if the generality problem arises already in the general theory of performance normativity, it will not be a specifically epistemological problem. As far as VR-JB* is concerned, then, the generality problem is not only a problem that every theory of justified belief encounters, but also a problem that finds its proper home in the general theory of abilities. What's more, I will provide reason to believe that, as a problem in the general theory of performance normativity, it is not particularly worrisome in the first place.

²³ It may be worth noting that a knowledge first version of VR will have an even better explanation of clairvoyant cases. Even if one is tempted to say that, in the above case, you do have an ability to form true beliefs about distant events, it is hard to deny that you don't have an ability to know things about distant events. After all, while your way of belief formation does dispose you to acquire true beliefs here, it does not dispose you to acquire knowledge. Without a disposition to acquire knowledge, however, you simply cannot have an ability to know. Given that, on a knowledge first version of VR, a belief is justified only if it is produced by an ability to know, clairvoyant beliefs will not be justified. In this way, adopting a knowledge first version of VR may serve to strengthen the argument here.

To begin with, let's take a look at how the generality problem arises for VR-JB*'s account of justified belief as competent belief: According to VR-JB* competent beliefs require exercises of abilities to form true beliefs and exercises of abilities to form true beliefs require uses of ways of belief formation. More specifically, they require uses of tokens of types of ways of belief formation. But, of course, each such token instantiates indefinitely many types. For instance, my belief that I am sitting at my desk is produced by a token way of belief formation that instantiates the types cognitive process, visual belief-forming process, etc. To make determinate predictions about whether a given belief is competent, then, we need a specification of which of these types is the one at issue in VR-JB*'s account of competent belief. This is the generality problem as it arises for VR-JB*.

It is now fairly easy to see that the generality problem arises for competent performances in general. To see this notice that competent SGP moves in general require exercises of abilities and exercises of abilities require uses of ways of move production. More specifically, they require uses of tokens of types of ways of move production. But, of course, each such token instantiates indefinitely many types. Consider, once more, a case in which you take a shot at a target in ARCH. Your token way of shooting instantiates all of the following types: shooting, shooting with a bow, shooting with a bow and arrow, etc. To make determinate predictions about whether a given shot is competent, then, we need a specification of which of these types is the one at issue in the general account of competent performance.

In this way, it comes to light that the generality problem arises for competent performances in general. The generality problem, insofar as it arises for VR-JB*, is a problem in the general theory of performance normativity, rather than a problem that affects VR-JB*'s account of justified belief in particular.

Finally, here is why the generality problem as it arises for competent performances in general is not particularly worrisome. Notice first that we typically have no special difficulties in discriminating agents who have certain abilities from those who do not. For instance, it's not difficult to distinguish agents who have the ability to hit the target in ARCH from those who don't. The same goes, once again, for the epistemic case. It is not too difficult to find out who has the ability to recognise apples, BMWs or Picasso's and who doesn't. Crucially, we can do so without being able to offer a precise account of how the process at issue in the ability is typed. For instance, we can identify that you have the ability to hit the target in ARCH without being able to pinpoint the exact process type at issue in your ability. And the same goes for your ability to recognise apples.

As a result, VR-JB* can defuse the generality problem. According to VR-JB*, justified belief does require reliable belief forming processes of sorts. However, it does so because justified beliefs are beliefs produced by epistemic abilities and epistemic abilities involve reliable belief forming processes. Since we can identify abilities, including epistemic ones, without being able offer a general account of how the process at issue in the ability is typed, we do not need a solution to the

generality problem in order to have a workable version of VR-JB*.²⁴

5 Objections and Replies

Before closing, I would like to discuss a couple of issues that VR-JB* might be thought to encounter.²⁵ The first two effectively take the form of objections to process reliabilism, which VR-JB* is then thought to inherit, thanks to its concession to the former. In contrast, the third is less of an objection than an attempt to help Sosa to a comeback.

5.1 Modus Ponens Inferences

The first objection concerns cases of immediate inference. To facilitate discussion, I will work with the following example: you infer q from p and *if p , then q* by an application of modus ponens. Now here is the critic's remark:

Processes are essentially temporally extended, whereas an immediate inference must be punctiform in time, so this modus ponens inference is not a process.

I take it that what the critic is driving at here is that PR will run into trouble when it comes to cases of justified belief based on modus ponens inferences. According to PR, a belief is justified only if it is produced by a reliable belief-forming process, which is essentially temporally extended. However, your modus ponens inference is punctiform—that is, not temporally extended—and hence cannot be a process. Hence, PR cannot allow for the possibility of justified belief based on modus ponens inference.

