Mark Schroeder
University of Southern California
March 20, 2015

is knowledge normative?

Epistemology is, at least in part, a normative discipline. Epistemologists are concerned not simply with what people believe, how confident they are, or the conditions under which they believe it or are so confident. They are also – and more interestingly – concerned with what people *should* believe, how confident they *should* be, and the conditions under which they should believe it and be so confident. But epistemology is, paradigmatically, the study of *knowledge*. Is knowledge itself a normative concept? Or is it merely normatively important? That is the question that will interest me in this paper. I will attempt to make it more precise as we go along.

I Preliminaries

Philosophers sometimes distinguish concepts that concern *norms*, such as those expressed paradigmatically by words such as 'must', 'ought', 'obligation', and 'reason' in appropriate contexts, with concepts that concern *values*, such as those expressed paradigmatically by words such as 'good' and 'better'.[1] Others sometimes contrast the 'genuine' normativity of *reasons* with the pseudo-normativity of rules.[2] Neither of these is the distinction that I have in mind, when I say that epistemology is at least in part a normative discipline. When I say this, I mean to contrast 'normative' with 'descriptive' and with 'positive', rather than with 'evaluative' or to contrast reasons with rules.

This point is important. When epistemologists question what agents ought to believe, or how confident they ought to be, they are typically interested in what agents ought *rationally* to believe, or how confident they *rationally* ought to be. Or sometimes the question is even more direct: what beliefs or levels of confidence do *rational* agents have – and hence what beliefs or levels of confidence must other agents have *in order to be* rational? When we ask, in epistemology, what an agent must or ought to believe in her situation, we are using some of the words, 'must' and 'ought', that fall paradigmatically on the 'normative' side of the

[1] Compare Tappolet [2013].
[2] See especially Parfit [2011, 144-145].

'normative'/ 'evaluative' distinction. But if all that we mean in asking this question is what the agent must believe in order to be rational, then the word 'must' cannot receive the credit for making our question normative. For we could also ask what an agent must believe in order to be member of the Church of Latter-Day Saints. That is not a normative question; it is merely a question about the LDS membership criteria, and being a member of the LDS church is a social property, rather than a normative one. So if the question of what an agent must believe in order to be rational is a normative one, it must be because being rational is a normative property, not because we can use the words 'ought' or 'must' to ask this question.

So if we are interested in whether epistemology is even in part a normative discipline, it is important to be clear on whether being rational is indeed a normative property. In a famous paper, Niko Kolodny [2005] has called this into question, arguing that rationality is not normative. But Kolodny means something very specific by 'normative' in his paper. The question that he is asking is not whether rationality is better classified along with properties and relations such as *obligation*, *better*, and *virtue* or with properties and relations like *yellow*, *tall*, and *more common*, and hence whether rationality falls on the normative side of the 'normative'/ 'descriptive' distinction. Rather, he is asking whether rationality is a property that there are always reasons to have. In my broader sense of 'normative', on which 'normative' contrasts with 'descriptive' or 'positive', I know of no good reasons to doubt that rationality is a normative property. So I infer that insofar as epistemologists are concerned with what agents ought rationally to believe, they are engaged in normative inquiry in the sense in which I am interested.

Still, even once we are clear that inquiry about which beliefs or degrees of confidence are rational counts as normative inquiry, it remains true that much of epistemology is not directly concerned with the rationality of belief or confidence. Central among the topics of concern in epistemology is knowledge: who knows, what they know, in virtue of what they know, and the consequences of that knowledge. In contrast to the case of rationality, I do not believe that it is obvious or incontrovertible that knowledge is a normative property. Indeed, many analyses of knowledge seem to be best interpreted as committed to the consequence that knowledge is not itself a normative property.

For example, Timothy Williamson [2000], who denies that knowledge can be analyzed, claims that knowledge is a mental state. Now it could very well be that knowledge is both a mental state and a normative property. But Williamson's view is that knowledge is a mental state *on a par with belief*. Since on most views, belief is not a normative property, it is natural to interpret Williamson as holding that knowledge is not itself a normative property, either.

Moreover, many analyses of knowledge have been offered from which it follows that knowledge reduces to facts about the truth values of counterfactual conditionals, causal facts, or facts about tracking or

information-carrying relationships.[3] These views can maintain the thesis that knowledge is normative only if we assume that normative properties are not just naturalistically reducible, but heterogeneously so, so that the natures of various normative properties and relations may have little, if anything, at all to do with one another. This is because there is no reason to think that the naturalistic analyses of rationality, value, and obligation are exhausted by similar facts about the truth values of counterfactual conditionals, causation, and tracking or information-carrying relationships. I do not know of any way of ruling out this possibility definitively, but I do not understand it.

