

Introduction
Virtue Theoretic Epistemology
Christoph Kelp

There are three leading theories of normativity: teleology, deontology and virtue theory. All three types of normative theory countenance values, norms and virtues. What they disagree on is the order of explanation. Teleology takes values to be the fundamental normative kind and explains norms and virtues in terms of them. Deontology takes norms to be the fundamental normative kind and explains values and virtues in terms of them. And, finally, virtue theory takes virtues to be the fundamental normative kind and explains norms and values in terms of them.

To get a better feel for how this is meant to work, let's look at some dummy versions of teleological, deontological and virtue theories of normativity. According to our dummy teleological theory, one ought to ϕ if and only if ϕ -ing maximises value. And, similarly, S is a virtuous person if and only if S is disposed to maximise value. What makes this view into a distinctively teleological theory of normativity is that the direction of explanation goes from right to left with the result that norms and virtues are analysed in terms of value. What we ought to do and what makes for a virtuous person are explained in terms of value.

Our dummy deontological account has it that X has value if and only if one ought to favour X. And S is a virtuous person if and only if S has a disposition to do as S ought. This dummy view is deontological in virtue of the direction of explanation, which once again goes from right to left. What's of value and what makes for a virtuous person are explained in terms of what one ought to do.

Finally, and most importantly for present purposes, let's look at a dummy virtue theory of normativity. According to our view, X is valuable if and only if X would characteristically be favoured by a virtuous person, and one ought to ϕ if and only if ϕ -ing is what a virtuous person would characteristically do. What's important here is again the way in which the direction of explanation proceeds. What is of value and what one ought to

do is analysed in terms of virtue. This makes our dummy view a distinctively virtue theoretic approach to normativity.

As characterised so far, the debate between the three theories of normativity is extremely abstract. Things get a little more concrete once we apply these views to more familiar normative domains. The kind of domains I have in mind here include ethics, aesthetics and, most importantly for present purposes, epistemology. It is hard to deny that most of the research has been done in normative ethics.¹ While the debates with teleological and deontological approaches to ethics are of course important, it is worth mentioning that there are a variety of important strands of virtue ethics on the market.

According to the perhaps most popular form of virtue ethics, and the one I will be focusing on here, the idea of eudaimonia takes centre stage. Eudaimonia is often translated as flourishing, happiness or well-being. What is important to note, however, is that it is itself a value-laden notion. It is the kind of well-being that is worth having (Hursthouse and Pettigrove 2018). Consider, by way of illustration, the following case: eating lots of chocolate might give you lots of pleasure. Say that you spend your entire life doing nothing but eating chocolate. At the end of your life you look back with great satisfaction. You take yourself to have led a happy life. Even if in your chocolate-eating life you achieved a form of happiness, you did not achieve eudaimonia. The reason for this is that the happiness that derives from only ever eating chocolate is not a form of well-being worth having, at least not for a normal adult human being.

Now, the thought is that moral virtues play the central part in normative ethics that they do because of their relation to eudaimonia. Of course, the thought cannot be that eudaimonia is identified as a value that's independent of the moral virtues and that moral virtues are understood in terms of their relation to eudaimonia. After all, this would turn the view right into a teleological normative ethics. Rather, the idea is that the moral virtues are constitutive of and essential to eudaimonia. This means that we cannot fully understand what eudaimonia is without understanding the moral virtues. In this way, virtues play the foundational role in ethics that virtue ethics takes them to play.

For champions of this kind of virtue ethics, one central question will be which moral virtues are constitutive of and essential to eudaimonia. While

¹ Classical defences of teleological approaches to moral normativity include Bentham 1961; Mill 1998; Sidgwick 1907. For classical defences of deontological approaches, see e.g. Kant 1964; Ross 1930 and for virtue ethics, e.g. Aristotle 1985; Plato 2000.

there is of course ample room for debate, it is hard to deny that there are a number of character traits that are widely recognised as virtues, including charity, courage, generosity, honesty, justice and kindness, to name but a few. The morally perfectly virtuous agent possesses all the moral virtues to the highest degree. Less than morally perfectly virtuous agents are virtuous to the extent that they approximate the morally perfectly virtuous agent.

