



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions

Citation for published version:

Chrisman, M 2012, 'The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions', *Journal of Philosophy*, vol. 109, no. 10, pp. 588-612.
<<http://www.journalofphilosophy.org/articles/>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Published In:

Journal of Philosophy

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



The Normative Evaluation of Belief and the Aspectual Classification of Belief and Knowledge Attributions^{*}

Matthew Chrisman
University of Edinburgh

It is a piece of philosophical commonsense that belief and knowledge are states. Some recent virtue epistemologists have been tempted to ignore this common sense because they think doing so is *the* key to some of the open and difficult questions in epistemology. In my view, however, they are wrong to do so, especially when it comes to two important questions about the normative evaluation of belief.

The basic problem, I shall argue, is that denying that belief and knowledge are states offends not only against philosophical commonsense, which is a popular whipping boy, but also against ordinary common sense, at least as far as this is manifested in the meaning of the words we ordinarily use to talk about belief and knowledge. More specifically, I think it is obvious that ordinary belief and knowledge attributions would be classified aspectually as state descriptions in any adequate semantic analysis. This can be shown relatively easily with some linguistic tests, and it illustrates the way philosophical commonsense on this topic is an outcropping of ordinary commonsense rather than mere dogma.

Hence, moves against the view that belief and knowledge are states threaten to simply change the topic rather than answer open and difficult questions in epistemology. I do not know how to answer fully the two questions about the normative evaluation of belief that will be my focus here, but I pursue this critical point against those who would deny that belief and knowledge are states in service of a positive proposal about the general framework in which these questions should be answered.

In brief, the general framework is one that recognizes an important place for what I call state-norms, beside the action-norms that are much more familiar from discussions of normativity in ethical theory. My claim will be that norms of belief are state-norms and this is an important point often elided in recent discussions of normativity in epistemology. I believe there are epistemically relevant action-norms, but in my framework, these are viewed as governing the kinds of actions which are parts of the practices of inquiry and instruction. In my view, these action-norms are connected in complicated ways to the norms governing beliefs, but they are also importantly different in their presuppositions about capacity for choice. This is why I think several dominant approaches to epistemology make an important mistake in assuming that the norms governing beliefs must either be action-norms or be very much like action-norms by

^{*} For helpful feedback on earlier versions of this material, I would like to thank J. Adam Carter, Davide Fassio, Georgi Gardiner, Allan Hazlett, Graham Hubbs, Conor McHugh, Michael Ridge, Sam Wilkinson, editors for this journal, and audiences at the University of Bled, University of Edinburgh, University of Georgia, Northwestern University, and the Australian National University. I appreciate support for research on this paper from the Arts and Humanities Research Council.

applying analogically to a genus of which actions are a species: performances. By contrast, I defend a framework here where the norms governing belief are treated as state-norms. This, as you might suspect, is not only consistent with but indeed underwrites the philosophical and ordinary commonsense that belief and knowledge are states.

Below, I shall have more to say about my proposed framework for thinking about the normative evaluation of belief and its connection the aspectual classification of belief- and knowledge-attributions, but I start with a brief characterization of the two open and difficult questions in epistemology that motivate some to ignore the fact that belief and knowledge are states, and for which I would like to propose a new epistemological framework.

I

The first question is how knowledge could be better than belief that falls short of knowledge. It is rightly popular to think that knowledge is better than false or unjustified belief, because a belief does not count as knowledge unless it is true and justified, and these are good properties for a belief to have. However, as Gettier taught us, a belief may be true and justified and yet not be knowledge.¹ Nevertheless, it still seems that knowledge is better from an epistemic point of view than even such "Gettierized" beliefs. Although Gettierized subjects believe a truth and their belief in this truth is justified, there is still something less than ideal about their epistemic situation; the justification is not well connected to the truth. This is puzzling. For it means that it cannot be the truth and justification that fully explains why knowledge is better than belief that falls short of knowledge. But if it is not that, then what is it? I shall call this the *axiological question*. I see it as one of the open and difficult questions in contemporary epistemology.²

The second question is how it could be intelligible that we make deontological evaluations of beliefs, given that our beliefs unlike our actions seem to be largely automatic and not objects of conscious choice. Pascal is widely thought to have taught us how we might exercise a sort of *indirect* control over some of our beliefs; we can control them roughly like we control our weight. However, that seems to be an exception that proves the rule: from when we open our eyes every morning to when we go to sleep, all sorts of mundane beliefs flood into our minds. Indeed, it even seems that we probably have a very large number of beliefs about things we are not consciously thinking about, even things we have never consciously thought about. At least, it seems to make sense to say, for example, that my mother knows the rain in Spain is not purple – this, even though my mother is not now thinking about the color of the rain in Spain and maybe never will. Yet, assuming knowledge requires belief, this means that she believes this. Indeed, we might want to say that she knows this in part because it's a belief that she *ought* to have given her evidence. But all of that flies in the face of the philosophically popular idea that deontological evaluation of a subject (e.g. with a particular species of 'ought') requires something like the sort of voluntariness or control (e.g. that we would describe with a particular species of 'can') manifested in conscious choice. So, why do we so blithely evaluate what people ought or ought not believe, given that their beliefs are mostly automatic and unchosen? I

¹ Edmund Gettier, "Is Justified True Belief Knowledge?," *Analysis* 23 (1963).

