

Lecture 4

Epistemic Normativity

We humans are zestfully judgmental across the gamut of our experience: in art, literature, science, politics, sports, food, wine, and even coffee; and so on, across many other domains.¹ We love to evaluate even when no practical interest is in play. We judge performances, whether artistic or athletic; grade products of craft or ingenuity; evaluate attitudes, emotions, institutions, and much more.

Any such domain of human experience admits values of two sorts: the derivative, and the fundamental—that is to say, the derivative or fundamental *for that domain*. A value might be irreducible to other values distinctive of a given domain, without being fundamental absolutely, since reducible to other values *beyond* that domain.

According to epistemic truth monism, truth is the fundamental epistemic value.² The epistemic justification of a

¹ In a narrow, praxiological, sense, normativity pertains to choice or action, and to rules or standards for choosing or assessing conduct. In an extended, axiological, sense, it pertains to the evaluative more generally, whether the objects of evaluation be actions or not. The two are intimately related, since conduct that brings about something intrinsically good is to that extent apt, even if it must meet further requirements of rational control and intentional guidance. Here we take the broader view.

² One might defensibly define generic reliabilism as a truth-monistic view. But there are various equally defensible ways to understand that important, but flexible, epistemic term “reliabilism.” On some of these, externalism is not crucial, and Descartes counts straightforwardly as an extreme, *infallibilist*, reliabilist.

belief, its epistemically positive status beyond that of being true, is held to involve truth-conducive reliability, however conceived, whether as tracking the truth, or as deriving from a reliable process, or competence, or virtue. A true belief is said to constitute knowledge only through some such connection with the truth.

In calling a belief knowledge, we evaluate it positively by epistemic standards. Within the domain of epistemic assessment, knowledge has a standing higher than that given to its constitutive belief by its mere truth. But how can this be, if truth is the fundamental epistemic value? Suppose a belief is epistemically justified if and only if it derives from a truth-reliable source, because what matters essentially and distinctively in epistemology is whether and how we are in touch with the truth. In that case, once true, a belief would seem to gain nothing further from being thus justified.

Our worry goes beyond the Platonic worry of how *knowledge* that a certain road leads to Larissa could be better than *true belief*, if either will get you there equally well. The Larissa worry is assuaged by a distinction between epistemic and pragmatic varieties of justification. The efficiency of our belief in getting us to Larissa is to be distinguished from its distinctively *epistemic* normative status. Knowledge could of course have a value beyond such efficiency. But the belief's epistemic status must then concern more than just how well it guides you to your objectives. Not even by restricting ourselves to *epistemic* objectives, such as that of gaining truths, can we reduce epistemic justification to instrumental value. You may lack justification for trusting a certain book, for example, even if it would reveal a trove of truths. This sort of epistemic efficacy does not even protect the belief that the book is trustworthy from being wildly unjustified. Despite

overwhelming evidence that the book is not to be trusted, believing it trustworthy might still give you that trove of truths.

Our present worry abstracts from such Platonic issues of epistemic normativity. Truth may or may not be intrinsically valuable absolutely, who knows? Our worry requires only that we consider truth the *epistemically fundamental* value, the ultimate explainer of other distinctively epistemic values.

Our issue for truth-centered epistemology, beyond the Larissa problem, is the “value problem,” as follows:

How can the truth-reliability of an epistemic source give to the beliefs that it yields any additional epistemic worth, over and above any that they already have in virtue of being *true*?³

Compare this. A cup of coffee is not a better cup of coffee for being yielded by a good, reliable machine! It will be useful here to compare (a) the realm of the epistemic, with (b) the world of coffee. Are these critical domains similarly structured? A tabular comparison is shown in Table 4.1.