My lack of understanding is not based on skepticism about the naturalistic reducibility of normative properties and relations; on the contrary, I am quite sympathetic to the project of reducing the normative in non-normative terms which are likely to be broadly naturalistic. Rather, my lack of understanding derives from the heterogeneity. If knowledge reduces to facts about causation or reliability, but rationality, goodness, justice, and obligation are either irreducible or reduce to some quite different kind of properties, then I find it hard to see what the class of normative properties has in common in virtue of which they are *normative*. Without some common nature to tie them together, this looks to me like skepticism about genuinely normative properties, rather than like a way of preserving the view that knowledge is normative. So I infer that the most natural way of understanding analyses of knowledge according to which it wholly reduces to facts about counterfactuals, causation, or information-carrying relationships, is as denying that knowledge is a normative property.

As I said, I do not think that it is obvious or incontrovertible that this commitment of such views is false. That it is not incontrovertible is demonstrated by the fact that such views controvert it. That it is not obviously false is supported by the same fact, together with the fact that it is not a terribly surprising way of interpreting the commitments of such views. But a more important reason why it is not obvious that knowledge is itself normative, is that it is sometimes difficult to distinguish between concepts that are themselves normative, and those that are merely greatly normatively *important*.

To see the distinction between normative content and normative importance, it helps to compare the property of death.[4] Death, I take it, is a biological property. It is possible that the boundaries of death are vague, or that we have come to think of it differently over time, due to our interests and the availability of better medical technology, or that there are important distinctions to be made among closely related concepts of death that would never have been distinguishable before the advent of modern life-support

[3] Classical sources for such views include Armstrong [1973], Goldman [1967], [1988], and Dretske [1981], among many others.
[4] See also Schroeder [2007, 80] for discussion.

technology such as artificial respirators. But those observations, though important, do not change the fact that death, however conceived, is a biological property and not a normative one. Nevertheless, death is a normative property of immense normative importance. The fact that an action would result in one's own death is at least nearly always some significant reason not to perform that action. When moral philosophers are interested in the normative question of what people ought to do, their answers are rarely uninformed by assumptions about who would die on each course of action.

In the case of death, I think there is a good case to be made that we are concerned with a property that is not itself normative, but is clearly and generally normatively important. But other cases are I think not as obvious. Some candidates for 'thick' normative concepts, for example, may be intelligibly thought of as genuinely normative or as merely normatively important. For example, consider the property of being *kind*. On one natural view, being kind is a non-normative, purely dispositional property, which someone has just in case they have a strong predilection to strongly empathize with other people and to act on that empathy. On this view, kindness is highly normatively important, because it is a virtue to be kind, and it is good to surround oneself with kind people, but kindness is not itself normative.

On another view, kindness is a normative property. Perhaps it is the property of being *appropriately* empathetically responsive. Or perhaps, less plausibly, it is the property of being empathetically responsive *and that's being good*.[5] On views like these, kindness is itself a normative property. Or a third possibility is that kindness – the primary property that we attribute when we call someone 'kind' – is itself not a normative property, but that the word 'kind' is a loaded or positive word for attributing this property, which gives away that the speaker sees kindness in a positive light, or presupposes that it is good to have this property. It may be that one of these views of kindness is much better than the others, but determining which is correct, or whether they are each correct of overlapping properties, each of which are typically attributed in ordinary talk about kindness, is a substantive matter – it is not obvious which is correct, in the absence of more careful reflection.

Similarly, it is not obvious whether knowledge seems to matter so much for what we may do and what it is rational to believe because it is itself normatively loaded, or whether it is so important simply because knowledge is one of the important non-normative factors that matters greatly for what we may do and what it is rational to believe. I think that it is clear, and I will assume throughout this paper, that

[5] This view is less plausible because it carries the controversial commitment that thick concepts like *kind* can be decomposed into separate descriptive and normative components. See Hurka and Elstein [2009] for exploration of how thick concepts can be analyzed in terms of thinner normative concepts without being decomposable in this way, and Väyrynen [2013] for more extended discussion of plausible models for thick language.

knowledge is *at least* normatively important. In section 2 I will describe one particular aspect of knowledge's normative importance that I think is particularly compelling and worthy of explanation. But the theorist who denies that knowledge is itself normative will say that normative importance is *all* knowledge has. Once we understand the norms governing rational belief, and see how knowledge is relevant for those norms, she claims, we will see why knowledge had the importance that it has, without need to advert to the idea that there is anything intrinsically normative about knowledge itself. To determine whether knowledge is really normative, we will need to test the limits of this idea.