² The question, of course, traces back to Plato, "The Meno," in *Five Dialogues*, ed./trans. G. M. A. Grube (Indianapolis, IN: Hackett Publishing Co., 1981). In some ways, it was revived in contemporary theory of knowledge by Jonathan Kvanvig's *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003). See also the helpful discussions of the various value problems in epistemology in Duncan Pritchard, "Recent Work on Epistemic Value," *American Philosophical Quarterly* 44 (2007), "The Value of Knowledge," in *The Nature and Value of Knowledge: Three Investigations*, ed. Millar, Haddock, and Pritchard (Oxford: Oxford University Press, 2010).

shall call this the *deontological question*. I see it as second open and difficult questions in contemporary epistemology.³

These two questions seem to me to generate much of the torque in the oscillation between two deeply different approaches one witnesses in contemporary epistemology. On the one hand, there is an approach that we inherit from Descartes, wherein belief formation is treated very much like action. That is, at least in the ideal case, one weighs up the options of what to believe and then chooses rationally what to believe. Because of this, Cartesians can seem to be in a pretty good position to answer the deontological and axiological questions. They will say that we evaluate belief deontologically because they *are* chosen, or at least can be; and the ones which count as knowledge are better than the ones which do not because they represent better choices (somehow). However, a central problem with this approach is that it runs roughshod over the observation that most of our beliefs are not (and seemingly cannot be) chosen.

There are surely things that a staunch Cartesian will want to say about this problem,⁴ but my point is simply that it is one of the factors that have tended to rotate epistemology towards a second prominent approach: reliabilism. Here, belief-formation is treated very much like the processes of digesting food or acquiring immunity to a disease. That is, it is viewed as a kind of biological process, which like other processes of our biological systems can achieve its aims more or less reliably. In light of the popular idea that the aim of belief is truth, the familiar thought is that knowledge is something like true belief that has been formed by a reliable process. Obviously, reliabilists take very seriously the observation that most of our beliefs are not (and seemingly cannot be) chosen – most of the processes involved in digestion and immunity acquisition are not (and seemingly cannot be) chosen. Then, however, it becomes mysterious why we deontologically evaluate what people believe. We can of course say how a faulty digestive or immune system ought to work, but we are not deontologically evaluating the person whose digestive or immune system it is, unless we think they are indirectly responsible for its malfunctioning. Yet, in the epistemological case, the evaluations seem to be person-level rather than system-level – otherwise, Cartesianism would have never been a tempting view. Moreover, if we accept reliabilism, it becomes hard to see why the beliefs of Gettierized subjects are worse, from an epistemic point of view, than beliefs that count as knowledge. After all, they're true and reliably formed; what more could we want?⁵

³ See especially Bernard Williams, "Deciding to Believe", in *Problems of the Self: Philosophical Papers, 1956-1972* (Cambridge: University Press, 1973), William P. Alston, "The Deontological Conception of Epistemic Justification," *Philosophical Perspectives*, vol. 2, *Epistemology* (1988), Alvin Plantinga, *Warrant: The Current Debate* (New York: Oxford University Press, 1993). Their various versions of "doxastic involuntarism" are disputed, of course, and I am not attempting to settle the dispute here. It is enough for my purposes that very many of our beliefs are unchosen, even if not all beliefs are. Some philosophers distinguish between involuntary belief and voluntary "acceptance". See especially Jonathan Cohen, *An Essay on Belief and Acceptance* (Oxford: Clarendon Press, 1992), Michael Bratman, "Practical Reasoning and Acceptance in a Context," *Mind* 101 (1992), and Keith Lehrer, "Justification, Coherence and Knowledge," *Erkenntnis* 50, no. 2-3 (1999). Insofar as this distinction is cogent, here I mean to be talking about the involuntary belief side of it here. Whatever exactly acceptance is, I take it belief is required for knowledge, and this is borne out by the fact that we will ascribe belief and knowledge of a proposition to people who have never consciously considered it.

⁴ See Carl Ginet, "Deciding to Believe," in *Knowledge, Truth, and Duty*, ed. Matthias Steup (Oxford: Oxford University Press, 2001), and Brian Weatherson, "Deontology and Descartes' Demon," *Journal of Philosophy* 105, no. 9 (2008), for some of these.

⁵ Compare the objection to reliabilism developed in Richard Swinburne, *Providence and the Problem of Evil* (Oxford: Oxford University Press, 1999), pp. 57-66, Kvanvig, *The Value of Knowledge and the Pursuit of Understanding*, pp. 44-52, and Linda Zagzebski, "The Search for the Source of Epistemic Good,"

Some believe these difficulties should rotate epistemology back towards the Cartesian approach: perhaps we should reject the conception of belief as something that is not (typically) chosen. Others are motivated to bunker into the reliabilist position: perhaps the axiological and deontological questions themselves are implicitly derived from an anti-naturalistic worldview, which reliabilists should stand firm against rather than pander to.

There is, however, a more subtle response that promises to occupy some interesting middle ground. There are several appropriate names for this response, but *virtue epistemology* seems to have caught on, especially under the influence of Sosa's careful virtue-theoretical improvements on the reliabilist position.⁶ So, in what follows, I want to begin by outlining a central tenet of Sosa's virtue epistemology and explaining how it can be seen as a more nuanced approach than Cartesianism and reliabilism – an approach, which seems to bring important new resources to the axiological and deontological questions. (Ultimately, however, Sosa's view will be the principal foil for my objection against those tempted to ignore the fact that belief and knowledge are states.)