³ Our issues of epistemic normativity are discussed in a growing literature that includes the following early contributions: Ward E. Jones, “Why Do We Value Knowledge?,” *American Philosophical Quarterly* 34 (1997): 423–39; Jonathan L. Kvanvig, “Why Should Inquiring Minds Want to Know? *Meno* Problems and Epistemological Axiology,” *The Monist* 81 (1998): 426–51; Linda Zagzebski, “From Reliability to Virtue Epistemology,” in G. Axtell (ed.), *Knowledge, Belief, and Character* (Lanham, MD: Rowman & Littlefield, 2000); Marian David, “Truth as the Epistemic Goal,” in M. Steup (ed.), *Knowledge, Truth, and Duty* (Oxford: Oxford University Press, 2001); Michael DePaul, “Value Monism in Epistemology,” in Steup, op. cit.; Wayne Riggs, “Reliability and the Value of Knowledge,” *Philosophy and Phenomenological Research* 64(2002): 79–96. In the present lecture I develop an approach sketched in “Beyond Skepticism, to the Best of Our Knowledge,” *Mind* (1988), and in “Reflective Knowledge in the Best Circles,” *Journal of Philosophy* (1997), reprinted in Steup, op. cit.

Table 4.1 Epistemic comparison

<i>Epistemology</i>	<i>Coffee criticism</i>
Beliefs	Liquid coffee
Concepts	Ground coffee, coffee beans, coffee makers
Believers and their ways of forming and holding beliefs	Baristas and their ways of making liquid coffee
Schools, modes and methods of teaching, laboratories, modes of inquiry, epistemic communities, criteria of assessment and promotion	Coffee plantations, harvests, ways of grinding, kinds of land, climate, relation to the sun, etc.
True beliefs and theories	Delicious, aromatic liquid coffee

Consider the world of coffee—of its production, elaboration, and consumption. One central value organizes the critical assessment distinctive of that domain. I mean the value of liquid coffee that is delicious and aromatic. Think of the assessment of coffee beans, fields, coffee machines, baristas, ways of making liquid coffee, plantations, harvests, etc. What organizes all such evaluation, the value at the center of it all, from which the other relevant values are derivative, is the value of good coffee, of liquid coffee that is delicious and aromatic. (We leave aside the use of coffee in recipes, liqueur, ice cream, etc.; these may also be fundamental, but they are peripheral.)

The world of coffee is a “critical domain,” a set of interrelated entities evaluable through correspondingly interrelated values. Paradoxically, one can be an adept critic within such a domain even while discerning in it no domain-transcendent value. Thus, someone knowledgeable about guns and their use for hunting, for military ends, and so on, may undergo a conversion that makes the use of guns abhorrent. The good shot is thus drained of any real value that he can discern. Nevertheless,

his critical judgment within that domain may outstrip anyone else's, whether gun lover or not. Critical domains can be viewed as thus *insulated*, in ways suggested by our example.⁴

Instrumental value is among the kinds of value that derive from more fundamental values. When a barista does a "good" job, he does things that produce some good coffee. The barista does "well" because his actions result in the good coffee. But extrinsic value takes many other forms besides the instrumental. What makes for the goodness of good fields, for example, or good beans, good coffee makers, or good baristas, is not exactly their efficiently causing something good at the time when they are "good" exemplars of their kind. Nevertheless, efficient production of the fundamentally good is also implicated in these other varieties of goodness. Given its importance, not only on its own, but also in explaining other forms of evaluation, instrumental goodness deserves a closer look.

If your taking some aspirin relieves your migraine without side effects, then it is probably a good thing you do, good in virtue of its good effects, including the removal of the headache. Let's here assume that the relief is itself a good, so that the taking of the aspirin produces something good (or, if that is found implausible, suppose you down the pill with a good Scotch). On that assumption, let us now abbreviate as follows.

T = the taking of the aspirin (with the Scotch).

P = the pleasure produced.

When we say to the subject, about T, "that's a good thing you did," we base ourselves on what T brought about. What

⁴ A chess master, such as Bobby Fischer, might become depressed and disenchanted with chess, while retaining his preeminence at the chessboard, and his peerless critical judgment.

made it a good thing to do is its causing something good, the relief (or pleasure) that ensues.

Compare now the following two facts:

The causal fact: <T produces P>.

The temporal fact: <T precedes P>.