2 Searching for Leverage: Knowledge and the Rationality of Belief

Knowledge has many important connections to rational belief and rational action. If any of these are 'built in' to the nature of knowledge, so that they flow from what it is for someone to know, then it follows that knowledge is itself normative. For example, according to some philosophers, knowledge can be transmitted by testimony, and so it is rational for you to believe something, if you are told so by someone who knows. If what it is to know that P is, *inter alia*, just to be in a position to make it rational for someone to believe that P by telling them so, then it is not surprising where this fact about testimony comes from – it flows directly from the nature of knowledge. If this really were part of what it is to know, we would have an explanation of this norm on testimony. And this explanation would make knowledge out to be normative – part of *what it is* to know, is for some normative condition to obtain – for you to be in a position to make it rational for others to believe that P, by telling them so. If this view of knowledge were true, I take it, then that would suffice for knowledge to be normative.

However, I don't believe that this view of knowledge is true. It may very well be that when we triangulate, as philosophers, on which property we mean to be studying when we theorize about 'knowledge', one of the earmarks that we keep in mind is that knowledge is a state that can be transmitted by testimony. But I don't think that it is particularly plausible that part of what it is to know is to be in a position to make it rational for someone else to believe that P by telling them so. This view explains why knowledge *entails* the power to make it rational for someone else to believe by testimony, but it gets the explanation the wrong way around. If you have the power to make it rational for someone else to believe that P by telling them so, that is plausibly *because* you know, rather than what *makes it the case* that you know. So this analysis, which builds the testimonial power of knowledge into its nature, gets things the wrong way around.

Much more plausible, in contrast, is the view that the testimonial power of knowledge flows from important facts about the rationality of belief in the face of testimony. On this view, the power of knowledge

to engender rational belief by testimony is a way in which knowledge is normatively *important*, but not in a way that reveals it to have a normative *nature*. If we are to look for a dimension of knowledge's normative importance that reveals it to have a normative nature, we must look somewhere else.

A more revealing fact about the normative importance of knowledge, I believe, is that knowledge entails rational belief. If you know that P, then it is rational – that is, not irrational – for you to believe that P. My claim that it is rational for you to believe that P is helpfully contrasted with the claim that you are rational *in* believing that P. Epistemologists often make a similar distinction between *propositional* and *doxastic* justification: intuitively, whether you are propositionally justified in believing that P depends only on whether it is a justified thing for you *to* believe, but whether you are doxastically justified in believing that P depends, further, on the reasons *for which*, or *on the basis of which* you believe that P.[6] What it is rational for you to believe corresponds to propositional justification, and what you are rational *in* believing corresponds to doxastic justification. I happen to believe that it is also clearly true that if you know that P then not only is it rational for you to believe that P, but you are rational *in* believing that P. But it is only the weaker claim – that it is rational *for* you to believe that P – that I will be assuming, here, and trying to explain.

On one natural interpretation, talk about justification *just is* talk about rationality. On this interpretation, my assumption that knowledge entails rational belief is just the assumption that knowledge entails justification. This interpretation is natural, because historically, justification has been assumed to be a condition that is clearly required for knowledge, but exhibited by Gettier-like cases to be insufficient, even together with belief and truth, for knowledge, and rationality seems to be clearly required for knowledge, and fits closely with standard judgments about Gettier cases. But my assumption that knowledge entails rationality of belief does not require that we identify justification with rationality.

On some views, for example, there are some propositions that we are *entitled* to believe, but which do not require justification.[7] These views do not seem to be best described as views according to which we are entitled to believe things that it is irrational for us to believe. Rather, they seem to be better described as views according to which justification requires more than rationality. On such views, knowledge may not turn out to entail justification, depending on whether they allow that we can know propositions to which we are merely entitled. But they should still allow that knowledge entails rationality of belief. This assumption is weaker, and harder to reject, than the assumption that knowledge entails propositional justification.

[6] For one discussion of this distinction, see Conee [1980].
[7] Compare especially Wright [2009].

Other, externalist, conceptions of justification are also clearly stronger than rationality. For example, Clayton Littlejohn [2013] defends a conception of justification on which a belief is not justified unless it is true. Littlejohn's way of using the word 'justification' is natural, because in moral philosophy, justifications are contrasted with excuses, and an action has a justification only if it was the right thing to do. Since it is natural to think that a belief is not right unless it is true, it follows from this conception of justification, by strict analogy with moral philosophy, that a belief is not justified unless it is true. Though Littlejohn's conception of justification makes sense, my claim is not that knowledge entails justification in Littlejohn's sense, but rather the weaker claim that it entails rationality.

On a very traditional sort of analysis of knowledge, knowledge entails rationality of belief because it is part of what it is to know that P, that you believe that P and that it is rational for you to believe that P. It follows from this view that knowledge is normative, in the sense in which I am interested in this paper, because rationality – a normative condition in the broad sense in which I am interested – figures in its analysis. In contrast to the view on which part of what it is to know that P is, *inter alia*, to be in a position to make it rational for someone to believe that P by telling them so, the class of views that endorse this thesis about knowledge are not intrinsically implausible. It is not directly implausible that you count as knowing that P in part because it is rational for you to believe that P. I take it, therefore, that the class of views on which knowledge is to be analyzed in some way that builds in rationality of belief can be thought of as *motivated* by the explanation that they give of the fact that knowledge entails rationality of belief. In contrast, views on which this is no part of the nature of knowledge must provide some alternative explanation of why this entailment obtains.