II

Sosa's key idea is that beliefs are *performances*. He writes, "Beliefs are a special case of performances, epistemic performances. When a belief is correct attributable to a competence exercised in its appropriate conditions, it counts as apt and as knowledge of a sort..."⁷ Similarly, he begins his most recent book writing, "Belief is a kind of performance, which attains one level of success if it is true (or accurate), a second level if it is competent (or adroit), and a third if its truth manifests the believer's competence (i.e., if it is apt). Knowledge on one level (the animal level) is apt belief."⁸ The reason Sosa thinks this is important that he thinks all performances can be evaluated along the three different dimensions: for success, skill (i.e. "competence" or "adroitness"), and whether success is because of (i.e. "manifests") skill rendering it "apt".

To explain this and how it helps with the axiological question, he uses an analogy of archery

Metaphilosophy 34, no. 1-2 (2003). But for some response see also Erik J. Olsson, "In Defense of the Conditional Probability Solution to the Swamping Problem," *Grazer Philosophische Studien* 79, no. 1 (2009), and Alvin Goldman and Erik Olsson, "Reliabilism and the Value of Knowledge," in *Epistemic Value*, ed. Haddock, Millar, and Pritchard (Oxford: Oxford University Press, 2009).

⁶ Ernest Sosa, "Beyond Internal Foundations to External Virtues," in *Epistemic Justification: Internalism Vs Externalism, Foundations Vs Virtues*, ed. L. BonJour and E. Sosa (Malden, MA: Blackwell, 2003), *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, vol. 1 (Oxford; New York: Oxford University Press, 2007), "Value Matters in Epistemology," *Journal of Philosophy* CVII, no. 4 (2010), *Knowing Full Well* (Princeton, NJ: Princeton University Press, 2011). Here, I focus Sosa's recent works because they develop one of the clearest and most sophisticated version of the basic idea. Similar ideas have been pursued by others broadly sympathetic to virtue epistemology, e.g. Linda Zagzebski, *Virtues of the Mind* (Cambridge: Cambridge University Press, 1996), John Greco, "Knowledge as Credit for True Belief," in *Intellectual Virtue: Perspectives from Ethics and Epistemology*, ed. DePaul and Zagzebski (Oxford: Oxford University Press, 2004), "The Value Problem," in *Epistemic Value*, ed. Haddock, Millar, Pritchard (Oxford: Oxford University Press, 2009), and Pritchard, "The Value of Knowledge." Here I will be assuming with Sosa and these other virtue epistemologists that knowledge is a subspecies of belief. This may be wrong, but it is such a widespread assumption in epistemology that I will not take the space to consider how the debate over the normativity of beliefs looks from the point of view of someone who denies it.

⁷ Sosa, *A Virtue Epistemology*, p. 93.

⁸ Sosa, *Knowing Full Well*, p. 1.

shots. These may succeed or not at hitting a target – this is one way to evaluate an archery shot. But whether or not they hit their target, they may manifest skill or they may be luckily successful – this is a second way to evaluate an archery shot. Moreover, when they manifest skill and hit their target, this may or may not be because of the skill manifested – this is a third way to evaluate an archery shot. A performance that is good in the first way may fail to be good in the second way, and vice versa; but a performance that is good in the third way is good in the first way because it is good in the second way. Intuitively, being good in the third way – i.e. being “apt” – is better than being good in the first and second ways, even when these are concurrent but not related.

Sosa’s leading idea then is to apply this framework to belief. If he is right that belief is a kind of performance, his virtue epistemological framework has the potential to answer both the axiological and deontological questions in ways that are sensitive to some of the motivations for the Cartesian and reliabilist approaches but better than both of them. For the axiological question, Sosa suggests that we think of belief as a performance that is successful when it is true and skillful when it is justified, but then insist that it does not amount to knowledge unless it is true *because* it is justified – i.e. successful because skillful. Then, assuming that apt performances are better than merely successful and skillful performances, Sosa’s view correctly predicts that knowledge is better than belief that falls short of knowledge.⁹ He writes, “...it would always, necessarily be proper for one to prefer one’s knowing to one’s merely believing correctly. This is just a special case of the fact that, for any endeavor that one might undertake, it is always, necessarily proper for one to prefer that one succeed in that endeavor, and indeed succeed aptly, not just by luck.”¹⁰

Moreover, this answer to the axiological question does not require the assumption that beliefs are chosen. After all, many performances are automatic, unchosen, and even unconscious. For example, understanding a sentence in one’s native language could be thought of as a performance that one does, even though (typically) this is not plausibly thought to be chosen or under one’s conscious control. Assuming this is right, it underwrites an answer to the deontological question as well. If it makes sense to talk about whether and when someone ought to understand some sentence (e.g. as an evaluation of the person’s linguistic competence), it is going to make sense to talk about whether and when someone ought to engage in the performance of believing something. That is to say that conceiving of beliefs as performances also promises to remove the mystery in our practice of deontologically evaluating beliefs.