Only the causal fact gives rise to a widely recognized, distinctive sort of value, the instrumental value that T inherits from it. The temporal fact does not bestow on T any recognized sort of value. But why should that be? Isn't it just arbitrary to distinguish thus between the causal relation and the temporal relation? Why not recognize a kind of "precedence value" related to the temporal fact as is instrumental value to the causal fact? Why not even allow a "coincidence" value that T can have through its mere *coexistence* with the valuable P?

Distinguishing instrumental value from these other definable forms of "value" is hardly arbitrary, surely, though we cannot stop now to detail why. For some reason, causation enters into the proper determination of values beyond the intrinsic in a way denied to mere precedence and coexistence. Causation helps induce other values, instrumental value prominent among them, in a way that mere temporal relations and bare coexistence are powerless to do.

Moreover, causation can work in either direction, by helping induce further value in a cause, or alternatively by helping induce further value in an effect, as in the following example.

A ballerina's body moves across a stage with utmost grace. We take great pleasure in those graceful movements; we admire and applaud them. Then we learn that the dancer was drugged, her movements mere stumbles, nothing more. How now do we assess what happened, however improbably, on that stage?

One might of course still admire the movements, as we admire the swelling flow of Niagara, or falling snow in bright sunlight, as natural phenomena with immediately appreciable beauty. Many would react that way. For others, the disappointment would be so great as to obscure any natural beauty to be found in the movements themselves, however produced, whether through artistic control or through stumbles. Why disappointment? What we had paid for, surely, was not just to see stumbles, however graceful. We had paid to see a performance, the product of artistic excellence and control. We take pleasure in seeing the grace of the movements, true enough, but we take special pleasure in knowing it to be grace due to the ballerina and, more particularly, to her art. It is her actions that we normally admire and value.

One and the same item does not in that example clearly gain a kind of value through how it is caused, however, just as the taking of the aspirin might or might not have a kind of value, instrumental value, depending on what it causes. In our example, causation plays a crucial role. It's the movements under the dancer's *control* that we admire. But it is not clear that causation is there adding value to something that might still have existed but would not have had that value without its causal liaison.

The plastic arts provide a better example. We come across a canvas in a museum and admire the design on it. What if the canvas turns out to be a patch covering a wall under repairs? What if the design got there when the canvas rubbed against something on the way to the wall? What if a still un-removed sign below the now removed painting had attributed the design to Picasso? Do we still value and admire that design just the same once we learn its true origin? Surely not, but what makes the difference?

It is not the action of Picasso painting a painting that we admire when we view a real Picasso, as we admire the ballerina's dance. The artist's action might be admired when he is seen at work in his studio. But it is not what we admire now as we view the painting in a museum. Nor do we admire the painting's having been painted by Picasso. Our admiration need not even grow much when we learn the artist's name and biographical profile.⁵ What does normally matter is whether the design is owed to a mere accidental rubbing or to an artist's genius. What is valued, however, what is assessed as valuable, is the work of art itself, the painting, not the fact of its having been created by someone or other, someone whose identity we may not even know. What is admired is that design, those lines and colors on the canvas, not the fact that they were authored. Nevertheless, its origin does matter, as noted: an accidental rubbing, again, is not a work of art.

Something similar holds for epistemology. To begin, Table 4.2 compares two domains of normative criticism, one epistemic, one athletic.

The good shot is the central value that organizes the sport of archery and the criticism proper to it. Think of how we grade other things distinctive of the world of archery. All such evaluation is dependent on the value constituted by the good shot. Consider, for example, what determines the quality of bows, arrows, archers, archery schools, methods, training camps, and so on. This all depends on the value of the good shot, which is fundamental to

⁵ The value of a painting *can* of course depend on its specific authorship. This is reflected in its market value.