3 A Simple Norm

It is straightforward to imagine simple views that can explain why knowledge entails rationality of belief. The simplest holds that there is a single rule governing when it is rational to believe, and it says that it is rational to believe that P just in case you know that P. On this view, knowledge entails rationality of belief because knowledge makes a difference for what it is rational to believe. Given this view about the only rational norm governing belief, knowledge is not just of some normative importance, but is of *central* normative importance, because it tells us everything about what it takes for a belief to be rational. Nevertheless, the entailment flows from the rational norm governing belief, rather than from the nature of knowledge. Knowledge is just the condition that is mentioned by this norm, which is compatible with it not

being intrinsically normative in any way. This is the view endorsed by Williamson [2013], in response to Cohen and Comesaña [2013].

Although this is a simple view that explains how our entailment could be true without knowledge being intrinsically normative, however, I don't think that this is a particularly plausible direction to go. If there is a rational rule governing belief which makes it irrational to believe what you do not know, then there are no cases of rational belief that fall short of knowledge. This rules out not only the possibility of Gettier-like cases, where you believe something truly and rationally but fail to know, but also rules out the possibility even of rational but *false* belief.

The resulting picture of rationality is so distorted that I find it hard to understand how it could be a picture of the very same notion of rationality on which it is so compelling to think that knowledge entails rational belief. When I drew attention to the plausibility of the claim that if you know that P, then it is rational for you to believe that P, I suggested that rationality is something entailed by, but perhaps weaker than, justification, on the traditional conception on which justification is compatible with false belief. So even if there is *some* norm on belief which forbids believing something that you do not know, we cannot aptly use the word 'rational' in the same breath to describe this norm and to describe the entailment of rational belief by knowledge that we are setting out to explain. But it is only if we can use the word 'rational' univocally both to describe the entailment and to describe the norm, that we may take the norm, by itself, to explain the entailment.[8]

There is another problem with the proposal that the only rational norm governing belief says to not belief what you do not know. Sometimes the only reason why you do not know that P is that you have not yet inferred it from other things that you know. In these cases, you do not know that P, and you are not rational *in* believing that P (because you do not believe that P), but it *is* rational *for* you to believe that P. Recall that when I spelled out the entailment between knowledge and rationality, I focused on the fact that knowledge entails not just that you are rational *in* believing what you do, but that it is a rational thing *for* you to believe. It is an important feature of the concept of what it is rational for you to believe that it is in principle *prospective*. It applies not only to what you do believe, but to what you do not yet believe, but may. A norm that restricts belief to knowledge cannot make sense of prospective rationality of belief.[9]

[8] See Cohen and Comesaña [2013b] for further discussion.

[9] Nor does it help to weaken the norm to allow only beliefs in what is *entailed* by one's knowledge. This view faces a dilemma between disallowing non-deductive inference and having nothing to say about prospective rationality. Cohen and Comesaña [2013b, 409-411] have a particularly good discussion of this point.

Finally, though least importantly, I find there to be something peculiar about Williamson, in particular, accepting a norm like this one on belief. For on Williamson's view, there are two different mental states, knowledge and non-knowledge belief. Though both are cases of belief, that is not because knowledge is belief plus a further condition. So to believe that P is just to be in one or the other of these two states.[10] But then to say that the only rational norm governing belief is to believe only if you know, is just to say that there are no rational constraints at all on one of these states, but that the other is always irrational.

I conclude, therefore, that though this norm is the right *kind* of thing to explain how knowledge could entail rationality of belief without assuming that knowledge is itself normative, it is not the right way of doing so. If we are to have a view that competes favorably with the explanation offered by the traditional view that part of what it is to know is to believe rationally, then we must do better.

4 The Difference-Making Argument

On the traditional view, you only *count* as knowing that P *because* it is rational for you to believe that P. So the entailment follows from the analysis of knowledge, but what makes this true, is that the order of explanation goes from rational belief (*inter alia*) to knowledge. On the view that knowledge is merely normative important, however, the order of explanation goes in the other direction. It is because you know that it is rational for you to believe. Presumably, if knowledge is to explain why it is rational for you to believe that P, then knowledge of P must sometimes be able to make a difference in whether it is rational for you to believe that P. But I am skeptical about how knowledge that P could ever be the difference-maker in making it rational for you to believe that P. In this section I will explain why.

If knowledge that P is ever the difference-making in whether it is rational for an agent to believe that P, then there must be cases where it makes this difference. In such a case, there must be an agent A who knows that P, and for whom it is rational to believe that P, but it would not have been rational for A to believe that P, but for her knowledge. In other words, if she did not know, then it would not have been rational for her to believe that P. The difference-making argument challenges whether there are any cases with this feature.