So, it appears that Sosa’s virtue epistemology gains traction against both of our original questions about the normative evaluation of belief without succumbing to the extremes of the Cartesian and reliabilist approaches. In light of the tension and oscillation between Cartesian and reliabilist approaches, this seems to put Sosa’s virtue epistemology and frameworks like it

⁹ There are putative counterexamples to this explanation of the value of knowledge in the literature, which I will not discuss here, except to note that the objection I want to develop is much more general. Jennifer Lackey, “Why We Don’t Deserve Credit for Everything We Know,” *Synthese* 158 (2007), §2, has stressed that we get at least some knowledge from testimony, where the success is more because of the cognitive skill of the informant than the knower, which calls into question whether knowledge requires manifesting epistemic skill. And Pritchard, “The Value of Knowledge,” *op. cit.* has argued on the basis of Gettier-style barn façade cases (see Alvin Goldman, “Discrimination and Perceptual Knowledge,” *The Journal of Philosophy* 73 (1976)) that not all instances of successful-because-skillful belief are instances of knowledge.

¹⁰ Sosa, “Value Matters in Epistemology,” p. 189.

on very attractive middle ground in contemporary epistemology.

The problem, however, is that the most important idea for achieving this middle ground is that belief (and by extension knowledge) is a type of *performance*, and this idea seems to be false. That it is false is a corollary of the piece of philosophical commonsense with which I began this paper: belief and knowledge are states. And states and performances are fundamentally different kinds of things; one is by its nature passive and unchanging, whereas the other is by its nature active and changing. Virtue epistemologists may be inclined to respond by rejecting this thought as nothing better than philosophical dogma standing in the way of important progress on the axiological and deontological questions in epistemology. But, like I said above, I think this thought is not only philosophical commonsense; it is enshrined in *ordinary* commonsense via the meanings of the words ordinarily used to talk about belief and knowledge. Next, I want to demonstrate this by taking a detour through a bit of linguistic theory, which turns out to be relevant not only for appreciating what I see as a systemic problem in virtue epistemology but also for motivating the alternative framework I want to propose for working towards better answers to the axiological and deontological questions.

III

It is common in linguistics to characterize the occurrence of verbs in various phrases and sentences in terms of tense (e.g. past, present, future), mood (e.g. indicative, subjunctive), and voice (e.g. active, passive). Linguists also characterize the occurrence of verbs in terms of aspectual categories. Here the terrain is more complex and less familiar, and we see correlatively less agreement about what to call the aspectual categories, and how they are linguistically marked in English. However, it is widely agreed that there are interesting interactions between grammatical categories, such as the simple, perfect, and progressive form of verbs, and semantic features of verb phrases, such as they way they include ideas of change, duration, or telicity (i.e. whether or not they have an internally implied end-point).¹¹ Linguistic theories of aspect seek to explain this interaction.

The classical typology of English verb phrases is due to Kenny¹², who distinguishes between “statives” and “non-statives” – that is, verb-phrases referring to states such as “He owns the house” and verb-phrases not referring to states but rather something that happens or occurs. Then, harking back to Aristotle¹³, Kenny distinguishes within the latter category between “activities” and “performances” – that is, verb-phrases referring to occurrences that have no necessary endpoint such as “He is tinkering on the house” and verb-phrases referring to occurrences that have some necessary endpoint such as “He is building the house” or “He

¹¹ L. J. Brinton, *The Development of English Aspectual Systems: Aspectualizers and Post-Verbal Particles* (Cambridge: Cambridge University Press, 2009), ch. 1. Brinton distinguishes usefully between the grammatical category “verbal aspect” and the semantic classification of a verb-phrase in terms of what is known as an “aktionsart”. His argument is that these are intimately related but only confusion comes from running them together. The classical typology of English verb-phrases that I discuss below is a semantic (aktionsart) classification, which can be explained and motivated by some linguistic tests picking up on grammatical categories such as the progressive.


¹² Anthony Kenny, *Action, Emotion, and the Will* (New York; London: Routledge & K Paul, 1963), p. 171-186.

¹³ See Aristotle, *Metaphysics* 1048b; however, it is doubtful whether the distinction between *enérgeiai* (“activities”) and *kinéseis* (“movements”) plays exactly the same role for him as it does for Kenny or linguists who follow his classification of verb-phrases.

completed the house". Under the influence of Vendler¹⁴ and the latter category is now commonly divided into "accomplishments" and "achievements" based on whether the verb-phrase refers to something non-punctual such as "He is building the house" or punctual such as "He completed the house."¹⁵ (See diagram 1.)

Kenny's and Vendler's philosophical interests were in the philosophy of mind and agency, and this manifests itself in the terminology they chose for each of the sub-categories under non-statives. Words like 'activity', 'performance', 'accomplishment', and 'achievement' have their primary homes in discussions of agents of some sort or another. However, the types of linguistic tests we can use to aspectually classify a verb phrase in one of these categories apply to verbs other than verbs of agency. This is what motivated Mourelatos¹⁶ to propose the following topic-neutral terminology for the Kenny/Vendler typology (with the rejected agency-related labels in parentheses):

Diagram 1



Whatever the precise terminology¹⁷, the general idea behind the Kenny/Vendler/Mourelatos typology is that verb-phrases falling into one or another of these categories can be distinguished in terms of a number of semantic features, such as whether their referents are conceived to be dynamic, durative, or telic.¹⁸ Moreover, we can use various linguistic tests, which turn on the interaction between syntactic and semantic features of a sentence, to determine which of these

¹⁴ Zeno Vendler, "Verbs and Time," *Philosophical Review* 56 (1957) p. 97-121.