Table 4.2 Comparison of two domains of normative criticism

<i>Epistemology</i>	<i>Archery criticism</i>
Beliefs	Shots
Concepts	Bows, arrows
Believers and their ways of forming and sustaining beliefs	Archers and their ways of shooting
Schools, methods of teaching, modes of inquiry. Epistemic communities, criteria for rating and promoting	Communities that preserve, supplement, and transmit the lore of archery, honoring accomplishment in accordance with criteria
True beliefs	Accurate shots
True beliefs might vary in epistemic respects; for example, some have more content than others, being more specific	Accurate shots might vary in respects relevant to archery; for example, some come closer to the bull's-eye than others

archery criticism, through the ways in which it is constitutive of other values proper to that field: the quality of bows, for example, or of arrows, archers, methods, etc.

Truth is similarly a fundamental value of epistemology. Evaluation is distinctively epistemic when it is concerned with truth. Granted, belief is pragmatically evaluable. Hospital patients and competitive athletes might be helped to prevail through confidence, which is thus well placed pragmatically, even when evidentially baseless.

A shot might be accurate without being adroit. It might hit the bull's-eye aided by a gust of wind without which it would have missed the target altogether. In that case, the shot is accurate though unskillful. Contrariwise, a shot might be inaccurate though skillful. Perhaps it would have hit the bull's-eye but for the gust that diverts it.

A shot is adroitly accurate, then, only if it is accurate while manifesting the archer's skill. This much is good as far as it

goes, but it might still fall short. An archer might manifest sublime skill in a shot that does hit the bull's-eye. This shot is then both accurate and adroit. But it could still fail to be accurate *because* adroit. The arrow might be diverted by some wind, for example, so that, if conditions remained normal thereafter, it would miss the target altogether. However, shifting winds might then ease it back on track towards the bull's-eye. Though accurate and adroit, this shot would still fail to be "apt," that is to say, accurate *because* adroit.

Aptness depends on just *how* the adroitness bears on the accuracy. The wind may help some, for example; it may even help enough that the arrow would otherwise have bounced off the side of the target on its way to the ground. Only with the wind's help does it bury its tip near the bull's-eye. If the shot is difficult, however, from a great distance, the shot might still be accurate sufficiently through adroitness to count as apt, though with some help from the wind.

An index of sufficiency seems required, with some threshold, probably contextually determined, so that we can affirm this:

A shot is apt if and only if its accuracy is due "sufficiently" to the archer's adroitness.

What does such "sufficiency" depend upon? This is a difficult and interesting question that we must here postpone. Better yet, we might do well to abstract from the threshold-requiring classificatory concept to its presupposed comparative, as follows:

How apt a shot is varies in direct proportion to the adroitness manifest by the archer and to how much its accuracy is due to that adroitness.

Alternatively, we might understand success due to an agent's competence as success that *manifests* that competence, a special case of the manifestation of a disposition. But we cannot tarry over this promising alternative.

Epistemology too, like the aesthetics of dance, reverses the import of causality found in instrumental value. The distinctively epistemic evaluation of a cognitive performance can depend substantially on its source, unlike the instrumental evaluation that depends on effects rather than sources. Consider thus the justification of a belief derived from a good inference, as when a detective figures out who did it, or when you determine how much you owe a shopkeeper. Something is then believed because it is concluded from prior information already in the thinker's possession. To draw it as a conclusion and to believe accordingly for that reason is, moreover, a broadly causal matter. It is a matter of believing such and such *because* of so and so, or *on the basis* of a prior belief that so and so. Accordingly, the conclusion belief **gains its epistemic status through being *based* on the premises inferentially.** One believes the conclusion at least in part on **that basis, for the reason** that, as one can see, it follows from **the already accepted information.** The fact that one's belief in **the conclusion** is thus "motivated rationally" helps to make it epistemically appropriate, a rationally justified believing.

How must a belief be related to a fact if it is to be knowledge of that fact? According to one proposal, the belief must be safe, **that is, one that the believer *would then hold only if correct.*** A **second proposal** requires the belief to be apt, correct in a way **creditable to the believer,** as determined by how salient is the **believer's competence in the explanation of his being right.**

Although both conditions still seem defensible in some form,⁶ each needs qualification. Although the right requirement in this vicinity is one of aptness, this is not to be explained just by appeal to explanatory salience, or to avoidance of luck or accident. The reasons emerge already with the problems canvassed earlier⁷ for any unrestricted requirement of safety.

