

The Spectra of Epistemic Norms

Matt Weiner, University of Vermont

Abstract

I argue that there is a wide variety of epistemic norms, distributed along two different spectra. One spectrum runs from the ideal to the practical and concerns the extent to which it is possible to follow the norm given our cognitive and epistemic limitations. The other spectrum runs from thin to thick and concerns the extent to which the norm concerns facts about our beliefs over and above the content of the belief. Many putative epistemic norms, such as truth and various conceptions of justification, can be found at different points on the spectra. There is no single obvious privileged point from which to say any of these norms is more fundamental than the others, though there may be some reason to doubt that some of the norms are intrinsically interesting.

Epistemology is the study of our beliefs and whether they are good beliefs. This may be a bit tendentious; there are those who would argue that beliefs are not part of our mental furniture, and even relentless naturalists such as Quine who argue that belief is not normative at all.¹ But supposing we do take our mind to represent the world through beliefs, and supposing we do take there to be norms on those beliefs, we may ask

¹ Stich (1983) is one person who has argued that our mind is not best described in terms of belief. Toward the end of this essay I will consider the views of Williamson (2000) that we should consider knowledge as a mental state that is more fundamental than any state of belief. The locus classicus for the naturalized non-normative view of epistemology is Quine (1969).

ourselves about those norms: What makes a belief a good belief? What should we believe?

It might seem that these are effectively the same question. We should believe whatever it takes to make our beliefs good; so if we can figure out what makes a belief good, that is the end of the story about epistemic norms. The epistemic norm is whatever property makes a belief good. I will argue that this is mistaken. Even if we have one conception of the epistemic good, there will still be many norms that flow from that conception.

An answer to the question “What makes a belief a good belief?” gives us an epistemic ideal. Whatever our answer is, it tells us what our beliefs would be like if we were in the best possible epistemic situation. An answer to the question “What should we believe?” or “How should we form beliefs?” may be more practical. Answering this question might tell us what actions or processes we actually should be using to form beliefs, given our actual epistemic situation and cognitive capabilities; or it might tell us what beliefs would result from these actions or processes. Each of these answers can be seen as yielding an epistemic norm, one more idealized, one more practical. But there will also be intermediate norms. We need not make a simple binary choice between an ideal or a practical point of view; rather there is a spectrum of points of view between these extremes, and from different viewpoints different norms seem appropriate.

The ideal-practical spectrum is not the only way in which epistemic viewpoints can vary. We can also consider what I will call thinner and thicker viewpoints. A thin viewpoint is one from which we consider individual propositions in isolation as something to be believed or not. Confronted with a proposition, we ask “Should I believe

this?” or “Would belief in this proposition be a good belief?” A thicker viewpoint takes into account more than an isolated proposition. It allows for questions such as “How should I believe this proposition?”, “Would belief in this proposition make sense in light of my beliefs in other propositions?”, or “How might belief in this proposition fit in with my entire system of beliefs?” Thickness and thinness also admits of degrees: We might consider a proposition in isolation, or against the background of a narrowly circumscribed set of related propositions, or against our entire belief system; and we might consider more or less about the way that we believe a proposition, over and above the fact that we do believe the proposition.

So there are at least two spectra along which we may locate a viewpoint from which we ask epistemic questions.² We may take up a viewpoint that is more practical or more ideal, that is thicker and thinner. Any viewpoint along these two spectra may be associated with a different epistemic norm. So there will be not just a plurality of epistemic norms but potentially an entire coordinate system of them, each appropriate to a different epistemological project.

1. Advice And Practicality

The best way to understand the practicality-ideality axis is by analogy to advice about actions. A norm can be thought of in terms of an associated property *P*, such that only beliefs with property *P* satisfy the norm. Endorsing a norm associated with *P* is like

² There may also be a spectrum from what we might call static to dynamic norms, where static norms take our current epistemic situation for granted and dynamic norms consider how our epistemic situation may change or what we might do to change it. And perhaps there are other ways in which our epistemic viewpoints might vary.

advising “Hold all and only beliefs with property *P*.” But, as I have argued elsewhere (Weiner 2005), advice about actions falls along a spectrum from the ideal to the practical, so epistemic norms plausibly fall along a similar spectrum. Advice about actions can be rooted in the practicalities of the advisee’s capabilities or by idealized to what she would be able to accomplish in the best of all possible worlds, and neither kind of advice is necessarily superior to the other. Epistemic norms can similarly be rooted in practicalities or aimed at an ideal.