¹⁵ There are other possible typologies than the one I am going to use here. Brinton, *The Development of English Aspectual Systems*, *op. cit.* p. 33-36 canvasses three others and compares these usefully to the ones derived from Kenny and Vendler. See Zelda Boyd and Julian Boyd, "To Lose the Name of Action: The Semantics of Action and Motion in Tennyson's Poetry," *PTL: a Journal for Descriptive Poetics and the Theory of Literature* 2 (1977), Ronald Taylor, "The Aspectual Structure of the English Sentence," *Doshisha University English and English Literature Research* (1977), and David Dowty, *Word Meaning and Montague Grammar: The Semantics of Verbs and Times in Generative Semantics and in Montague's Ptq* (Norwell, MA: Kluwer Academic Publishers, 1979), 180-6. Brinton also discusses and refines the refinements suggested by Bernard Comrie, *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems* (Cambridge: Cambridge University Press, 1976), ch. 1. The optimal typology will to some extent depend on the purposes for which it is designed. All of these typologies make a distinction between statives and non-statives, which is all that really matters for my argument below. I have adopted the Vendler/Kenny approach because it is most familiar to philosophers and because it makes the stative vs. non-stative distinction at the highest-level.

¹⁶ Alexander Mourelatos, "Events, Processes, and States," *Linguistics and Philosophy* 2 (1978), p. 423.

¹⁷ Mourelatos's term 'process' is potentially the most misleading insofar as it suggests a sequence of discrete steps, which he later recognizes; see Alexander P. D. Mourelatos, "Aristotle's Kinêsis/Energeia Distinction: A Marginal Note on Kathleen Gill's Paper," *Canadian Journal of Philosophy* 23, no. 3 (1993).

¹⁸ Compare Comrie, *Aspect*, *op. cit.* ch. 2 for more extensive discussion of each of these categories.

features are possessed by the relevant verb-phrase.

This means that we can devise linguistic tests to aspectually classify arbitrary sentences of our language (at least the logically atomic ones) as descriptions of states, processes, developments, or punctual occurrences.¹⁹ The way this works is made clearer by working through some tests applied to specific examples. Considering these will also provide several points where we can call into question Sosa's claim that belief and knowledge are performances. (My basic contention will be that descriptions of paradigmatic performances always fall under a different aspectual category from attributions of belief and knowledge.)

According to the typology illustrated above, states are thought to differ from occurrences in that they are non-dynamic. And we can test for whether a verb phrase refers to something dynamic by considering the semantic contrast between simple and progressive uses of the verb. Some of the literature on aspectual classification suggests that the mere availability of the progressive (-*ing*) form in English is a marker for occurrences as opposed to states. This is encouraged by the fact that the progressive is available in sentences such as

- (1) a. He is jogging. (process)
- b. He is building the house. (development)
- c. He is winning the race. (punctual occurrence)

which convey something happening or occurring, i.e. something *dynamic*, while it is not available in paradigmatic state-descriptions, such as

- (2) a. He is owning the house. (?)
- b. He is desiring something to drink. (?)
- c. He is being tall. (?)

However, although the unavailability of the progressive seems to be a sufficient condition for a verb-phrase to be non-dynamic²⁰ and so aspectually classified as a state description, we should recognize that there are intuitive state-descriptions deploying the English progressive, e.g.

- (3) a. Currently, we are living in Sydney.
- b. The statue is standing in the parking lot.
- c. A photograph of Grandma is sitting on the mantelpiece.

The traditional explanation of state-descriptions deploying the English progressive is that these sentences refer to states that considered somehow temporary. Goldsmith and Woisetschlaeger argue (convincingly in my view) that a deeper analysis is needed.²¹ But whatever analysis that

¹⁹ See Dowty, *Word Meaning and Montague Grammar*, *op. cit.*, pp. 51-55 for useful discussion of this typology, its history, and a compendium of many more linguistic tests than I will discuss here. Anita Mittwoch, "Aspects of English Aspect: On the Interaction of Perfect, Progressive and Durational Phrases," *Linguistics and Philosophy* 11, no. 2 (1988) contains helpful criticism and refinement of several of these tests involving durational phrases and the adverb 'already'. Also, Jan Van Voorst, "The Aspectual Semantics of Psychological Verbs," *Linguistics and Philosophy* 15 (1992), contains a helpful discussion in a linguistic context of specifically psychological verbs such as "to believe".

²⁰ All significant proposals for articulating the semantics of the progressive in English predict this. See especially, Terence Parsons, "The Progressive in English: Events, States and Processes," *Linguistics and Philosophy* 12, no. 2 (1989), Fred Landman, "The Progressive," *Natural Language Semantics* 1, no. 1 (1992), and James Higginbotham, "The English Progressive," in *The Syntax of Time*, ed. Jacqueline Gueron and Jacqueline Lecarme (Cambridge, MA: The MIT Press, 2004).

²¹ John Goldsmith and Erich Woisetschlaeger, "The Logic of the English Progressive," *Linguistic Inquiry* 13, no. 1 (1982) p. 84-85. Sentences (3b) and (3c) are adapted from their paper. Their argument consists of cases where such sentences are expanded with clauses indicating that the state is not temporary. Their deeper analysis has to do with the distinction between features of a situation presented as

is, we already have some linguistic evidence against Sosa's idea that beliefs are like archery shots in being performances. For contrast

- (4) a. He is shooting the arrow with skill.
b. It is hitting the target because of his skill.
- (5) a. He is believing that p. (?)
b. He is knowing that p. (?)