A virtuous performance, whether a correct belief due to intellectual virtue or a right action due to practical virtue, will involve both the agent's constitution and his situation. If the act is due to a competence exercised in its appropriate conditions, its success may be due to luck in various ways. It may be just an accident that the agent retains his relevant competence, for example, or that the conditions remain appropriate. Either way, the act fails to be *safely* successful, since it might too easily have failed, through lack of the required competence or conditions. It might still be apt, nevertheless, indeed attributably, creditably apt.

A certain archer's shot hits the mark through a normal exercise of skill, let us suppose, in normal circumstances. What if the archer might easily have been disabled, having just taken an unadulterated drink at random from a collection nearly all adulterated? What if a gust of wind or stroke of lightning *might easily* have denied him his propitious situation for part at least of the relevant period, by affecting the arrow on its way to the target? Even so, the shot might have been apt, surely, still accurate because adroit, and creditable to

⁶ More recently the idea has been developed insightfully by John Greco; see his "Knowledge as Credit for True Belief," in Michael DePaul and Linda Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford: Oxford University Press, 2004).

⁷ In Lecture 2.

the agent, so long as the competence remained in place, and the conditions appropriate, even if only by luck. What matters is that the conditions remain appropriately normal (or better) along dimensions relevant to the agent's retained competence.

The same goes for any apt belief accurate due to the subject's competence. A belief can be unsafe because the subject might too easily become disabled, or because the conditions might too easily become inappropriate. Suppose you see a surface to be red, for example, by exercising your visual competence in good light. Perhaps you took an unadulterated drink at random, however, from amongst many containing a color-blinding drug. Or perhaps the surface might too easily have been lit with red light, making it impossible to see red patches in their true colors. In such circumstances, where too easily you might have lacked the requisite competence or conditions, your color belief would have been unsafe. Nonetheless, it might have been apt even so, still correct due to competence. Or so one would think, if the archer's shot can be apt despite clear and present threats to his competence, and to the appropriate normalcy of his situation.

Our reasoning distinguishes between (i) factors because of which the circumstances *might* now easily have failed to be normal, without *already* being abnormal, and (ii) factors that *do* already preclude normalcy. Factors of sort (i) make a belief unsafe without precluding its being apt, that is, correct because adroit. Factors of sort (ii) deprive the belief not only of safety but also of aptness.⁸

⁸ More strictly, if the subject's performance is to be apt, then the conditions must be normal or better. But this requires a prior or at least coordinate conception of the relevant adroitness or skill or competence. What is required in the conditions is that they be normal or better for the exercise of the relevant competence.

Outright safety is not a requirement for apt performance in any case. A performance can still be apt when its safety depends on a circumstantial contingency, provided it is guided by that contingency. Even if the light alternates quickly and randomly between being bad and being good, one can still acquire perceptual knowledge so long as the deliverances of one's color vision are accepted, not just at face value, but guided by the ringing of a bell bound to ring steadily when and only when the light is good.

Think back, however, to the easily possible gust of wind or stroke of lightning, or to the lighting conditions that might easily have been spoiled. One's performance could then still be apt despite failing to be safe outright, since safe only dependently on a factor that might too easily have been absent. Here the agent's performance need not even be guided by the factor dependently on which it is apt. Question: what distinguishes such contingencies, those on which the performance can depend for its safety and aptness without benefit of guidance?

Our purposes in evaluating people plausibly help determine that distinction, given our need for coordination and mutual reliance, and hence for keeping track of strengths and weaknesses, our own and others'. We value, as "aptitudes," certain abilities relative to certain background conditions. Such abilities are relative to distinctive correlated parameters at the time of their exercise. That you have the ability means that your relevant performances tend to succeed when they fall within those parameters. This applies to athletic prowess, intellectual abilities, crafts, skills, and so on. Failed attempts

This is why the lucky gust does not make the conditions better than normal, enabling the archer to earn credit for a great shot. The gust creates conditions wherein it is not any skill manifest by the archer that is operative. The gust takes over independently of anything that could count as such a skill.

in abnormal circumstances do not show lack of the ability. Despite such failures we might still depend on you in normal circumstances. What is required is only that your attempts tend to succeed when circumstances are normal. (And the like seems true of dispositions generally, not just abilities or competences, including the likes of solubility and fragility. Normality, moreover, is not just statistical normality; it is a kind of normality determined implicitly, for any given dispositional concept by those who share that concept.)