To begin with, for the purposes of this paper, I will grant the critic the assumption that your modus ponens inference (henceforth MPI) is indeed punctiform. Moving on to my response to this objection, I'd first like to focus on a version of the case in which you instantiate MPI but do not believe q based on MPI. Rather, you believe q based on wishful thinking. It is widely agreed that, in this version of the case, your belief in q is not justified. As a result, instantiating MPI is not sufficient for you to justifiably believe q .

Let's now ask what is missing. The overwhelmingly plausible answer is that, in addition to instantiating MPI, you must acquire a belief in q based on MPI, where this requires that your belief in q is causally related to MPI in the right way. That

²⁴ An anonymous referee suggested that PR also faces a separate problem concerning the individuation of the length of the processes at issue in PR. While I am not sure whether this problem is indeed a separate problem, suffice it to say that the above defusing strategy for the generality problem will also work for the problem of how to individuate the length of the relevant processes. After all, first, the problem also already arises in the general theory of normativity. And, second, the problem is no more worrisome than the generality problem. After all, in order to find out who has relevant abilities, we don't need to be able to offer a precise account of just how long the process at issue in the ability is supposed to be.

²⁵ Thanks to a set of anonymous referees for pressing me on these issues.

is to say, what's needed for you to justifiably believe q once you have instantiated MPI is that your belief in q is the result of a causal process that involves MPI in the right way.

Given that this is so, the fact that MPI is punctiform does not constitute a problem for process reliabilism. Here's why. The critic is of course right in claiming that MPI, since punctiform, cannot be *identical* to the kinds of causal process that, according to process reliabilism, are required for acquiring justified beliefs. Crucially, however, there is independent reason to think that acquisition of a justified belief in q requires more than just instantiating MPI. More specifically, it is independently plausible that what's needed in addition is that your belief in q is produced by just the kind of causal process that process reliabilists claim is required for the acquisition of justified beliefs. And even though MPI cannot be identical to a causal process of this kind, it is entirely compatible with process reliabilism that MPI is part of such a kind of process. In fact, it is overwhelmingly plausible that in any version of our toy case in which you do arrive at a justified belief in q based on MPI, this is exactly what happens. Once we are clear that justified beliefs in the conclusions of modus ponens inferences involve punctiform inferences, without being exhausted by them, the supposed problem for process reliabilism disappears. Since VR-JB* is said to inherit this problem from process reliabilism, if the latter walks free, so does the former.

5.2 The Cogito

Let's move on to the second problem, then. Whilst involving a different case, the thrust of the objection is similar to the first:

[Process reliabilism] would reject the cogito, since the important competence exercised in Descartes's belief is an infallible one that is instantaneous in its operation, since relevant competences pertain to the very moment of the belief <I think> or <I exist>. What happens prior to that moment might have been missing with no effect on the quality of the process and its reliability. What matters is that at the very moment when the believer believes a proposition with that content the belief must be true. The connection requires the absolute temporal coincidence. What happens prior to the believing is irrelevant to the infallible reliability required for certainty.

Unsurprisingly, my response here is similar to the one offered above. First, consider a version of the case in which Descartes believes <I think>, say, but not based on reliable introspection (or perhaps rational intuition). Rather, his belief is produced by an evil demon who causes Descartes to switch back and forth between believing <I think> and <I do not think>. Here, Descartes's belief is not justified.

This means that simply hosting a belief in <I think> is not sufficient for Descartes to justifiably believe <I think>. Rather, his belief in <I think> must be properly based, e.g. on reliable introspection (rational intuition). However, this requires

that his belief in <I think> is causally related to reliable introspection (rational intuition) in the right way. What's needed for Descartes to justifiably believe <I think> is that his belief in <I think> is produced in the right way by the right kind of causal process, e.g. one involving reliable introspection (rational intuition).