With these points in play, we can see how our virtue ethics might account for moral norms and values, at least on the dummy version outlined above. According to one straightforward way of implementing the view, we get the result that one morally ought to ϕ if and only if ϕ -ing is what a morally virtuous agent would characteristically do. To take a more concrete example, you morally ought not to lie if and only if this is what a morally virtuous person would characteristically do. Since lying is not something that a morally virtuous person would do (as the virtue of honesty would characteristically lead them in this direction), you morally ought not to lie. Crucially, what we find here is the direction of explanation that is key to virtue theories of morality. What explains why you morally ought not to lie is not, for instance, that it would violate a moral norm, as deontologists would have us think, but that it is something a morally virtuous person would not do. Moreover, X is morally valuable if and only if X would characteristically be favoured by a morally virtuous agent. For instance, organisations such as Doctors without Borders and UNICEF have moral value if and only if they are the kind of entity that a morally virtuous agent would characteristically favour. Since they are the kind of organisation that a virtuous person would characteristically favour (as the virtue of charity would characteristically lead them to do so), these organisations have moral value. Again, what's important here is the direction of explanation. It is not the fact that these organisations bring about a lot of moral value in the world that explains why they are morally good. Rather, it is the fact that they are likely to be favoured by a morally virtuous agent.

Let us now move on to virtue epistemology. If what we have just described captures the contours of virtue ethics at least roughly, it is tempting to think that the contours of virtue epistemology can be captured in much the same way. For instance, it is tempting to think that virtue epistemology will adopt a virtue-based account of epistemic normativity. It is also tempting to think that a virtue epistemology (will, or at least might) hold that intellectual virtues, like moral virtues, play the central part in normative epistemology they do because, just like moral virtues, they are partly constitutive of eudaimonia. After all, it is plausible that eudaimonia is constituted not only by moral virtues but also by intellectual virtues.

Again, there is ample room for discussion, but in the intellectual sphere as well, there is a wide range of character traits that are widely recognised as virtues, including attentiveness, curiosity, intellectual courage, intellectual humility, open-mindedness and intellectual thoroughness.

Unfortunately, there is reason not to adopt this straightforward way of characterising virtue epistemology. Or, to be more precise, there is reason not to do so if we want all paradigm cases of virtue epistemologies to be classified as virtue epistemologies. To see this, note that it is customary to distinguish between two different kinds of virtue epistemology: reliabilist and responsibilist.² The trouble is that reliabilist virtue epistemologies will simply not fit with the general theory of normativity that takes virtues to be fundamental. This is because reliabilist virtue epistemologies are very clear that at least some (perhaps all) epistemic virtues are virtues at least in part because they are conducive to producing some independently specified epistemic good or other. In other words, they offer a distinctively teleological account of epistemic virtues. And, of course, it is hard to see how this could still be compatible with the approach to general normativity that is distinctive of virtue theories of normativity and the approach to moral normativity that we found in virtue ethics.

Again, an example may help to make this point clearer. Consider a disposition that takes the look of a certain object as trigger and beliefs about the presence of chanterelles as manifestations. Reliabilist virtue epistemologists will typically allow for this disposition to count as an epistemic virtue. Crucially, the reason why they do so is that the disposition reliably produces true beliefs rather than false ones, at least in suitably favourable conditions. But, of course, if this is how we explain why certain dispositions count as epistemic virtues, we are effectively offering a teleological account of these epistemic virtues. It is hard to see how this can be squared with the approach to normativity distinctive of virtue theory according to which virtues play the foundational part.³

² Among the most influential defences of virtue epistemology are Greco 2010; Pritchard, Millar, and Haddock 2010; and Sosa 2015 on the reliabilist side and Code 1987; Montmarquet 1993; and Zagzebski 1996 on the responsibilist side.

³ It may also be worth noting that these virtues cannot be accounted for except in teleological terms. To see this, note that the very same disposition may not count as an epistemic virtue for agents who inhabit environments in which too many things that have the relevant look are not chanterelles (e.g. environments in which jack-o'-lanterns abound) and, as a result, the disposition does not produce true beliefs about the presence of chanterelles reliably enough. What comes to light is that whether one and the same disposition is an epistemic virtue may turn only on a difference in reliability with which it produces true belief. At the same time, it is hard to see how it can be the difference in the degree of reliability that determines whether one and the same disposition counts as an epistemic virtue unless we adopt a teleological account of epistemic normativity.