Accordingly, there are at least two ways in which your attempt might be unsafe while still apt: first, although not safe *outright*, the attempt might be safe dependently on a circumstantial contingency, awareness of which guides your performance or motivates it rationally; second, although not safe outright, the attempt might be safe dependently on circumstantial normality in various respects, even despite being unguided by any such respects, and rationally motivated by none such.

What is required for aptness is that the performance succeed through the exercise of a competence in a situation **appropriately normal** for that exercise. A performance that is safe **only** dependently on a certain contingency, will be apt **only** if that contingency either guides it or is constitutive of the relevant normalcy of the situation.

The archer's skill is a state that reliably yields accurate shots **when applied** in normal circumstances. The exercise of the **archer's skill** is a rational activity in that the archer is guided **by reasons**. The archer is motivated by reasons to release the **arrow** when the bow and arrow are held just so. He may **be unable** to articulate these reasons, but we cannot plausibly **require** that our reasons must always be articulable, lest we **deprive** ourselves of reasons that matter to us as much as

anything. We identify a loved one, for just one example, in ways that we could not articulate fully. We recognize the loved one on sight based on reasons whose full verbal articulation we cannot plausibly require. Similarly, the archer has reasons for releasing the arrow depending on just how it is then positioned. Of course, if asked to explain why he released it at that point, he could say little beyond "It just felt right to do so." Depending on the target's size and distance, however, it will feel right to release the arrow when positioned in quite different ways. When the archer takes a shot, things relevantly appear quite differently from how they do when he takes another shot. Yet in all such cases it will tend to feel more or less equally right to release the arrow at just the point where the archer does so. However, the specific feeling that qualifies as feeling right will be different from situation to situation, and it is the specific feeling that guides the excellent archer to release the arrow when he does.

The commitments that inform different archers' skill, with respect to various apparent situations, to release the arrow when things appear *thus*, can vary vastly in reliability. The better the archer, the more reliable his commitments; and vice versa.

Some of our relevant commitments come courtesy of Mother Nature and her evolutionary ways, but many others must be learned. The archer's learning requires practice, and seems tantamount to inductive learning. An aspiring archer tries various things, takes note of the degree of resulting success, adjusts accordingly, and with luck remembers what he has learned. What he has learned is not articulable, not fully, but that seems epistemically inessential. He has learned something. When he next takes a shot, he does so based on what he has learned, provided he remembers it well enough.

The same goes for intellectual practice. Much of our intellectual competence comes with our brains, but much is due to learning. The commitments constitutive of learned competences are a varied lot. Some come through explicit instruction, but many come through life experience of enormous scope and variety. Much of our competence, whether practical or intellectual, requires memory. We need to retain it, and it can be lost in various ways, from the localized losses due to lack of practice to the ravages of Alzheimer's.

When a success, practical or intellectual, is creditable to an agent, it is due to an aptitude (to a competence or skill or virtue) seated in that agent, whose exercise is rewarded with success in his act or attitude. Concerning such success, how are we to understand its being *due* to the agent's aptitude? According to one promising proposal, its explanation must saliently involve that aptitude.

We face problems here that mirror problems encountered earlier, by the safety proposal. Perhaps, for example, what in the circumstances is explanatorily most salient concerns why the agent retains his competence, or why the situation remains normal. Thus, the evil demon in charge may systematically spoil the competence of agents in an archery competition, or the circumstances of their shots, while making an exception of our successful archer for one of his shots. For that one shot he does not disable the competence or spoil the circumstances. **Against that background, what is then explanatorily salient, when we ask why that shot was successful, concerns more the doings of the demon than those of the archer. Despite that, the archer does surely hit his target aptly: his shot is accurate because adroit. It seems irrelevant that only luck accounts for his retained competence and propitious circumstances.**

Somehow it is the exercise of competence in a normal situation that makes the shot apt: that is, accurate because adroit. That it is apt by luck makes it no less apt.