Consider an advisee with a fixed overarching goal but limits on her ability to achieve it. Suppose that Alice has the overarching goal of scoring well in bowling; any advice we give her is directed to that end and only that end. In one way, “Knock down all the pins” would be excellent advice. If she manages to follow this advice, she will certainly score well. Nevertheless, the advice does not seem very helpful, for we have not told Alice how to knock all the pins down, and in fact it may not always be within her power to do so. More helpful advice may be “Hit the head pin on the right.” If Alice always can hit the head pin on the right she may not score as high as if she always knocks down all the pins, since sometimes she will succeed in hitting the pins in this way without knocking them all down. But the advice is more helpful because it will be easier for Alice to do what we have advised her to do.

Still, even this may not always be within her power to do; she may bowl for that spot and fail. If we want to give Alice advice that she can follow without fail, perhaps we should say “Try to hit the head pin on the right.” Alice may not be able to hit the head pin on the right at will, but she can always try to. Yet succeeding in following this advice is still less likely to result in a high bowling score than succeeding in following the advice

to hit the head pin on the right, for when Alice does try and fail to hit the head pin on the right she will likely knock down fewer pins.

Here we have a spectrum of different advice, all in the service of the same goal of a high bowling score. The advice varies in how much successfully following the advice conduces to Alice's goal, and in how much control Alice has over whether she does successfully follow the advice. Conduciveness and control yield one spectrum, not two, because good advice must balance any loss of control by a gain in conduciveness. Consider the advice "Roll the ball in the gutter"; this may be something that Alice can accomplish every time, but accomplishing it will not conduce to a high bowling score even to the extent that accomplishing "Try to hit the head pin to the right" will. If we give Alice advice that she can successfully follow with a certain rate of success, we should give her the advice that best conduces to a high score among all advice that she can successfully follow with that rate. So less control requires more conduciveness and vice versa.

Note further that no particular place on the control-conduciveness spectrum need be privileged. "Knock down all the pins" seems like obviously unhelpful advice to an ordinary bowler, but it might be just the thing to say to a perfect athlete who does not know the rules of bowling. The advice that is most conducive to the goal expresses the ideal goal of the advisee's action. It idealizes away her limitations. Advice that can be followed readily but may not achieve the goal when successfully followed is practical with respect to the advisee's capabilities. But neither the idealized end nor the practical end of the spectrum is necessarily superior to each other, or to any place in between.

Epistemic norms are subject to similar considerations of practicality and ideality. Let us begin by considering thin norms, which apply to individual propositions in isolation and consider only whether believing that proposition would be epistemically good for a subject (without considering anything else about the belief). Suppose that there is some property G such that believing in propositions with G is an unconditional epistemic good. Clearly it is some sort of epistemic norm that we should believe propositions with G . If every proposition you believe is G and you believe every G proposition, then you have attained this epistemic good to its fullest extent, and (since it is a good) in some sense you should do this. Yet it might be the case that most non-ideal subjects could not unfailingly ensure that they believed all and only the G propositions. Either their cognitive limitations prevent them from evaluating whether propositions are G ; or, leaving their limitations aside, they do not have evidence that puts them in a position to distinguish every G proposition from a non- G proposition. In that case, the norm “Believe the G propositions” is like the advice to knock down all the pins. It is ideal rather than practical; it expresses what would be best if we could attain it, while abstracting away from the limitations that may prevent us from unfailingly attaining it.

2. Truth And Justification

I mentioned two kinds of limitation, evidential and cognitive. Let us begin by focusing on the evidential limitations. Any given believer will have certain information by which she can judge how the world is. Remaining as noncommittal as possible about exactly what the nature of that information is, let us call the state of having that information her *evidential state*. For instance, if one’s evidence consisted entirely in

having had certain sense-experiences, one's evidential state would be the state of having had those sense-experiences.

In general, a norm of belief can be expressed in terms of a relation between believers and beliefs; if relation R holds between subject S and belief b , then S ought to hold belief b according to the norm that corresponds to R . We can classify these relations in terms of the facts that determine whether R holds of b for a given S . If whether S bears R_1 to b is determined entirely by S 's evidential state, and whether S bears R_2 to b depends in part on facts that go beyond S 's evidential state, then P_2 is more idealized than P_1 ; it abstracts away from some of the limits of S 's evidence.³ Intuitively, the more facts that R may depend on other than those determined by S 's evidential state, the more idealized the norm corresponding to R is. There may be no rigorous way to define "more facts" in this sense, but it should be intuitively clear at least in some cases; in particular, the most idealized relations R will be those that can depend on any fact that there is.