Sentences with the form of (5a) and (5b) are generally anomalous. There *may* be contexts where one could say, e.g. 'He is believing that the story is true, but as soon as Sally appears he'll know it's a lie.' However, if these are legitimate sentences of English, they clearly occur in contexts where the belief is temporary or otherwise not part of the structure of the situation. This means that they are like the sentences in (3), and we should not let marginal uses of the progressive 'is believing' dissuade us from the view that belief-attributions are state-descriptions. Moreover, it does not seem to me that (5b) is ever a grammatical sentence of English (indeed linguists often use knowledge-attributions as a paradigm example of a state-description). This means that it is correct to aspectually classify belief- and knowledge-attributions as statives. Importantly, however, that contrasts with the correct aspectual classification of paradigmatic descriptions of performances and achievements such as those in (4). This suggests strongly *pace* Sosa that it is part of the meaning of belief- and knowledge-attributions that they are about states rather than performances.²²

This result can be confirmed by a related test for whether a verb-phrase refers to something conceived of as dynamic or involving change. With non-statives, we typically use the simple present form of the relevant verb only if some form of repetitiveness, seriality, or habitualness is suggested. For example, contrast

- (6) a. He is running.
b. He is writing a book.
c. He is (just now) winning the race.
- (7) a. He runs.
b. He writes a book.
c. He wins the race.

The sentences in (6) are naturally interpreted as referring to a single occurrence, whereas the sentences in (7) are typically used in contexts suggesting repetitiveness, seriality, or habitualness, such as: 'He runs *every day*', 'He writes a book *whenever he has a sabbatical*', or 'He wins the race every year.' The opposite is true for statives:

- (8) a. We are living in Sydney.
b. The statue is standing in the parking lot.
c. A photograph of Grandma is sitting on the mantelpiece.
- (9) a. We live in Sydney.
b. The statue stands in the parking lot.
c. A photograph of Grandma sits on the mantelpiece.

"structural", in the sense that they structure experience of the situation, as opposed to features presented as "phenomenal" in that they are experienced.

²² Could a sentence deploying a stative verb-phrase be about the same thing as another sentence deploying a non-stative verb-phrase? Perhaps. This is a difficult question at the interface of metaphysics and the philosophy of language, which I don't hope to settle here. But two relevant points: First, there doesn't seem to be intuitive alternatives for attributing belief and knowledge that deploy non-stative verb-phrases. Second, the important issue here is the implications of the conception of belief and knowledge reflected in the meaning of the sentences we use to talk about these things. As I'll point out below (footnote 28), one could engage in error-theory with respect to these concepts, but that does not seem to be the way Sosa or other epistemologists want to go.

To be sure, the sentences in (9) *can* be expanded to convey some form of repetitiveness, seriality, or habitualness but they do not sound strange absent such context either.

Now consider the sentences which should be analogous from the point of view of Sosa's virtue epistemology:

- (10) a. He shoots the arrow.
b. It hits the target (because of his skill).
- (11) a. He believes that p.
b. He knows that p.

Sentences (10a) and (10b) are analogous to sentences (7b) and (7c). To describe a one off occurrence, one would not normally use these sentences, which are in the simple present, but another form of words, such as 'He is shooting the arrow' and 'It is hitting the target (at this precise moment)'.²³ The same is not true of sentences (11a) and (11b). These are instead analogous to sentences the sentences in (9). So, again, we get the result that paradigmatic descriptions of performances refer to something dynamic and so should be aspectually classified with non-statives, whereas belief and knowledge attributions refer to something non-dynamic and so should *pace* Sosa be aspectually classified with statives.

The basic result is that belief and knowledge attributions seem, by virtue of their meaning, to be about something non-dynamic, whereas paradigmatic performance-descriptions (e.g. of arrow-shootings) seem to be about something dynamic and so non-stative. I think this shows that Sosa's suggestion that belief is a performance – which when successful (true) because skillful (justified) is apt and so a kind of knowledge – to involve him in a sort of metaphysical category mistake in the way he uses these words. Below I'll consider some responses to this objection, but it is both interesting and relevant for what follows to first consider how to distinguish among the three different kinds of non-statives.

In the Kenny/Vendler/Mourelatos typology, tests for telicity and duration distinguish among the three kinds of non-statives – processes, developments, and punctual occurrences. That is, we can separate three categories based on two features:

- a. processes (e.g. running): durative and atelic
- b. developments (e.g. writing a book): durative and telic
- c. punctual occurrences (e.g. winning the race): non-durative and telic

The basic thought is that a verb-phrase like 'is running' allows us to talk about something that takes time but does not imply a specific endpoint, whereas a verb-phrase like 'is writing a book' allows us to talk about something that takes time and implies a specific endpoint, and a verb-phrase like 'is winning the race' allows us to talk about the point in time at which some implied endpoint is reached. Of course, we may say 'He is running' when we know that he plans to run precisely a mile, but the endpoint is not incorporated into the meaning of the verb-phrase we use. Similarly, we may say 'He is winning the race' when we know that winning the race takes a small but not infinitesimal amount of time, but, as far as the meaning of the phrase goes, the occurrence is conceived of as point-like.