Let us return, finally, to the value problem for reliabilist conceptions of epistemic justification. The problem is put in proper context only with a distinction between generic reliabilism and virtue epistemology. More specifically, we need to understand that, according to an instrumental conception of justification, a belief is justified by deriving from mechanisms or processes of belief acquisition that reliably deliver true beliefs. True belief seems here installed as a fundamental value of epistemology, while instruments and processes of belief formation are then assessed in respect of how well they deliver that value, and beliefs are granted epistemic status depending on how reliable are the sources from which they derive.⁹ This whole way of thinking of epistemology invites the comparison with good coffee and the conundrum as to how true belief could be better for deriving from good sources, if good coffee is no better for deriving from the adroit use of a good coffee maker.

One part at least of the solution to the value problem lies in a point central to virtue epistemology: namely, that the value of apt belief is no less epistemically fundamental

⁹ I speak of "mechanisms" or processes of belief formation, and sometimes of "input/output mechanisms," but I want to disavow explicitly any implication that these are simple or modular. The process can of course be as subtle and delicate as that of determining whether the butler is guilty, or whether a remark was a deliberate insult, or whether vagueness is epistemic. The epistemic virtues involved need not be as simple as the perception that one seen line is longer than another, though even here there is more than meets the eye, as is shown by the Müller-Lyer illusion. Thus, a mechanism can be something close to a reflex, or it can be a very high-level, central-processing ability of the sort that enables a sensitive critic to "decide" how to assess a work, based on complex and able pondering.

than that of true belief.¹⁰ For this imports a way in which epistemic virtues enter *constitutively* in the attainment of fundamental value, not just instrumentally.¹¹ Virtues are thus constitutive because the aptness of a belief is constituted by its being accurate because competent. This means that the relevant competences of the agent enter into the constitution of something with fundamental epistemic worth: namely, the apt belief, true because competent.¹²

Our subject has been epistemic normativity, a kind of normative status that a belief attains independently of pragmatic concerns such as those of the athlete or hospital patient. Epistemic normativity is a status by having which a true belief

¹⁰ And, perhaps, no more: the true belief, like the accurate shot, is a “success” in its own right. The accurate shot is surely, to that extent, a “successful” shot.

¹¹ A full account would need to recognize in addition the epistemological value of good omissions, as when one suspends judgment. Cf. Sosa 1991, p. 17.

¹² I am here opting for something like the ballerina’s graceful *pas* as correlate of the knower’s believing. The apt believing is a performance-immanent value, unlike the cup of coffee vis-à-vis the doings of the barista. Plenty of room remains in the critical domain of epistemology, however, for the kind of value that we find in the work of art that hangs on the museum wall. Take those magnificent works of art, powerful scientific theories. These are often the work of many hands acting collectively. Yet they are obviously important objects of epistemic evaluation. Nor is the only evaluation possible that which is capturable through abstract data-hypothesis relations. The history of science shows how much room there is for further evaluations involving why the theory holds sway in a certain community, including our own, why it is a “live” theory, etc. Some of the relevant factors that explain why it is live, and even held, are only pragmatically relevant, but the abstract observational data/hypothesis relations are unlikely to exhaust those that bear epistemically on the theory’s being *properly* (epistemically) live, or *properly* held, etc. Other virtues would seem epistemically pertinent, some of them seated in institutions and not just in individuals distributively. (Developing this point satisfactorily would require distinguishing the theory as a cultural artifact whose “life” depends on how it relates to the community in which it is live from potential theories as just abstract truth bearers. This would be analogous to the distinction between the abstract sequence of propositional contents that constitutes the abstract “play” *Hamlet* throughout eternity, and *Hamlet* as the living cultural artifact, the great play created by the genius of Shakespeare.)

constitutes knowledge. We must distinguish the normative status of knowledge as knowledge from the normative status that a bit of knowledge may have by being useful, or deeply explanatory, and so on. Something might be known far better, with greater certainty, and better justification, than something else, while yet the latter knowledge is intellectually finer by far. Compare knowledge that one’s back hurts now with some deeply illuminating knowledge about a friend, or a historical period, or a novel, or a scientific theory.