(Now a bit of terminological simplification. For any relation R between subjects and beliefs, holding the subject S fixed yields a property of beliefs: being a belief b such that S bears R to b . Since it is generally obvious who the subject of beliefs is, I will often speak of norms as corresponding to a property of beliefs rather than to a relation between subjects and beliefs. I will also tend to identify the norm "Hold all and only beliefs that have property P " with the property P itself, so that I can speak of a property occupying a place on the spectrum of epistemic norms. And, as mentioned above, for the thinnest

³ More explicitly, suppose that for any two subjects S_1 and S_2 who have the same evidential state, S_1 bears R_1 to a given b if and only if S_2 bears R_1 to b . And suppose that R_2 lacks this property; there is some b such that S_1 bears R_2 to b and S_2 does not, even though S_1 and S_2 have the same evidential state. Then the norm corresponding to R_2 is more idealized than the norm corresponding to R_1 .

norms the only fact about a belief that matters is the believed proposition, so these norms may be seen as a property of propositions rather than beliefs, if we are holding the subject fixed.)

The considerations about which facts a property of beliefs depends on tell us where it falls along the practical-ideal axis. In the analogy to advice, this corresponds to our chance of successfully following a certain piece of advice. As good advice must be the most conducive to the goal given how easy it is to follow, an epistemic norm must be the property that best conduces to the epistemic good, given its place on the practicality-ideality axis; which is to say, given the facts that it depends on. Once we have decided the limitations on the facts we may take into account when deciding whether our subject ought to believe something, the norm that we place on beliefs ought to be whatever property produces the most epistemic good within those limitations.⁴

Consider thin norms again, represented as properties of the propositions that are to be believed. What thin epistemic norms will we find at various points along the practical-ideal axis? The norm at the ideal extreme corresponds to a property that may depend on any facts whatsoever. This should be the property of being a proposition that we would want to believe irrespective of the limits on our epistemic position; what we should believe with all the facts at our disposal and unlimited mental power to process them. Clearly these are all and only the true propositions. If we could help ourselves to all the

⁴ In (Weiner 2005) I argue at greater length for a similar conception of epistemic norms.

facts, then we could help ourselves to the facts about whether any given proposition is true, and that is enough to determine that we should believe it.⁵

Hence truth is the epistemic norm that we find at the thinnest and most idealized point on the epistemic spectra. This idealized viewpoint lets us avoid the standard objection to truth as a norm of belief, that true beliefs may result from lucky guesses or otherwise be unjustified by our evidence. Since we are considering the most idealized epistemic norm, the limits imposed by our evidence are completely irrelevant, and the objection gains no traction.

Not too far away from the extreme ideal end, we find familiar internalist and externalist notions of justification. Take a reliabilist notion of justification on which a belief is justified if and only if it is attained through an actually reliable process.⁶ Assuming that this property is well-defined, it depends on the facts concerning how certain beliefs were arrived at and the facts that determine the reliability of those processes. This makes it a somewhat thick norm, as it depends not only on the proposition that is believed but on the process that the believer used to arrive at that belief. A completely thin concept of externalist justification would have to abstract from the actual way that beliefs are arrived at, perhaps as follows: *S* is thinly externally justified in believing *p* if and only if there is some reliable method *M* by which *S* could come to believe *p*. Supposing that this property is well-defined, it is less idealized than

⁵ It might be thought that we could say the same thing about knowledge: If we could help ourselves to all the facts, then we would want to believe all the propositions that we knew to be true. But this analysis does not motivate any special role for the propositions that we actually do know. It does motivate a special role for the propositions that we would know if we had all the facts, but those are all and only the actually true propositions. (Thanks to John Turri for pressing me on this point.)

⁶ See, e.g., Goldman (1979).

truth as a norm, because it does not take into account every fact that there is; only facts concerning the availability of belief-forming methods, the reliability of those methods, and what beliefs those methods would yield.⁷

Compare now an internalist notion of justification: Belief in proposition p is internalistically justified if p is likely to be true given the believer's evidential position. (Here we must set some threshold for how likely it must be, as the reliabilist must set some threshold for how reliable the methods involved must be.) Reliabilist justification depended only on a limited set of facts beyond the subject's evidential situation; internalist justification so defined depends on no facts beyond the subject's evidential situation. So internalist justification is more practical than reliabilist justification, which is in turn more practical than truth. Each should be the most truth-conducive property given the facts that it is allowed to depend on. Truth is of course the most truth-conducive property of all; if one is allowed to consider only facts about the reliability of available methods and their output, the most truth-conducive way to form one's beliefs is to believe the outputs of the reliable methods; and if one is only allowed to consider one's evidential state, the most truth-conducive beliefs will be those that that state makes likely.

⁷ The original thicker notion of externalist justification might fail to be well-defined because of the generality problem: whether the process through which a belief was attained is reliable can depend on how that process is described (Feldman 1985). The thinner notion may share this problem, and may also encounter the problem of defining which methods are available. I will not worry about whether these problems are solvable, as my concern is not with the workability of any particular flavor of reliabilism but with the place of externalist notions of justification on the practical-ideal spectrum. Even if externalist justification as described in the text is not well-defined, there should be well-defined externalist notions of justification that occupy similar places on the epistemic spectra.