In general, knowledge is equivalent to the conjunction of four conditions. An agent A knows that P just in case she believes that P, it is true that P, it is rational for her to believe that P, and it is not the case that she believes P truly and rationally but does not know. To say this is not to offer an analysis of knowledge; it is just to assert that knowledge entails belief, is factive, and entails rationality of belief. Nevertheless, it is

[10] Compare Williamson [2000, chapter 1].

a helpful observation, because it gives us some more helpful things to think about, in assessing the counterfactual that if A did not know that P, it would not have been rational for her to believe that P. For convenience, I will stipulatively use 'Gettiered' as a name for the condition of believing truly and rationally without knowing, and correlatively, 'unGettiered' as a name for its contradictory. I mean my use of this term to be evocative, but not to build in the assumption that justification may be identified with rationality.

Since knowledge is necessarily equivalent to believing truly and rationally and being unGettiered, we know that the closest possibility in which an agent A does not know that P is one where she fails to satisfy one of these four conditions – *and* we know that there is no closer possibility where she fails to satisfy any of the four conditions. So there is at least one condition out of these four, such that the closest world in which A does not know that P is also the closest world in which she does not satisfy that condition. So the counterfactual that if A did not know that P, it would not have been rational for her to believe that P will be true only if at least one of the four component counterfactuals is true: that if A did not satisfy condition C, then it would not have been rational for her to believe that P.

We may safely rule out being unGettiered. The closest world in which A is not unGettiered is one in which she is Gettiered. But by definition, being Gettiered is just believing truly and rationally. So it is trivially false that had A not been unGettiered, it would not have been rational for her to believe that P. Hence we may conclude that at least one of the other three counterfactuals must be true.

What about truth? Is it plausible that the closest world in which it is not true that P is one in which it is not rational for A to believe that P? Some choice of propositions P make this counterfactual plausible. For example, if P is the proposition that you exist, then the closest world in which you do not exist is indeed one in which it is not rational for you to believe that you exist. Nevertheless, if the only cases in which knowledge makes a difference for the rationality of belief are cases involving special choices of propositions like this one, I take it that that would not be enough to explain how knowing *any* proposition entails that it is rational for you to believe it. So we may conclude that at least in the general case, at least one of the other two counterfactuals must be true.

Next let us try the belief condition. Is it plausible that the closest world in which A does not believe that P is one in which it is not rational for her to believe that P? Again, some choices of propositions P make this counterfactual plausible. For example, if P is the proposition that you believe that P, then it is plausible that the closest world in which you do not believe that P is one in which it is not rational for you to believe it. Again, however, special cases like this are not enough to underwrite the idea that knowledge can make an important enough difference to explain how knowing any proposition entails that it is rational for you to

believe it. In general, however, believing something does not make it rational for you to believe it. So we may conclude that at least in the general case, it is the fourth counterfactual that must be true.

The fourth counterfactual says that had it not been rational for A to believe that P, it would not have been rational for A to believe that P. This counterfactual is clearly true, and trivially so. But the truth of this counterfactual cannot illustrate how knowledge makes a difference for the rationality of belief. This is because it being rational for you to believe that P cannot be what *makes* it rational for you to believe that P – something else must make this so. A trivial counterfactual like this one is the wrong sort of counterfactual to illustrate an explanatory relationship.

As I have been arguing, if knowledge ever makes a difference for the rationality of belief, there must be cases in which had an agent A not known that P, it would not have been rational for her to believe that P. Separating knowledge into four conditions, though it is not an analysis, helps us to see that except in very special cases where the lack of truth or belief could explain the lack of rationality of belief, the only cases in which it will be true that if A did not know, it would not have been rational for her to believe will be cases in which had she not known, the only condition on knowledge that she would have lacked would have been that it was not rational for her to believe. But the truth of these counterfactuals do not support the idea that it is knowledge that makes the difference for what it is rational for the agent to believe. On the contrary, the fact that they are the only cases in which such counterfactuals are true supports the contrary view that being rational to believe is part of what makes a proposition knowledge.

5 A Different Approach: Indiscriminability

So far in this paper, I have tried to articulate my question about whether knowledge is a genuinely normative property, like rationality, or is only of great normative importance. I have taken it to be a sufficient condition for knowledge to be genuinely normative, that facts about rationality are part of the nature of knowledge. This is what is claimed by the traditional approach to the analysis of knowledge, on which what it is to know that P is to be in some state which involves, *inter alia*, that it is rational for one to believe that P. This view has, moreover, the virtue of offering an easy and elegant explanation of why it is that in general, knowing that P entails that it is rational for one to believe that P. The proponent of the view that knowledge is merely normatively important must offer an alternative explanation of this entailment. The simplest and most natural is to appeal to the idea that the only rational norm on belief is to not believe what you do not know. But we saw that it is deeply implausible that this is the only rational norm on belief. More troublingly, I used the difference-making argument to argue that knowing that P can rarely, if ever, be what makes the

difference in making it rational for you to believe that P. This casts very general doubt on the idea that the explanatory direction can go *from* knowledge *to* rationality of belief.