It might be thought that the take-home lesson is that reliabilist virtue epistemology is not *virtue* epistemology proper and that virtue responsibilism is the real thing. On the face of it, this may sound promising. And it gets even more promising once we appreciate the fact that virtue responsibilists are in practice often much closer to virtue ethicists than virtue reliabilists. Reliabilist virtue epistemologists are for the most part interested in offering accounts of the nature and value of core phenomena in epistemology such as knowledge and justified belief. Virtue responsibilists, in contrast, are often less interested in the analysis of knowledge and justified belief. What they often focus on instead is the nature of the wide range of intellectual virtues that we find among humans, including the likes of curiosity, open-mindedness and intellectual humility. It is much less clear that these virtues must be understood in terms of the good epistemic consequences they tend to produce. Rather, an account in terms of their relation to eudaimonia, perhaps along much the same lines as virtue ethicists envisage for the moral virtues, may still seem quite promising.

Does this mean that, as true virtue epistemologists, we will have to turn to virtue responsibilism? The answer to this question is no. This is because there is independent reason to think that a virtue theory of normativity isn't all that plausible for the epistemic domain in the first place. To see why not, note that on a virtue theory of normativity, virtues will *essentially* have positive normative status and not possessing a given virtue will necessarily reflect negatively on a given moral agent (or at least show that there is room for improvement). Note also that this is fairly plausible in the case of the moral virtues. That is to say, it is fairly plausible that, for instance, kindness necessarily has positive moral status and that not possessing this virtue will reflect badly on a moral agent. The trouble is that the same does not appear to be true in the case of epistemic agents. Perhaps the easiest way to see this is by considering infallible and omniscient epistemic agents. Infallible omniscient agents may fail to possess a range of core responsibilist virtues without this reflecting badly on them qua epistemic agents, nor even meaning that they would be improved qua epistemic agents if they had these virtues. Most importantly, they may not be curious, open-minded, intellectually humble or intellectually courageous. And yet, for all that, they may not be any worse qua epistemic agents, nor even better if they did have them. If this is right, there are intellectual virtues, even paradigm ones, which do not enjoy positive epistemic status essentially. And that does not fit well with a virtue theory of epistemic normativity.

By way of further support for the thesis that a virtue theory of normativity isn't all that plausible for the epistemic domain, consider an epistemic phenomenon that is widely believed to enjoy positive epistemic status essentially, to wit, knowledge. Given that knowledge enjoys positive epistemic status essentially, we may expect a virtue theory of epistemic normativity to explain why this is. The trouble is that the prospects for a satisfactory account are rather dim here. Knowledge is not a virtue in itself. And an account in terms only of the kinds of virtues that the responsibilist countenances doesn't seem promising either. You could be entirely intellectually non-virtuous (in the responsibilist sense) and still acquire perceptual knowledge, say that there is a chanterelle before you, at least in suitably favourable environments. Since your perceptual knowledge will still have positive epistemic status, an explanation in terms of responsibilist virtues seems unpromising.

What comes to light, then, is that a virtue theory of epistemic normativity seems unpromising. This may be a surprising result to find defended in the introduction to a volume on virtue epistemology. On the upside, it once again levels the playing field between the two most prominent kinds of virtue epistemology, reliabilism and responsibilism. The concern that only virtue responsibilism is a virtue epistemology in the proper sense because only virtue responsibilism can accord virtues the fundamental normative role required by a genuine virtue theory of epistemic normativity can be laid to rest.

What's more, there is another way of interpreting the central thesis of virtue theoretic approaches to epistemology, which is plausibly more promising. Virtue theories of normativity venture to analyse norms and values in terms virtues. For instance, according to our dummy view X is valuable if and only if X would characteristically be favoured by a virtuous person, and one ought to ϕ if and only if ϕ -ing is what a virtuous person would characteristically do. Even if we abandon an ambitious virtue theory of epistemic normativity like our dummy view, we might still hold out hope for a more modest research programme in epistemology. Specifically, even if values and norms in the epistemic domain cannot be accounted for only in terms of epistemic virtues, it may be that what it is for someone to believe or do something *well* epistemically is to be accounted for in terms of virtues. That is to say, it may be that believing or doing something well epistemically is believing or acting from virtue. And, *mutatis mutandis*, for a wide range of positive epistemic statuses for epistemically assessable ϕ -ing. To mark the distinction between the ambitious and the modest programme, I henceforth reserve the term 'virtue epistemology' for the

ambitious programme. This is in recognition of the fact that the programme that is designated by the term ‘virtue ethics’ in moral philosophy is of the ambitious variety. In addition, I will use the label ‘virtue theoretic epistemology’ for the modest programme. My proposal is that the bulk of research that has been carried out in recent epistemology is really best understood as falling within the purview of the modest programme. Given this terminological distinction, it is thus most aptly characterised as contributing to virtue theoretic epistemology.