So all these properties are norms that occupy different points on the practical-evidential axis.⁸

3. Working Within Our Limitations

Even though internalist justification does not idealize the believer's epistemic situation, it ignores any limits there may be on her processing power. Ordinary human beings will not in general be capable of ensuring that we believe all and only what is justified by the totality of our evidence; we lack the computational power to reevaluate all our beliefs anew as new evidence comes in. So internalist justification is not the most practical thin epistemic norm. The injunction "Believe only what is internalistically justified" would be one that we could not always follow, even though the norm it embodies depends only on our evidential situation. More practical epistemic norms will be those that take some account of our cognitive limitations; the more they take them into account, the more practical they will be.

Some way down the practical end of the spectrum, we might be interested in a belief revision theory of the sort put forth by Harman (1986). Harman points out (p. 27) that we do not have the mental resources to constantly update our degrees of belief in response to new evidence. Though his immediate target is Bayesian epistemology, the point applies to any theory of internalist justification, in which whether a belief is justified can depend on any aspect of our evidential situation. If we are to conform to the epistemic norm that

⁸ I present a similar argument in more detail in (Weiner 2005), in particular arguing that internalist justification is a viable epistemic norm even though it is less truth-conducive than externalist justification.

we believe whatever is justified by our evidential situation, we might have to update all our beliefs whenever our evidential situation changed.

In contrast, Harman's belief revision principles are based on the idea that we cannot clutter our minds with infinitely many beliefs and that we cannot constantly reevaluate our beliefs without some specific reason to think they are mistaken, among others. Hence his norms of belief differ greatly from those of internalist justification. To take one example, the Principle of Conservatism is that "One is justified in continuing to fully accept something in the absence of a special reason not to" (p. 46). If we idealize away our cognitive limitations, there is no need for the Principle of Conservatism, since we could always evaluate our beliefs anew with every new piece of evidence. From a less idealized perspective that takes into account some aspects of what we can actually do given our cognitive limitations, the Principle of Conservatism can be seen as part of an epistemic norm.⁹ The injunction "Continue to hold your beliefs unless there is a specific reason not to" would be easier to conform to in practice than "Believe whatever is justified by your total evidential situation," even if both operate within our evidential limits; so the first injunction is a more practical epistemic norm. Note that this, like standard reliabilist justification, is a somewhat thick norm; it depends not only on the proposition believed but also on how that belief was arrived at.

There are more practical norms, however. Harman's theory takes into account some of our cognitive limitations, but (as Harman acknowledges, p. 9) it still idealizes our cognitive powers somewhat. For instance, to follow the Rule of Conservatism we must be able to recognize when we do have a reason to revisit our beliefs. Sometimes some of us

⁹ The epistemic norm, stated in full, would be to follow the entire set of Harman's belief revision principles.

will not be able to recognize this. In contrast Bishop and Trout (2005) aim to offer epistemic principles that respect all our cognitive limits, making epistemology “a discipline that offers practical, real-world recommendations for living” (p. 6). Bishop and Trout argue that we should seek out and employ Statistical Prediction Rules, evaluating certain questions in terms of simple linear combinations of a limited number of measurable variables; for instance, to diagnose a patient as psychotic if a combination of five factors on a personality test adds up to 45 or more (see p. 14, citing Goldberg 1965). These prediction rules, they argue, are more likely to yield correct judgments than attempts by experts to evaluate the totality of available evidence. Hence a simple norm that we should believe in accordance with prediction rules would be, on their telling, the most truth-conducive of all injunctions that we could invariably carry out given our actual cognitive limitations as well as our evidential situation. Like Harman’s, this is a somewhat thick norm that depends partly on how beliefs are formed.

Note that Trout and Bishop’s prescriptions depend specifically on a refusal to idealize our cognitive capacities. Besides their explicit citation of our cognitive limits (p. 39), they suggest that some prediction rules apply only to problems that people tend to get wrong (p. 33). They also, like Harman, concern themselves with how we may efficiently allocate our cognitive resources; Bayesian updating and deductive closure, they argue, may be more reliable than our normal reasoning practices but also more expensive in terms of cognitive resources (p. 68). This expense would not matter if our cognitive resources were not limited. Indeed, Bishop and Trout acknowledge that the costs involved in learning and carrying out their some of the strategies they discuss may be high enough that some people might reason better by sticking to their old reasoning

patterns.¹⁰ As we would expect from a norm that is farther toward the practical end of the axis, they are trading off some truth-conduciveness to consider properties that people can actually evaluate.