Still, the defender of the idea that knowledge is a non-normative property that merely happens to be of great normative importance is not down and out. In section 3, we considered how she might appeal to the idea that the only rational norm on belief forbids believing what you do not know. This idea has two parts: that this *is* a rational norm on belief, and that it is the *only* one. Importantly, the assumption that it *is* a rational norm on belief is not necessary, in order explain why knowledge entails rationality. All that is required, is the assumption that there is no *other*, more *restrictive* rational norm on belief. But my objections in section 3 were objections to the thesis that there is any such norm on rational belief.

So the defender of the view that knowledge is merely of normative importance may deflect these objections by accepting only a *weaker* rational norm on belief. Indeed, this is precisely what my objections push her to do. A proper rational norm on belief should not forbid false belief, and it should not make it impossible to be Gettiered. So it must be less restrictive than the norm that forbids believing what you do not know. But if the only rational norms on belief are *less* restrictive than this norm, then knowledge will entail rational belief, after all.

One way in which this could be true is if there are no rational norms on belief at all – and so every belief whatsoever is rationally permissible. This view validates the entailment from knowledge to rational belief, but it does not preserve the idea that knowledge has any normative importance. On this view, knowledge is not genuinely normative, but it is not clearly normatively important, either, or at least, not in this way. Another view that shares this feature is the view that the only rational norm on belief is to not belief what is false. Since knowledge entails truth, this norm is strictly less restrictive than the norm that forbids believing what you do not know. So this view also explains why knowledge entails rational belief, but it fails to make good on the idea that knowledge is normatively important. What is normatively important, on this view, is just truth.

What gave the simple but implausible rational norm its claim not only to explain why knowledge entails rational belief, therefore, but to explain why this makes knowledge particularly normatively important, is that it held that it is knowledge itself that makes the difference in order to make a belief rational – and doesn't just entail that some other condition is satisfied, which itself makes the difference. But this is what the difference-making argument suggests is impossible, except at best in a very restricted range of cases.

Still, I think that there are better views of the rational norm governing belief to suit the purposes of the defender that knowledge is merely normatively important. Here is one: perhaps the only rational norm

governing belief is to believe only what you are not in a position to know that you do not know.[11] Call this the *indiscriminability* norm, since it differs from the simple but implausible norm by permitting belief in just those cases that you cannot discriminate from those that are permitted under the implausible norm. Since knowledge is factive, if you know that P, you are not in a position to know that you do not know that P. So if you know that P, you are permitted by this norm to believe that P. So if the indiscriminability norm is the only rational norm governing belief, it follows that knowledge entails the rationality of belief.

But the indiscriminability norm is more permissive. There are some propositions that are false, and hence which you do not know, but which you are not in a position to know that you do not know – for all you know, you know them. So this norm rationally permits belief in some false propositions. If it is the only rational norm governing belief, it is sometimes rational to believe false propositions. Similar remarks go for Gettier cases, in which you do not know, but are not in a position to know that you do not know. This norm rationally permits belief in such cases, and so it does not rule out Gettier cases by fiat. Hence it is more permissive in both of the ways that we earlier noted that the simple but implausible norm was too restrictive. This makes it a much more promising account of the conditions of rational belief.

Like the norm that forbids false belief, the indiscriminability norm explains the rationality of believing what you know by forbidding only beliefs that fail to meet a condition that is entailed by knowledge. Like that norm, therefore, it does not predict that knowledge is often or ever a difference-maker for the rationality of belief; all of the difference-making may be done by the weaker condition. And if we take Timothy Williamson's [2000] luminosity argument to show that there are always borderline case of non-knowledge where the agent is not in a position to know that she does not know, then we should infer that if the indiscriminability norm is the proper norm on the rationality of belief, then knowledge is *never* the difference-maker for the rationality of belief.