These considerations suggest that the critic is not quite right in claiming that what happens prior to the moment of believing doesn't matter to whether Descartes's belief in <I think> is justified. Rather, whether Descartes's belief in <I think> is justified depends crucially on whether it is produced by the right kind of causal process and thus on what happens prior to the moment of believing.²⁶ As a result, the cogito also doesn't produce an insurmountable problem for process reliabilism. If so, the same goes for VR-JB*.^{27,28}

5.3 Abilities and Competences

Finally, consider the following response on behalf of Sosa to the problem posed by cases like the vow-case from section 2.3:

Sosa very rarely *if ever* appeals to ability or abilities. He focuses rather on competence, and would argue that the agent who would too rarely shoot with the right hand is *not* really competent to attain success (reliably enough). He does intuitively have the *ability* to do so, but lacks the *competence*. It might be argued that the view is refuted because *if*, in the archer case, the agent happens to use his right hand exceptionally, he would surely succeed aptly, even though he exercises only an ability and not a competence. But this is not very plausible. Compare the would-be bank robber who has the ability to open the safe by entering the right combination (the *bare* ability, which he has simply by having the proper use of his fingers with the lock available to his operation). If he happens to exercise that *bare* ability on a certain occasion, that is not apt success. What he needs is something approaching a competence: i.e., the disposition to try in what is likely enough to be the right way, so that nearly all the many, many possible combinations are ruled out as ones he *would not* employ.

²⁶ Compatibly with that, it may well be that the process that actually produces Descartes's belief in <I think> is infallible in the sense that Descartes's belief in <I think> could not be produced by *this* process and produce a false belief in <I think>.

²⁷ See [Goldman 1979: section 1] for a similar line of argument.

²⁸ Note that other cases that may be thought to involve instantaneously operating competences and so cause trouble for PR can be treated in the same way. Consider, for instance, a case in which Descartes forms a basic a priori belief in <two plus two is four> that is justified by rational intuition where this is thought to involve an instantaneously operating competence. Note, next, that Descartes might also believe <two plus two is four> thanks to the workings of a demon who causes him to switch back and forth between believing <two plus two is four> and <two plus two is not four> in which case his belief is not justified. But given that this is so, we can now adapt the treatment of the cogito case sketched above to the case at hand simply by replacing all occurrences of '<I think>' by '<two plus two is four>' and of 'reliable introspection (rational intuition)' by 'rational intuition'. Thanks to an anonymous referee for pressing me on this issue.

I think it's quite common for folk to use 'competence' and 'ability' interchangeably and that's how I have been using 'ability' so far. Of course, this is not to say that there aren't important distinctions to be made here. Crucially, once we start distinguishing more carefully between various such properties, it's important to keep in mind that what we need is a property fit to feature in the accounts of competent and apt performance at issue in the framework for the normativity of performances that champions of VR invoke.

Now, I agree that competent and apt performance cannot plausibly be unpacked in terms of any property of ability so minimal that even the bank robber in the critic's case possesses it (following the critic I will henceforth use 'bare ability' to refer to this property of ability). I am also completely on board with the critic's argument for this. When the bank robber exercises his bare ability, he will not perform aptly, no matter whether or not he succeeds and what conditions he may find himself in.

At the same time, I'd insist that the property of competence fares no better. To see why, note first that Sosa takes competences to require a disposition to succeed were one to try. As the critic expressly acknowledges, this means that, in the vow case, once you have taken the vow never to shoot with your right again you no longer have the relevant arching competence. But now consider the version of the vow case the critic also gestures towards: at some point after having taken your vow you suffer a one-off lapse. Perhaps you think to yourself: "Whatever, it's so nice to take a proper shot and once doesn't count anyway!" and take a shot with your right. Suppose that conditions are favourable and your shot hits the target right in the centre. It is no less plausible that, in this case, your shot *is* apt, than it is plausible that, in the critic's bank robber case, the bank robber's attempt at opening the safe is *not* apt. Since your shot is apt but you don't have the arching competence, there is reason to believe that the property of competence is too strong for the purpose of analysing competent and apt performance at issue in our normative framework.²⁹ In this way, there is reason to believe that Sosa's proposal to unpack competence and aptness in terms of competences meets with the same fate as the proposal to do so in terms of bare abilities: neither proposal will work. Of course, the two fail for very different reasons. While bare abilities are too weak for our purposes, competences are too strong.

But wait! If you don't have the competence, isn't there something important that you are missing? Consider the vow case once more. Once you don't have the competence, isn't there a clear sense in which we cannot trust you with taking shots any more? And doesn't this show that what we should be interested in is a competence rather than an ability?