It is not hard to see that a lot of the work that has been done in reliabilist virtue theoretic epistemology is indeed best understood along these lines. What has crystallised especially in the last decade or so is just how central a general theory of the normativity of performances is to virtue theoretic epistemology in the reliabilist camp. According to this theory, performances with an aim in general can be assessed along three normative dimensions: success, competence and aptness (e.g. Sosa 2015). Thus, for any performance with an aim, we can ask whether the performance attains its aim, i.e. whether it is successful. We can also ask whether the performance is produced by the exercise of an ability (or virtue) on the part of the performing agent, i.e. whether it is competent. And finally, we can ask whether the performance is successful because competent, i.e. whether it is apt. Crucially for present purposes, each of these dimensions specifies a way in which the agent may be performing *well*.

An example may help to illustrate these ideas. Consider a free throw in basketball. A free throw in basketball is a performance with an aim. The aim is for the ball to go into the basket. Suppose you take a free throw. By the theory of the normativity of performances, there are three normative properties in terms of which we may assess your free throw. We may ask whether your free throw was successful. Given that its aim is for the ball to go into the basket, your free throw is successful if and only if it gets the ball into the basket. We may also ask whether your free throw was competent, i.e. whether it was produced by the exercise of an ability on your part to make free throws. Finally, we may ask whether your free throw was apt, i.e. whether it was successful because competent. This will be the case if and only if your free throw got the ball into the basket because of the exercise of an ability on your part to make free throws.

It is important to note that success, competence and aptness are different normative kinds. True, there are relations between them. Most importantly, aptness by definition entails success and competence. Recall that an apt performance is a performance that it successful because competent. Given that this is so, any apt performance must also be successful and

competent. For instance, an apt free throw is one that finds the basket because of the exercise of an ability to make free throws. It is easy to see that an apt free throw must also find the basket and that it must be produced by an exercise of an ability to make free throws. At the same time, the three also come apart in important ways. For instance, it is possible for a performance to be successful but not competent and vice versa. For instance, you may be drunk beyond comprehension, completely unaware that you are to take a free throw and just randomly throw the ball that you are being handed. As it so happens, the ball finds the basket. Your free throw was successful but not competent. And, conversely, it may be that you are taking a perfectly competent shot that would have gone right in, had it not been for the interference of a disruptive fan who threw an apple at your ball, thus bringing it off course. Finally, a performance can be competent and successful but not apt. For instance, you may take a competent shot that would have gone right in had it not been deflected by the interference of a gust of wind, which brings it off course. At the same time, unbeknownst to you, a helper with a wind machine brings your shot back on target and it does go in after all. In this case, your shot is both competent and successful but not apt.

One key idea in reliabilist virtue theoretic epistemology is that beliefs are performances with an aim. In particular, they have a distinctively epistemic aim. Given that this is so, the theory of the normativity of performances applies to the case of belief. That is to say, we can ask whether a belief is successful, i.e. whether it attains its epistemic aim. We can ask whether it is competent, i.e. whether it is produced by the exercise of an ability to attain the epistemic aim (henceforth also *epistemic ability* or *virtue*). And we can ask whether it is apt, i.e. whether it attains its epistemic aim because of the exercise of an epistemic ability. In this way, we get three normative properties of belief. And since these three normative properties specify three ways in which one may be performing well, the result that we get here is that these three normative properties specify three ways in which one may be believing well epistemically.

Finally, another key idea is to identify normative properties of beliefs as performances with an epistemic aim with familiar epistemic properties. According to the standard view, the epistemic aim of belief is truth. As a result, the standard view supports an identification of successful belief with true belief. It also offers more information on what epistemic abilities are, i.e. abilities to form true beliefs. Thereby, it also tells us more about what competent and apt belief amount to. Finally, and crucially, the standard view identifies competent belief with justified belief and apt belief with

knowledge. The result that we get, then, is that familiar epistemic properties are simply instances of more general normative properties of performances with an aim. Moreover, it comes to light that true belief, justified belief and knowledge also correspond to ways in which one may believe well epistemically. And, of course, this means that the relevant forms of virtue reliabilism are brought straight under the umbrella of what we call virtue theoretic epistemology.