Even a norm that we can invariably respect to given our cognitive limitations and epistemic situation may not be the most practical of epistemic norms. That we can respect a norm does not mean that we will respect it. To return to the analogy with advice, if the most goal-conducive advice that is within the advisee's power is not advice that the advisee would follow, it may be more practical to give advice that would be followed, even if following it is somewhat less conducive to the goal. Analogously, even if following strategic prediction rules is the most truth-conducive thing that people can actually do given their cognitive and evidential limits, people seem to have considerable resistance to adopting them and other belief-forming strategies that are suggested by results in social psychology;¹¹ and Bishop and Trout worry about how ineffective their counsel will be if it is restricted to specialized journals (p. 102). If epistemic norms that it is easy to get people to conform to count as more practical than norms that merely respect our cognitive limitations, the most practical epistemic norm might be one that was likely to be disseminated and followed. Exactly what that norm might be I do not know; I am not aware of any philosophers who have a theory of the most effective epistemic advice that people will actually listen to.

¹⁰ See in particular their discussion of frequency format probability analysis and consider-the-opposite strategies (Bishop and Trout 2005, pp. 68-9).

¹¹ For instance, Bishop and Trout mention (p. 150) that some philosophers insist on holding job interviews even though they are aware of experiments that suggest that they do not improve the quality of hiring decisions.

4. Thick Norms

The norms I have considered up till now have been more or less thin; if they have concerned any facts about a belief other than its content, it has been relatively narrow facts about how the belief was arrived at. We can also consider thicker epistemic norms, which may concern how a belief relates to other beliefs, the extent of these other beliefs depending on the thickness of the norm; or what sort of dispositions and virtues a belief reflects; or any other aspect of a belief that might contribute some sort of epistemic good. In the extreme case, where the most idealized thinnest norm expressed which individual propositions we would want to believe in the ideal epistemic situation, the most idealized thickest norm will express what overall total epistemic state we would want to have in the ideal epistemic state.

For this thick norm, believing all and only true propositions will not be enough. Once we consider aspects of a belief beyond the proposition that is believed, truth alone will not be the sole epistemic good of belief. For instance, it is epistemically good to believe a proposition on proper grounds. Beyond that, we will want to understand why the proposition is true.¹² Indeed, the most idealized and thickest norm will be the highest epistemic state of all: understanding everything.

The thickest norms of all may depend on any aspect of a belief, and the most idealized norms may depend on any facts. So for a belief to satisfy the most idealized thickest norm, it must be a belief that is ideal in every way, irrespective of any of our epistemic or cognitive limitations. Mere true belief will not be enough here; we will want the belief to have come about in the right way and to have the right sort of relations to our

¹² Zagzebski (1996) and Kvanvig (2003) are two authors who have emphasized the importance of understanding as an epistemic norm.

other beliefs. Part of this means knowing which other propositions make our belief true and explain it. But it is then in turn desirable to understand those propositions, and to understand the propositions that make them true and explain them, and so on until we have complete understanding of any truth there is. Since we are concerned with the thickest possible epistemic norm, there is no reason to stop short of complete understanding of everything. This may provide some motivation for a sort of unity of the epistemic good; we cannot completely understand one thing without completely understanding anything. Without any limits imposed by our epistemic situation or the generality with which we may consider our propositional attitudes, this is the epistemic ideal we will arrive at. (We may also need to have arrived at this understanding in the right way, but it seems as though complete understanding should be self-legitimizing; if you understand everything, you can confirm all your beliefs in a well-grounded way.)¹³

A less totalizing notion of understanding will be found away from the thick extreme of the thick-thin spectrum. When we consider a somewhat less thick norm, we will be evaluating a belief in a particular proposition in the context of the believer's attitudes toward certain related propositions without expanding our view to every proposition there is. If we are at the ideal end of the spectrum, we need not worry about limits imposed by our evidential situation, so we may arrive at the injunction "Believe the truth about this

¹³ John Turri suggests that having a true belief about everything will entail understanding everything, since some of the propositions one believes will be the truths about explanatory relations among one's beliefs. This may be true, in which case having complete understanding of everything will be extensionally equivalent to believing all and only the truths. Yet, since moving a little way away from the thick extreme of the epistemic spectrum yields a norm to do with understanding in a specific area rather than true belief in a specific area, I think it is more useful to think of the thickest norm as a norm of complete understanding rather than one of complete true belief.

proposition and about those propositions that explain its truth, and understand why the proposition is true in terms of related propositions.” We need not understand all the related propositions in turn, since we are not looking for the thickest possible norm, but the belief in question must be true, known, and understood in terms of its relations to further true propositions. The idealized slightly less thick norm is that of understanding with knowledge. This contrasts with Kvanvig’s argument for understanding as an epistemic norm; Kvanvig argues that understanding does not entail knowledge (2003, p. 196ff.). But this kind of understanding might perhaps be found further away from the ideal end of the epistemic spectra. Whether we know a proposition can depend on facts outside our evidential situation, so a conception of understanding that does not entail knowledge will be less idealized than one that does. If we are at a point on the practical-ideal axis that restricts us to properties that depend on our evidential situation, but we thicken our norm so that it takes into account our attitudes to related propositions, the resulting norm may be that we come to the best understanding we can of the proposition given the limits on our epistemic situation.