I do not mean to say that competence and aptness are the only normative property that are of interest when it comes to evaluating performances. There may well be others, including some that require the possession of a competence. However, the fact remains that apt and competent performance are two important normative

²⁹ Or, to be more precise, it will be too strong if it is unpacked in the way Sosa suggests. Note that in what follows I will take competence to be unpacked in just this way.

such properties. Since not even apt performance—the logically stronger one the two—requires possession of a competence, neither does competent performance. Given that champions of VR have proposed to identify knowledge with apt belief and justified belief with competent belief, this means that neither knowledge nor justified belief requires possession of a competence. Whatever normative property may be missing when we do not have a competence is at any rate of little relevance to the intended application of the normative framework to epistemology.

But isn't that too quick? Why think that knowledge and justified belief do not require possession of a relevant competence rather than that the relevant normative properties to be identified with knowledge and justified belief are not apt and competent belief after all?

By way of response, note that there is independent reason against identifying knowledge and justified belief with any normative property that requires possession of a competence. To see this consider the following case, which is an epistemic analogue of the vow case. For the longest time you have acquired information on a certain issue—Brexit, say—from a newspaper you know to be highly reliable—*The Space*, say. Suppose also that *The Space* is the only reliable newspaper available to you. For some reason, however, you recently took a vow never to read *The Space* again. Instead you have taken to acquire information from a newspaper that you know to be unreliable—*The Moon*, say. Since you are no longer disposed to form beliefs about Brexit that are likely enough true when forming beliefs about Brexit at all, you no longer have the relevant epistemic competence. Now suppose it so happens that you suffer a one-off lapse. Perhaps you think to yourself: “Whatever, it's so nice to read a reliable newspaper and once doesn't count anyway!” and pick up a copy of *The Space*. One of the articles states that citizens from other EU countries make up only 5% of the UK population and you form the corresponding belief. Since it is hard to deny that your belief qualifies as knowledge here, there is independent reason to think that knowledge does not require the possession of the relevant epistemic competence anyway. And since knowledge entails justified belief, the same goes for justified belief. What comes to light, then, is that normative properties requiring possession of a relevant competence, interesting as they may be, are at any rate not suited for our epistemological purposes of analysing knowledge and justified belief.

Still our critic may not be satisfied. After all, if the only options at our disposal for spelling out the accounts of competent and apt performance are in terms of either bare abilities or competences, we might still opt for competences as the lesser of two evils.³⁰ Fortunately, however, there is reason to believe that these are not our only options. In particular, as I will argue momentarily, there is a further sense of ‘ability’ such that the property of ability it picks out promises to be better suited for our normative purposes.

Let's first try and home in on the sense of ‘ability’ I have in mind by looking at a case in which it is in play. Suppose while deliberating on whom to ask to take

³⁰ In fact, it might be thought that this is exactly the point the critic is trying to make.

a shot, the following conversation ensues between *X* and *Y* (for simplicity, I will assume that your name is *You* and mine is *I*):

X: Let's ask *I*. He lives right downstairs. He might be able to help us.

Y: No. *I* is blind since his eye infection last month. He doesn't have the ability to make the shot.

X: I guess you are right. I didn't know that *I* had gone blind. Let's ask *You* then. He surely has the ability to make this shot. After all, he is the reigning world champion in archery.

Y: *You* has the ability alright. In fact, he is one of the finest shots ever—with his right. Sadly, he recently took a vow never to shoot with his right again. And he's a terrible shot with his left. So, it's pointless to ask him.

Here *X* and *Y* deny me the ability to hit the target, whilst attributing it to you. Since both denial and attribution are serious, natural and appropriate, whilst not being based on any false belief about some underlying matter of fact, this constitutes evidence that they are both true.³¹

But now note that if we grant that the bank robber has the bare ability to open the safe simply in virtue of having proper use of his fingers and so may, by sheer luck, punch in the right numbers and open the safe, then we must also grant that *I*, the blind man, have the bare ability to hit the target. After all, *I* have proper command of my arms and hands and so may, by sheer luck, fire a shot that hits the target. Since we have evidence that the denial of ability to me is true in this case, there is reason to think that 'ability' does not have the relevant bare ability as its semantic value here. Moreover, we have already seen that once you have taken the vow you no longer have the relevant arching competence. Since there is evidence that the attribution of ability to you is true, there is reason to think that 'ability' does not have the relevant competence as its semantic value either. Since the sense of 'ability' in play here has neither the property of bare ability nor the property of competence as its semantic value, there is reason to believe

³¹ Keith DeRose describes the methodology that supports this claim in the following passages:

Since the speaker is using the term seriously and properly, and this use isn't based on any false beliefs the speaker has about underlying matters of fact . . . this is a use I think a theory about the meaning of [the term] should try to make come out true, and it's a strike against a theory if it fails to do so.