What about virtue responsibilism? How does this important strand of virtue theoretic epistemology relate to the present idea that to believe or do something well epistemically is to believe or act from epistemic virtue? There are a number of ways in which virtue responsibilism can be squared with this idea. This is unsurprising, given that virtue responsibilism is a much more multifarious research project than virtue reliabilism. For instance, one important question that virtue responsibilists are divided on is whether they should even aim to offer accounts of familiar epistemic properties such as knowledge and/or justified belief. There are those who think that the answer is yes. For them, it will be very natural to adopt the idea that believing well epistemically is to be analysed in terms of epistemic virtues. After all, knowledge and justified belief continue to be paradigm cases of believing well epistemically. But even those who return a negative verdict need not therefore reject the core idea of virtue theoretic epistemology. To see why not, note that there is fairly widespread consensus among virtue responsibilists that intellectual character virtues play a key normative role in enquiry (Code 1987; Hookway 2003). In particular, virtues feature in norms of enquiry. For instance, we ought to enquire in a way that is open-minded, intellectually courageous and so on. Given that this is so, it is only plausible that to have conducted an enquiry well will be unpacked in terms of virtues. And, of course, this fits nicely with our core conception of virtue theoretic epistemology.

There is one complication here, however. Enquiry is an activity with a constitutive aim. One way to characterise this aim in a fairly neutral manner is by saying that enquiry aims at settling the question at issue (Kelp 2019). Now, while many would sign up to this neutral characterisation of the aim of enquiry, there is a lively debate among epistemologists about how to characterise the aim of enquiry in a more substantive manner. The leading candidates here are that enquiry aims at knowledge (e.g. Kelp 2014; Millar 2011; Williamson 2000) or else at true belief (e.g. Kvanvig 2003; Lynch 2005). Crucially, if enquiry has a constitutive aim, it is plausible that virtues of enquiry are unpacked along instrumentalist lines, in terms of their conduciveness to the attainment of enquiry's

constitutive aim. Note that this way of thinking has precedents outside the epistemic domain. For instance, being tall is a virtue in basketball players. The reason why this is so is arguably that it is conducive to attaining the constitutive aim of basketball, i.e. winning the game, for instance by generating more successful blocks. What's more, this instrumentalist construal of virtues of enquiry is independently attractive. After all, it is tempting to think that the reason why open-mindedness, intellectual courage and so on are virtues of enquiry is that they are conducive to attaining the aim of enquiry.

At this stage, one may wonder where the complication lies. Isn't all this good and well and perfectly acceptable for virtue responsibilists? In principle, it is. At the same time, note that this way of thinking about character virtues makes them appear very close to reliabilist virtues. And the trouble is that one of the most important divides among virtue responsibilists is between those who think that character virtues feature a reliability condition (e.g. Zagzebski 1996) and those who think that this would be a mistake (e.g. Montmarquet 1993). That is to say, there is a significant number of virtue responsibilists who think that character traits such as open-mindedness and intellectual courage can be virtues even though they don't produce beliefs with a reliable truth to falsity (knowledge to ignorance) ratio.

There is a bit of wiggle room for responsibilists here. It could be that character traits such as open-mindedness and intellectual courage aren't virtues because they produce true belief (knowledge) reliably. In fact, they might not produce beliefs at all. Rather, the reason why they are virtues is that they are productive in generating opportunities to form true beliefs or knowledge (e.g. Sosa 2015).

Even so, we imagine that many virtue responsibilists would be unhappy to sign up to this construal of responsibilist virtues. This is because they think that the already instrumentalist account of these virtues, according to which they are virtues because conducive to attaining the goal of enquiry, is mistaken. The question remains as to whether our construal of virtue theoretic epistemology can accommodate virtue responsibilists of this stripe.

Here, in outline, is one attractive attempt at pulling this off, which circles back once more to eudaimonia or well-being. Here are the key thoughts: there is such a thing as epistemic well-being and it is multiply realisable. In particular, what epistemic well-being amounts to for fallible agents like us humans who are also very far from being omniscient might be quite different from what it amounts to for infallible omniscient agents,