Another way in which a norm can be thicker rather than thinner is to involve the way that a belief is arrived at, as well as the proposition that is believed. I have already mentioned that reliabilist justification, Harman’s belief-revision, and Bishop and Trout’s statistical prediction rules all yield norms that require beliefs to have been formed in the right way. The concept of internalist justification I have discussed might similarly be thickened; a belief that *p* would satisfy this norm if the subject’s evidential situation makes it likely and if the subject believes it *because* her evidential situation makes it likely. This would add a grounding relation to justification, but might not be as thick as

full understanding, which extends not only to the way the belief is formed but also to the subject's attitudes to related propositions. (Or perhaps neither norm should be said to be thicker than the other; there may be no complete linear ordering of thickness.)

As with thin norms, more practical thick norms involve not only our evidential situation but our cognitive limits. Here we find norms that tell us what general habits of mind conduce to the epistemic good, given our limitations. For example, Montmarquet discusses virtues of impartiality, intellectual sobriety, and intellectual courage.¹⁴ These virtues are thick in that they concern how beliefs are formed and concern a variety of propositions; it seems impossible to exercise these virtues with respect to one proposition only. They are toward the practical end of the spectra because they presuppose cognitive limitations. Someone with infinite cognitive power would have no need for impartiality, for she would be able to evaluate the evidence for beliefs without regard to whether they were her own or another's. She would never need to consult another's evaluation of evidence they shared. Impartiality and intellectual sobriety are needed so that we can recognize our intellectual limitations and come to the best beliefs we can in spite of them. Even further to the practical end of the spectrum we may find general habits of inquiry rooted in the psychological literature discussed by Bishop and Trout, such as finding generally applicable statistical prediction rules, considering reasons that go against your view (Bishop and Trout 2005, p. 144ff), and considering Bayesian problems in terms of frequencies rather than probabilities in order to make it easier to avoid certain errors (p. 141ff.).

¹⁴ For an extended treatment of these and other intellectual virtues, see Zagzebski (1996).

5. Pluralism And Its Limits

So far my discussion has been sunnily pluralistic. I have aimed to show that there are many different epistemic norms, each valid from a different standpoint, and that no one of these standpoints need be better than another. Whether we ask what our ideal belief would be or whether we ask what is within our capacities to believe, whether we focus exclusively on the content of our beliefs or whether we consider other aspects, whether we consider beliefs one at a time or whether we evaluate a belief against the background of related beliefs; all these perspectives can be valid, and all yield different epistemic norms. But this does not mean that just any concept that might seem to be an epistemic norm is one.

The first way we might want to limit pluralism concerns the question of whether all points on the epistemic spectra are equal. We have been finding possible norms by choosing a point on the spectra and considering the most epistemically effective property that meets the constraints on epistemic properties set by that point. (For instance, that thin norms depend only on the proposition that is believed, and that at a certain point on the practical-ideal axis the property in question can depend only on the subject's evidential situation.) As I argued by analogy to advice, no one point on the spectra is the single right point for finding epistemic norms. There are reasons to consider the norms we get by idealizing our epistemic situation, and there are reasons to consider the norms we get by considering our evidence and powers exactly as they are. But this need not mean that every single point on the spectra yields an epistemic norm of equal interest or importance.

Some points on the spectra clearly are of interest, so that the norms they yield will be important epistemic norms. The most extreme idealized norms tell us what our ideal epistemic state would be. The most extreme practical norms prescribe give us ways to actually improve our epistemic state. But elsewhere on the spectra, we may wonder why we should be interested in the most epistemically effective property that satisfies these particular constraints. Perhaps the degree of idealization or thickness at that point on the spectra isn't particularly natural, and there is no reason to care about a property that is exactly that idealized, no more or no less.

My suspicions here are directed at the reliabilist notion of justification. As discussed above, this notion of justification is far toward the idealized end of the spectra; not only does it abstract away from our cognitive limits, it even abstracts away from some of the believer's evidential state. Whether a belief is obtained through a reliable process can depend on facts that are not part of the believer's evidential state and are in no way epistemically accessible to the believer. Yet reliabilist justification depends only on some of these facts that are beyond the believer's ken. The only such facts that are to be taken into account when determining whether a belief conforms to this norm are facts pertaining to the actual reliability of the believer's belief-forming process. Why idealize this far and no further? If our epistemic standpoint allows us to take into account facts that go beyond the believer's evidential situation, why take into account only the facts that pertain to the reliability of belief-forming processes, and no other facts? There may be answers to these questions, but they must be answered in order to show that reliabilist justification is an important epistemic norm. Otherwise reliabilist justification will be the

answer to an uninteresting question, “What would we want our beliefs to be if we had this much extra evidence and no more?”