Several linguistic tests have been proposed for determining whether a verb-phrase is durative/non-durative, telic/atelic. However, the one that is perhaps most widely known for

²³ Exception: we sometimes use the simple present to refer to one-off occurrences when we use a rhetorical device called the "historical present." This involves projecting our narrative into the present moment of some unfolding series of events, e.g. "He dribbles, he shoots, he scores!" But I think we can recognize the historical present when we see it and modulate for it. In any case, most of the belief and knowledge attributions that are at issue in epistemology are obviously not in the historical present.

English has to do with the interaction between the past progressive and the perfect tenses. For some verb-phrases, the past progressive implies the perfect, but for others it does not. For example, notice the following contrasts:

- (12) a. He was running (then), so he has run (at some point in the past).
b. He was writing a book (then), so he has written a book (at some point in the past).
(?)
c. He was winning the race (then), so he has won the race (at some point in the past).

The reason (12a) sounds fine is that as soon as someone is engaging in the activity of running, it is true to say of him that he has engaged in that activity. The phrase 'is running' refers then to something durative but atelic. By contrast, (12b) sounds odd because someone who was writing a book could be interrupted and never finish. Hence, the phrase 'is writing a book' refers to something durative and telic. The sentence (12c) is more difficult to adjudicate (and doing so not crucial to for my purposes here). The thought behind classifying this differently from (12a) and (12b) is that although there is an implied endpoint to 'is winning the race,' this phrase refers not to the progression towards that endpoint but only the point of achieving it.

I have already explained why belief- and knowledge-attributions should be aspectually classified as state-descriptions, while descriptions of paradigmatic performances should be classified in the opposite node of the typology as occurrence-descriptions. But it is interesting to ask: where, as we move lower in the Kenny/Vendler/Mourelatos typology, should we classify descriptions of paradigmatic performances? Consider the following:

- (13) a. He was shooting the arrow (then), so he has shot the arrow (at some point in the past).
b. He was skillfully hitting the target (then), so he has hit the target with skill (at some point in the past).

Linguistic intuitions about these sentences might reasonably differ, but it seems to me that (13a) is anomalous like (12b). Shooting an arrow doesn't take long, but it could surely be interrupted, such that one could be in the midst of shooting an arrow but not yet have shot the arrow. If that is right, it indicates that the phrase 'is shooting the arrow' refers to a development (and partially vindicates Kenny's original term "performances" for this node of the typology). By contrast, it seems to me that (13b) is like (12c). It is not completely clear whether it makes sense, but if it does that is because hitting the target with skill refers not to the progression towards the implied endpoint but the point of achieving it. If this is right, it indicates that the phrase 'is skillfully hitting the target' refers to a punctual occurrence (and partially vindicates Vendler original term "achievement" for this node in the typology).

IV

Before moving on to propose an alternative framework for thinking about the axiological and deontological questions in epistemology, I want to consider three responses to my argument so far. Perhaps the most obvious response is to concede that belief and knowledge are not performances but then to argue that the *formation* of beliefs is a performance.²⁴ This is, of course, right, but I think it does not rescue Sosa's view. This is for two independent reasons.

²⁴ A similar but ultimately less plausible suggestion is that the performance relevant for knowledge is not belief but *judgment*. I believe this notion is a bit of a mongrel in contemporary philosophy. On some of its uses, it may be a performance, but it will have many of the same problems in the present dialectical context as I go on to point out for appealing to belief formation as the relevant performance. On other common uses in philosophy, judgment is not plausibly conceived as a performance but rather as an activity or a state.

First, as we have seen, one's knowledge can be tacit in the strong sense that one has never consciously considered the proposition known to be true (e.g. my mother's knowledge that the rain in Spain is not purple). Although we may want to say that the relevant beliefs are *tacit* as well, it is considerably strained to think of their subjects as having skillfully formed them. But unless we ignore this strain, we cannot think of such tacit knowledge as a state that inherits its value from the performance of forming the relevant beliefs. Moreover, we cannot make sense of the deontological claims that one ought or ought not have such tacit beliefs. This means that, even if we focus on the formation of beliefs rather than the beliefs themselves, Sosa's answer to the axiological and deontological questions depends on the dubious idea that subjects skillfully form all of the tacit beliefs constituting their tacit knowledge.²⁵

Second, even if we accept the assumption that a state inherits the value attaching to its formation (which is dubious as a general claim but perhaps not so in the particular case of belief formation), we should recognize that sometimes one forms a true belief skillfully (e.g. by a reliable process or based on reasons), but then maintains this belief in a way that undermines its status as knowledge. This can happen when, e.g., one forgets the original reasons for believing that *p* and is faced with strong countervailing but misleading evidence and yet maintains the belief. In such a case, it might be true that the belief is true because formed for good reasons (i.e. that the belief formation is successful because skillful) but it would still be wrong to say that the belief is knowledge, since the subject has maintained a belief in the face of strong countervailing evidence without any idea of why it is true. In this sort of case, the subject's belief would have whatever value can be inherited from its formation; however, it would not be as good as knowledge, which means that value inherited in this way cannot answer the axiological question.

I wanted to respond at length to the first objection to my argument not only because it is the most natural for a proponent of virtue epistemology to make but also because I think it highlights an important aspect of the normative dimension of epistemology and part of why it is so natural to use a stative verb-phrase to refer to knowledge. Metaphorically speaking, what is epistemically good about a belief that amounts to knowledge is, at least in part, a factor of how it is maintained "in waiting" across time. This is due to the durative and non-changing aspect of states; performances simply do not have this. I will come back to this below in sketching my alternative conception of epistemic normativity. But in the meantime, let me anticipate how it might be thought to point the way to a further objection to my argument from the previous section.