[DeRose 2009: 18]

[The] 'methodology of the straightforward', as we may call it, takes very seriously the simple positive and negative claims speakers make utilizing the piece of language being studied, and puts a very high priority on making those natural and appropriate straightforward uses come out true, at least when that use is not based on some false belief the speaker has about some underlying matter of fact.

[DeRose 2009: 153]

that there is a further property of ability that's picked out by 'ability' in this sense. (To facilitate exposition, in what follows, I will use 'robust ability' to refer to this further property.)

Crucially, robust abilities still hold out the hope of being fit for featuring in the accounts of competent and apt belief we are interested in. On the one hand, unlike bare abilities, robust abilities are not present in cases like the bank robber's. As a result, the bank robber's attempt at opening the safe is bound to come out as falling short of aptness, even if it is successful. So, robust abilities can steer clear of the problem preventing bare abilities from being fit for featuring in the account of competent and apt belief. On the other hand, unlike competences, robust abilities may be possessed even by agents who are not disposed to succeed upon trying. As a result, when you take a shot with your right, you will still be exercising your arching ability, and are in the ballpark for performing aptly. So, robust abilities can steer clear of the problem preventing competences from being fit for featuring in the account of competent and apt belief also.³²

It comes to light that robust abilities are different from both the options the critic countenances, to wit, bare abilities and competences. Neither bare abilities nor competences are fit to feature in the accounts of competent and apt performance at issue in the framework for the normativity of performances that we are interested in. Robust abilities, in contrast, steer clear of the problems on both sides and so may very well be up to the task. And while we may have opted for unpacking competent and apt performance in terms of competences if our only options here had been bare abilities on the one hand and competences on the other, once it is clear that there is a third and more promising option available to us, there remains little reason for us to go down that route. As a result, the critic's attempt at helping Sosa to a comeback by proposing that competent and apt performance are to be unpacked in terms of competences fails.³³

6 Conclusion

This paper has developed a new virtue reliabilist account of justified belief as competent belief. This account takes on board a number of insights of extant virtue reliabilist contributions to the literature, including the idea that competent

³² It is not hard to see that the account of SGP abilities from Section 3.2 is fit to serve as an account of robust abilities for SGP performances. After all, SGP abilities behave relevantly like robust abilities. For instance, I, the blind man, do not have the SGP ability to hit the target in ARCH. After all, SGP abilities feature a reliability condition that I simply don't meet. In contrast, as we have seen, you, the expert archer, do have the relevant SGP ability and do so even after you have taken the vow.

³³ Note also that in order to do be fit for unpacking competent and apt performance, robust abilities must be taken to feature fine-grained dispositions, i.e. ones that are relativised to ways of performing. Were robust ability taken to feature a coarse-grained disposition of the kind favoured by Sosa, we'd be back to the problematic result that, in the version of vow case considered above, the performance is not apt. As a result, robust abilities would be too strong to be fit for unpacking competent and apt performance. The case for going more fine-grained thus stands as well.

performance requires exercise of ability and that abilities involve dispositions to attain success. Crucially, the view also departs from the general accounts of ability offered by virtue reliabilists in that it takes abilities to be relative to ways of performing. It is worth noting that the claim that abilities are relative to ways of performing means that the resulting virtue reliabilist account of justified belief incorporates a key process reliabilist idea. In fact, on the proposed view a version of a process reliabilist condition drops out of an independently plausible general account of abilities. In this way, the proposed account rationalises the core idea of process reliabilism.

At the same time, the question arises as to whether the proposed account thereby also inherits the problems that process reliabilism encounters. I have argued that the answer to this question is no. The proposed virtue reliabilist account of justified belief does have the resources to solve, or at least defuse, a number of key problems that have riddled process reliabilism. Since the account allows that epistemic abilities can be exercised in unfavourable conditions, it can accommodate the presence of justified belief in cases of radical deception. Since it countenances a grounding condition on epistemic abilities, it can successfully predict the absence of justified belief in clairvoyant cases. The generality problem is defused once it is recognised that the problem arises not only for epistemic abilities, but for abilities in general and that this version of the problem is not particularly worrisome.