The second prospective limit to my pluralism concerns the concept of knowledge. As the reader may have noticed, none of the points on the epistemic spectra that I have discussed yields knowledge as a norm. In particular, knowledge is not the norm that one finds at any of the most idealized points of the spectra. I have already argued that the most idealized thin norm is truth; if we consider which proposition we should believe without reference to any limits imposed by our evidential situation, we should believe all and only the true ones. Truth is the norm here, not knowledge, because knowledge incorporates restrictions from our evidential situation; in order to know something we must at least have some evidence for it. This is precisely the sort of restriction that we are abstracting away from when we consider the ideal extreme of the epistemic spectra. There is no reason to think that in the ideal situation we would want to deprive ourselves of some true beliefs by believing all and only the things we actually know.

Similarly, when we consider the most idealized thickest norm, we find a notion of understanding that goes far beyond knowledge (at least as conceived of by contemporary epistemologists, if not Plato). When we consider the ideal epistemic state as applied to a complex of beliefs rather than a single isolated belief, we will not merely want to know that *p* but to understand it as thoroughly as possible. We can know an isolated fact without understanding much about it, as when a trustworthy informant tells us one thing about a subject that is otherwise completely obscure to us.¹⁵ If our concern with an ideal

¹⁵ See Lackey (2007).

epistemic state goes beyond merely believing the right propositions, we need more than we are guaranteed when we know something.

This does not prove that there is no point on the epistemic spectra from which knowledge will be seen to be an epistemic norm. Perhaps there is some degree of thickness such that the ideal norm with that thickness is knowledge. But this would require finding some set of facts about a belief such that a belief that was ideal with respect to those facts always qualified as knowledge, without guaranteeing understanding or anything else beyond knowledge. It may be hard to find the right thickness for knowledge without resorting to special pleading; and even if we can find it, it might be subject to the objection I raised against reliabilist justification, that we need an account of why this particular point on the epistemic spectra is interesting.

In any case, I think that the best explanation for the value of knowledge is not that there is one particular standpoint from which knowledge is valuable, but that to know something is to guarantee the fulfillment of several different epistemic norms. If we know that p , p is true, we are justified in believing that p , there is a high chance that we have some understanding of p , and depending on our analysis of knowledge more norms may hold besides. It is useful to wrap these norms up in a single concept; even if these norms are all valuable from different points of view, we still may be interested in an easy way of expressing them all at once. Hence knowledge, even if it does not have its own place on the epistemic spectra, may have value as a combination of norms that do have their own place.¹⁶

¹⁶ In Weiner (2009), I argue for a similar conception of the value of knowledge as a combination of properties that on their own are intrinsically valuable from different

There is one more way to account for knowledge as an epistemic norm. I began this essay by claiming that epistemology is the study of our beliefs and whether they are good; a somewhat tendentious claim, because some would argue that our mental furniture is not best described in terms of belief. And indeed Williamson (2000) has argued that knowledge is conceptually prior to belief. Belief is best thought of, roughly, as a state that aspires to the condition of knowledge. If this is true, then it will be baked into our conception of epistemology that knowledge is an epistemic norm. Indeed, since Williamson holds that our evidence just is the propositions we know, then knowledge would be the epistemic norm that held everywhere from the ideal end of the spectrum to the point on the practical-ideal axis at which we find properties that depend only on our evidential state, since our evidential state would always distinguish propositions we know from those we don't.

But this defense of knowledge as an epistemic norm would require a radical reconception of what epistemology is. Unless the only way to make sense of what beliefs are is to conceive of them as directed at knowledge rather than truth, this conception of epistemology effectively assumes rather than proves that knowledge is an epistemic norm. And even this conception of epistemology admits of epistemic spectra. We could still consider thicker norms that require more complete understanding beyond knowledge, and more practical norms that concern what beliefs we ought to arrive at when we cannot know what we know. So, even if a "knowledge first" approach to epistemology radically alters our picture of what the epistemic norms are, it still does not entail epistemic

standpoints. There my focus is on different standpoints on practical reasoning rather than different epistemic spectra.

monism. There will be many standpoints from which we can ask different questions of our epistemology, and those questions will yield different epistemic norms.