But unlike the norm that forbids false belief, the indiscriminability norm has an intelligible claim to explain why knowledge is of genuine normative importance – despite never making a direct difference for the rationality of that very belief. There are several ways of seeing this observation. The first is that the

[11] Compare Williamson [2013, 95], who notes that in the idealized models that he is using to investigate Gettier cases, the beliefs of his idealized agents are identified with what they do not know that they do not know. According to Williamson, the agents in these models believe whatever they are *excused* in believing. Williamson seems not to think that being *excused* in believing something is the same as being *rational* to believe it, but call the view that the only rational norm governing belief says to believe only what you do not know that you do not know, the *undiscriminated* norm. The undiscriminated norm is very similar to the indiscriminability norm as I've characterized it in the main text, and the two may correspond in some or all of the idealized models that Williamson is discussing. The indiscriminability norm as I've characterized it is more restrictive, however, because it is possible to fail to know something that you are in a position to know, simply because you do not believe it. Hence, if you are in a position to know that you do not know that P, but do not believe that you do not know that P, the undiscriminated norm permits the belief that P, but it is forbidden by the indiscriminability norm. This comparison leads me to find the indiscriminability norm more plausible.

indiscriminability norm is itself formulated in terms of knowledge. If it is the proper rational norm governing belief, then it is not whether you know that P that is of ultimate normative importance for whether it is rational for you to believe that P, but it is still a fact about what you are in a position to *know*, about what you do not *know*.

Another reason to think that the indiscriminability norm makes good on the normative importance of knowledge is that it can plausibly be seen as a kind of reflection of a more fundamental norm that states not a standard of rationality, but rather an ideal. If the ideal of belief is knowledge, then a permissive take on how to attain that ideal might permit us to believe whatever we are not in a position to know will not lead us to that ideal. A defender of the indiscriminability norm might therefore claim that it is a reflection of a deeper picture on which knowledge is of normative importance because it is the *ideal* for belief, rather than because it is the *rational standard* for belief. This view fits well, I think, with Williamson's [2000] idea that belief *aims at* knowledge, though it may be too permissive to fit with his further views about the relationship between knowledge and evidence.[12]

Finally, a suggestion by Errol Lord [2010] offers yet a third intriguing way of thinking about why the indiscriminability norm helps to articulate a view on which knowledge itself, and not just some weaker condition, has normative importance. Lord defends the view that the relation that you must bear to some reason for belief or for action, in order for it to be one of the reasons that you *have* in the sense required for those reasons to matter for what it is rational for you to do or to believe, is that the reason be something that you are in a position to know. If your reasons are the considerations that you are in a position to know, then the indiscriminability norm says that the only reason not to believe something is the fact that you don't know it – which is always a conclusive reason. If the fact that you do not know something is always a conclusive reason not to believe it, and nothing else could be a conclusive reason not to believe something, then knowledge is indeed of central normative importance.

6 Prospective Indiscriminability

Considerations like these lead me to suspect that the indiscriminability norm is at least very close to the ballpark of what is needed by the defender of the view that knowledge is merely of normative importance,

[12] In particular, Williamson holds that your evidence is just what you know, and that you should proportion your confidence that P to the evidential probability that P, conditional on your evidence. But if you believe everything that is permitted by the indiscriminability norm, then you believe everything that you are not in a position to know that you do not know. But some of these things will have low probability, conditional on what you know. So if you believe them but proportion your confidence to your evidence, you will believe some things in which you have very low confidence. This gives some reason to think that Williamson's notion of excusable belief comes apart from what he can plausibly say about *rational* belief, and indeed Williamson [2013] seems to endorse, as I've noted, what I've called the *simple* norm on the rationality of belief.

rather than genuinely normative, in the sense in which I am interested in this paper. Nevertheless, I have one remaining qualm. In my discussion of the simple but implausible norm on belief, I noted not only that it ruled out the possibility of Gettier cases and predicted that it is irrational to have false beliefs, but that it failed to account for the *prospective* nature of rational belief. If it is rational to believe only what you know, then it cannot be prospectively rational to have any belief. For before you believe it, you do not know it, and hence you do not satisfy the condition for it to be rational for you to believe.

The same problem applies to the indiscriminability norm, as I have formulated it here. Before you believe a proposition, you are at least typically in a position to know that you do not believe it. But if you are in a position to know that you do not believe that P, then you are in a position to know that you do not know that P – and hence belief that P is rationally forbidden to you by the indiscriminability norm. So the indiscriminability norm can allow for prospective rationality of belief, at best, only in the kinds of borderline cases of non-belief in which though you do not believe, you are not in a position to know that you do not believe. So it fails to make sense of prospective claims about the rationality of belief, such as that it is rational for you to believe what follows by simple deductive inferences from other things that you know.[13]

One possible path might be to *revise* our indiscriminability norm, so that it forbids only belief in what you are not in a position to know that you are not in a position to know. The original indiscriminability norm ran into trouble with prospective rationality because the fact that you do not believe something entails that you do not know it, and so whenever you recognize that you are in a position to perform an inference, the indiscriminability norm forbids it. In contrast, you can be in a position to know something even if you do not already believe it. So you can recognize that you are in a position to infer that P, without being in a position to know that you are not in a position to know that P. So the revised indiscriminability norm is an improvement, with respect to accounting for prospective rationality.