Given that the proposed account virtue reliabilist account not only drops out of an independently plausible general account of abilities, their exercise and competent performance, but also avoids a number of the key problems that both process reliabilism and rival virtue reliabilist theories encounter, I submit that it is a promising view and deserves to be taken seriously.

References

- Beddor, B. 2015. Process reliabilism's troubles with defeat. *Philosophical Quarterly*, **65**, 145–159.
- Bishop, M. 2010. Why the generality problem is everybody's problem. *Philosophical Studies*, **151**, 285–98.
- BonJour, L. 1980. Externalist theories of empirical knowledge. *Midwest Studies in Philosophy*, **5**, 53–73.
- BonJour, L. 1985. *The Structure of Empirical Justification*. Cambridge/MA: Harvard University Press.
- BonJour, L., & Sosa, E. 2003. *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Malden/MA: Blackwell.
- Cohen, S. 1984. Justification and truth. *Philosophical Studies*, **46**, 279–96.

- Comesaña, J. 2006. A well-founded solution to the generality problem. *Philosophical Studies*, **129**, 27–47.
- Conee, E., & Feldman, R. 1998. The generality problem for reliabilism. *Philosophical Studies*, **89**, 1–29.
- DeRose, K. 2009. *The Case for Contextualism*. Oxford: Clarendon Press.
- Goldberg, S. 2010. *Relying on Others: An Essay in Epistemology*. New York: Oxford University Press.
- Goldman, A. 1979. What is justified belief? In: Pappas, G. (ed), *Justification and Knowledge*. Dordrecht: Reidel.
- Goldman, A. 1986. *Epistemology and Cognition*. Cambridge/MA: Harvard University Press.
- Goldman, A. 2012. *Reliabilism and Contemporary Epistemology*. New York/NY: Oxford University Press.
- Greco, J. 1999. Agent reliabilism. *Philosophical Perspectives*, **33**, 273–96.
- Greco, J. 2000. *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. Cambridge: Cambridge University Press.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Kelp, C. 2011. In defence of virtue epistemology. *Synthese*, **179**, 409–33.
- Kelp, C. 2016. Justified belief: knowledge first-style. *Philosophy and Phenomenological Research*, **93**, 79–100.
- Kelp, C. Forthcoming. Knowledge first virtue epistemology. In: Carter, A., Gordon, E., & Jarvis, B. (eds), *Knowledge First: Approaches in Epistemology and Mind*. Oxford: Oxford University Press.
- Kornblith, H. 2002. *Knowledge and Its Place in Nature*. Oxford: Clarendon Press.
- Kornblith, H. 2004. Does reliabilism make knowledge merely conditional? *Philosophical Issues*, **14**, 185–200.
- Kornblith, H. 2009. A reliabilist solution to the problem of promiscuous bootstrapping. *Analysis*, **69**, 263–67.
- Kvanvig, Jonathan. 2007. Two approaches to epistemic defeat. *Pages 107–124 of*: Baker, D. (ed), *Alvin Plantinga*. Cambridge University Press.
- Lehrer, K. 1990. *Theory of Knowledge*. London: Routledge.

- Lehrer, K., & Cohen, S. 1983. Justification, truth, and coherence. *Synthese*, **55**, 191–207.
- Lyons, J. 2009. *Perception and Basic Beliefs. Zombies, Modules, and the Problem of the External World*. Oxford: Oxford University Press.
- Millikan, R. 2000. *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.
- Pritchard, D. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy*, **109**, 247–79.
- Putnam, H. 1981. *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Sosa, E. 1980. The raft and the pyramid. *Midwest Studies in Philosophy*, **5**, 3–25.
- Sosa, E. 1991. *Knowledge in Perspective. Selected Essays in Epistemology*. Cambridge: Cambridge University Press.
- Sosa, E. 2010. How competence matters in epistemology. *Philosophical Perspectives*, **24**, 465–475.
- Sosa, E. 2011. *Knowing Full-Well*. Princeton/NJ: Princeton University Press.
- Sosa, E. 2015. *Judgment and Agency*. Oxford: Oxford University Press.
- Wedgwood, R. 2002. The aim of belief. *Philosophical Perspectives*, **16**, 267–97.
- Williamson, T. 2007. On being justified in one's head. In: Timmons, M., Greco, J., & Mele, A. (eds), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*. New York/NY: Oxford University Press.
- Williamson, T. 2011. Philosophical expertise and the burden of proof. *Metaphilosophy*, **42**, 215–29.
- Zagzebski, L. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.