6. Why Pluralism Works

Finally, to address the objection that the whole idea of pluralism about epistemic norms is incoherent. One might think that norms are supposed to be action-guiding, or perhaps in the epistemic case belief-guiding, so that to say that different norms apply from different standpoints is to deprive them of this role. For sometimes norms may conflict, as when a belief that would be true would not be justified. On my view a true but unjustified belief would satisfy one epistemic norm, one we are concerned with at the most idealized thin standpoint, but violate another, which we are concerned with at a less idealized thin standpoint. The norm we derive from one standpoint tells me to believe the proposition, the norm we derive from another tells me not to. How can it be the case that I ought to believe and not believe the same proposition?

My response is to deny that norms are always action-guiding. A norm can express an ideal, even if there is no action that could guarantee achievement of this ideal. It may even be that when we take up a more practical standpoint we will not attempt to achieve the ideal, so that the more practical norm is not related to the more ideal norm in any obvious way. To return to the bowling example, suppose that a bowler lacks enough skill to consistently bowl strikes, but can consistently roll the ball down the middle of the lane and knock some pins down. The most ideal advice for this bowler remains “Knock down all the pins,” but the most practical advice may be “Roll the ball down the middle of the lane” rather than “Try to knock down all the pins.” Knocking down all the pins would

lead to the highest score if it could be done, but trying to knock down all the pins may not lead to as high a score as simply rolling the ball down the middle of the lane. (Perhaps the bowler will push too hard and put the ball in the gutter.) Yet knocking down all the pins is still the ideal, and the bowler's bowling is lacking insofar as it falls short of this ideal.

In this case "Knock down all the pins" is a norm on this person's bowling, but it is not an action-guiding norm. At the idealized end of the spectrum a norm can express an ideal, even if we cannot derive the actions that anyone ought to take from ideals that they fall short of. The more practical norms will guide actions, and it is no surprise when an action that conforms to these action-guiding norms falls short of an ideal. In such cases the two norms conflict, because reality falls short of the ideal.

So it is with epistemic norms. Different epistemic norms may express the ideals to which our beliefs should aspire, or the ideals to which our belief systems should aspire, or what beliefs would be ideal given our evidential situations; or they may answer many other sorts of questions, at different levels of idealization. The more practical, action-guiding norms concerning how we should form beliefs, gather evidence, or form belief-forming habits are only some of the epistemic norms, and beliefs formed in accord with them may fall short of the ideal in many ways. Nor can the most practical norms be derived from the more ideal norms in any straightforward way. Each point on the epistemic spectra represents one way we might approach epistemology, one sort of problem we might be interested in. When we seek an epistemic norm we must choose a

point; we must think about just what epistemological problem we are interested in. There is no single norm that suits every perspective on epistemology.¹⁷

Works Cited

Bishop, Michael A. and J.D. Trout (2005). *Epistemology and the Psychology of Human Judgment* (Oxford: Oxford University Press).

Feldman, Richard (1985). "Reliability and Justification," *The Monist* 68, 159-74.

Goldman, Alvin (1979). "What Is Justified Belief?" In G. Pappas, ed. *Justification and Knowledge* (Dordrecht: D. Reidel), pp. 1-23.

Harman, Gilbert (1986). *Change in View* (Cambridge, MA: The MIT Press).

Kvanvig, Jonathan L. (2003). *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press).

Lackey, Jennifer (2007). "Why We Don't Deserve Credit for Everything We Know." *Synthese* 158, 345-361.

Montmarquet, James A. (1987). "Epistemic Virtue." *Mind* 96, 482-497.

Quine, Willard Van Orman (1969). "Epistemology Naturalized." In *Ontological Relativity and Other Essays* (New York: Columbia University Press), 69-90.

Stich, Stephen P. (1983). *From Folk Psychology to Cognitive Science: The Case against Belief* (Cambridge, MA: The MIT Press.)

Weiner, Matthew (2005). "Why Does Justification Matter?" *Pacific Philosophical Quarterly* 86, 422-444.

Weiner, Matthew (2009). "Practical Reasoning and the Concept of Knowledge," in Adrian Haddock, Alan Millar, and Duncan H. Pritchard, eds., *Epistemic Value* (Oxford: Oxford University Press, 2009), 163-182.

Williamson, Timothy (2000). *Knowledge and Its Limits* (Oxford: Oxford University Press).

¹⁷ Thanks to Ram Neta and Chase Wrenn for discussion of related topics and especially to John Turri for his comments and suggestions.

Zagzebski, Linda Trinkaus (1996). *Virtues of the Mind* (Cambridge: Cambridge University Press).

Biography

Matt Weiner is Associate Professor at the University of Vermont. Topics he has written on include the norms of assertion, the epistemology of testimony, epistemic values, the analysis of knowledge as an inconsistent concept, and conversational implicature.

Keywords

Epistemic norms, knowledge, justification, understanding, epistemic virtues, internalism, externalism, epistemic value pluralism