Note, however, that the indiscriminability norm and the revised indiscriminability norm come apart in cases that are not prospective. I take it that being in a position to know stands to knowing as having a propositional justification stands to being doxastically justified.[14] So just as if you are propositionally justified, you can be doxastically justified if you base your belief in the right way, if you are in a position to know, then you can know if you base your belief in the right way. But similarly, just as you can believe something for which you have a propositional justification but fail to be justified, because you believe it on

[13] Williamson [2013] idealizes by assuming that each agent already believes everything that they do not know that they do not know. This idealization eliminates prospective rationality.
[14] Compare Lord [manuscript].

the wrong basis, you can believe something that you are in a position to know and fail to know it, because you believe it on the wrong basis.

So suppose that that is the case – you are in a position to know that P, and you do believe that P, but on the wrong basis, so you do not know it. Since you are in a position to know that P, you are not in a position to know that you are not in a position to know that P. So the revised indiscriminability norm says that your belief is permissible, while the original indiscriminability norm forbids it.

It's a consequence of these two principles' disagreement about such cases that the revised indiscriminability norm is required to say something different about some of the reasons that I offered in section 5 for thinking that the indiscriminability norm could still make good on the claim to explain why knowledge is normatively important. Whereas the original indiscriminability norm could say that knowledge is the *ideal* for belief, and that is why beliefs that are indiscriminable from knowledge are permissible, the revised indiscriminability norm would have to say instead that believing what you are *in a position to know* is the corresponding ideal for belief. This is somewhat unsatisfactory; clearly believing something that you are in a position to know but for the wrong reasons is not exactly ideal.

7 Conclusion

I don't claim to have settled the question that I set out to answer, in this paper. The traditional idea that part of what it is to know that P is for it to be rational for you to believe that P offers a powerful and simple explanation of why knowledge non-trivially entails the rationality of belief – one which predicts that knowledge is itself a normative property. The difference-making argument and the failure of the simple but implausible norm together set tight constraints on how any view that tries to capture the idea that knowledge is merely of great normative importance can be shaped. I've argued that something like the indiscriminability norm offers at least close to the best prospects for making good on these constraints, but that it faces problems with prospective ascriptions of rationality of belief. I haven't shown that these problems cannot be solved, but they are the problems that must be solved, I believe, by a philosopher who wishes to plausibly maintain that knowledge, though not itself normative, is nevertheless of central normative importance for the question of what to believe.[15]

[15] Special thanks to Ram Neta, Shyam Nair, Baron Reed, Jennifer Lackey, and to the audience for an overlapping paper at Northwestern University in February 2015.

References

Armstrong, David [1973]. *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.

Cohen, Stewart, and Juan Comesaña [2013a]. 'Williamson on Gettier Cases and Epistemic Logic.' *Inquiry* 56(1): 15-29.

_____ [2013b]. 'Williamson on Gettier Cases in Epistemic Logic and the Knowledge Norm for Rational Belief: A Reply to a Reply to a Reply.' *Inquiry* 56(4): 400-415.

Conee, Earl [1980]. 'Propositional Justification.' *Philosophical Studies* 38(1): 65-68.

Dretske [1981]. *Knowledge and the Flow of Information*. Cambridge: MIT Press.

Goldman, Alvin [1967]. 'A Causal Theory of Knowing.' *The Journal of Philosophy* 64(12): 357-372.

_____ [1988]. *Epistemology and Cognition*. Cambridge: Harvard University Press.

Hurka, Thomas, and Daniel Elstein [2009]. 'From Thick to Thin: Two Moral Reduction Plans.' *Canadian Journal of Philosophy* 39(4): 515-535.

Kolodny, Niko [2005]. 'Why Be Rational?' *Mind* 114(3): 509-563.

Littlejohn, Clayton [2012]. *Justification and the Truth-Connection*. Oxford: Oxford University Press.

Lord, Errol [2010]. 'Having Reasons and the Factoring Account.' *Philosophical Studies* 149(3): 283-296.

_____ [manuscript]. *The Importance of Being Rational*. Book manuscript.

Parfit, Derek [2011]. *On What Matters*, volume 1. Oxford: Oxford University Press.

Schroeder, Mark [2007]. *Slaves of the Passions*. Oxford: Oxford University Press.

Tappolet, Christine [2013]. 'Evaluative vs. Deontic Concepts.' *The International Encyclopedia of Ethics*.

Väyrynen, Pekka [2013]. *The Lewd, the Rude, and the Nasty: A Study of Thick Concepts in Ethics*. Oxford: Oxford University Press.

Williamson, Timothy [2000]. *Knowledge and Its Limits*. Oxford: Oxford University Press.

_____ [2013]. 'Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic.' *Inquiry* 56(1): 77-96.

Wright, Crispin [2009]. 'The Perils of Dogmatism.' In Nuccetelli and Seay, eds., *Themes from G.E. Moore: New Essays in Epistemology*. Oxford: Oxford University Press.