Accordingly, we do well to distinguish between two parts of epistemology: (a) theory of knowledge, and (b) intellectual ethics. The latter concerns evaluation and norms pertinent to intellectual matters generally, with sensitivity to the full span of intellectual values. It is therefore a much broader discipline than a theory of knowledge focused on the nature, conditions, and extent of human knowledge.

That distinction between theory of knowledge and intellectual ethics has drawn a skeptical reaction. Suppose Paul forms a belief about the number of motes of dust on his desk by consulting an Ouija board. According to our proposal, we might then (a) from the point of view of a theory of knowledge, evaluate the belief poorly because the belief is so unsafe with respect to the truth of the matter, and (b) from the point of view of intellectual ethics, criticize the believer for even bothering with such things—from the point of view of his intellectual flourishing, we might say, there are better ways he should be spending his time.

One might be skeptical about this for the following reason:

It seems that we do more than just evaluate the belief negatively from an essentially *performance* point of view when we say that Paul’s belief is unjustified. Rather, we

subject him to normative criticism for forming a belief this way: we reproach him and find him in some sense blameworthy. Importantly, we do this not *just* because he should have been doing better things with his time, from the point of view of intellectual ethics, but rather because, even with respect to this trivial subject, he has in some sense let us down.

True enough, but the considerations involved in such criticism would still go beyond those involved in the normative status that is *constitutive* of knowledge. The situation seems rather like that of a hunter in a hunting society, in conditions of scarcity, who used energy and resources (arrows) shooting at trivial targets (not live game, say, nor enemies, etc.), and did so carelessly. He would face a loss of credit, and prestige, and might thus forfeit a good position for the next hunt or battle. But this all goes beyond what makes his shot a poor shot. **Much** later, when the society is agricultural, absent hunting, **absent** battles, absent even any sport of archery, if someone **nevertheless** dusts off a bow and some arrows and takes a **shot**, that shot might be equally poor as a shot, but would **incur** no such criticism. In an archery-internal sense the shot **falls short** equally. So, in theory of archery it falls short, and is a “**poor**” shot. But in respect of the larger ethics of archery, **or** of the use of bows and arrows, its evaluation would be **quite different**, and not just in respect of the importance of **the targets** chosen.

One might be skeptical, moreover, because it is so hard to **say** much in general about the kinds of questions that are **worth** pursuing from an intellectual point of view. Thus it is **very hard** to say why it would be appropriate, from the point of view of intellectual ethics, to criticize someone who spent

his time happily learning everything there is to know about some trivial thing (say, the history of bubble gum), while someone who spent their time learning everything there is to know about (say) the us Civil War would escape criticism.

It is worth noting, in response, that one may know that something is so without knowing why it is so. We might know it better if we did know why it was so. But we can know very well that we have a headache even without knowing its cause or explanation. The criticism based on the gum aficionado might indeed just be that he is spending too much time on something too trivial, so that *unless there is some compensating reason*, he can be criticized all things considered. Of course, if he is unable to muster any interest in anything else, and is for whatever reason obsessed with the history of bubble gum, and this is his only way to stave off deep depression, then it’s fine. And so on. In any case, recognizing the idea of intellectual ethics, the domain of such evaluations and critical judgments, does not entail any commitment to the idea that it would have to be a domain that we could ever come to understand with a powerful explanatory theory.

Most of the history of epistemology has had the narrower focus of the theory of knowledge. Interest did eventually shift from the focus on the nature of knowledge in the Theaetetus, towards an interest in how far we can be justified in our beliefs generally. But this latter question is that of how extensively we can attain the kind of epistemic justification and aptness that is constitutive of knowledge. It is therefore still a concern in the theory of knowledge, not in intellectual ethics, and so it remains to the present day.