One might suggest that the performance relevant to the distinctive value of knowledge is not forming a belief but *maintaining it*. For example, in response to the objection that belief does not require conscious activity like most paradigmatic performances, Sosa encourages us to consider "those live motionless statues that one sees at tourist sites." He takes these to provide a nice analogy for his view, since, as he puts it, "Such performances can linger, and need not be constantly sustained through renewed conscious intentions. The performer's mind could wander, with little effect on the continuation or quality of the performance."²⁶ Could, perhaps, focusing on belief-maintenance regain the structures of "performance normativity" for a virtue

²⁵ In Sosa, "Value Matters in Epistemology", *op. cit.*, p. 168, he explicitly restricts himself to cases where the subject consciously considers the proposition; but he does so without comment or justification. Tacit knowledge is better than tacit belief falling short of knowledge. So, if Sosa's account does not extend to this, then it is not a satisfactory answer to the axiological question.

²⁶ Sosa, *A Virtue Epistemology*, p. 23.

epistemological answer the axiological and deontological questions?

I think there is something deeply right about focusing on the maintenance of belief, rather than the formation of belief, when it comes to answering these questions. However, the problem with this as a response to my argument from above is that maintaining a belief is not a performance. At least, the verb-phrase 'is maintaining a belief' would be aspectually classified with verb-phrases used to talk about processes (because it's atelic like 'is running'), and – as we saw above – all paradigmatic performances are referred to by verb-phrases aspectually classified with event-descriptions (which are telic like 'is writing a book').

What about those live motionless statues? Their performance, in my view, is maintaining a pose for some period of time. It requires no great skill to stand still for 10 seconds, but doing so for 10 minutes might be a performance worth a tip. However, one can maintain a belief for 10 seconds or 10 minutes or for any length of time, and the belief can count as knowledge, no matter how long one maintains it. So, if our ordinary ways of talking about these things are indicative or our ordinary conception of these things, which I think it they are, then, on our ordinary conception of maintaining a belief, this is a process (Vendler's term: activity) rather than an event (Kenny's term: performance). But, since we use event-descriptions rather than process-descriptions to talk about all paradigmatic performances, maintaining a belief is not plausibly thought of as a performance. Again this means that the structures of performance normativity do not apply to it. This is not just an accident of grammar. The telic aspect of performances seems to be part of why we can distinguish along the three dimensions of performance normativity: success, skill, and aptness. Hence, something atelic, such as maintaining a belief, will never be the right sort of thing to which to apply these evaluative distinctions.

Finally, one might grant that neither beliefs nor maintaining beliefs are performances but argue that Sosa's theoretical structure of "performance normativity" applies more generally than just to performances. The idea would be to claim that performances are the easiest example to appreciate, but there are other things besides performances that admit of the distinction between success, skill, and success because of skill. Perhaps anything with an aim could be said to admit of this distinction by analogical extension. Then, even if beliefs are states rather than performances, as the aspectual data seems to indicate, it may still be true that beliefs, because they "aim at the truth" admit of the normative distinctions characteristic of performances. And when a belief successfully reaches its "aim" because of the skill of the believer, it is knowledge.

The problem I have with this response is that it strikes me as completely *ad hoc*. Unless there are other states that plausibly admit of the distinctions characteristic of performance normativity, it will be entirely theory-driven to think that the states of belief and knowledge fit into this normative structure. Moreover, beliefs aim at the truth only in a highly metaphorical sense.²⁷ This means, I suppose, that we could say that beliefs and knowledge are performances in an equally highly metaphorical sense. But if we do so, I would like to stress that this does not explain the distinctive value of knowledge and the practice of deontologically evaluating beliefs any more than Shakespeare's idea that Juliet is the sun explains her beauty.

Let me grant that there might be some other way to gerrymander the conceptual space here, so that knowledge can be seen as a state one cannot get into or be in without engaging in a

²⁷ Some philosophers suggest that desires aim at the good like belief aims at the truth. Even if this highly controversial claim is true, however, notice how strange it would be to evaluate desires for their success, skill, and aptness vis-à-vis this putative aim.

special kind of performance. I have thought of several other ways that strike me as so forced that they are not even worth mentioning here, but in all cases one gets the sense that the virtue epistemologists who favor this general framework for answering the axiological and deontological questions in epistemology are trying to force the round peg of knowledge into the triangular hole of performance normativity. In the following section, I seek to motivate an alternative framework.

V

In light of the preceding discussion, some virtue epistemologists might want to insist that it does not matter what is implied by the meanings of ordinary belief- and knowledge-attributions, because ordinary language is notoriously vague, imprecise, and sometimes even confused. However, I take it to be a desideratum on an account of the normative evaluation of belief that it is at least consistent with the aspectual data rehearsed above. Otherwise, we run a serious risk of simply changing the topic²⁸ rather than providing illuminating answers to the axiological and deontological questions that I take to be difficult open questions in contemporary epistemology and which seem, in different ways, to motivate the search for better general frameworks for epistemology than Cartesianism and reliabilism.

The reason I wanted to rehearse the aspectual data was not just to bolster the piece of philosophical commonsense that belief and knowledge are states in the face of the temptation to ignore it when considering the axiological and deontological questions. I also think the aspectual data provide an important clue to answering these questions. I shall not venture full answers to these questions in what follows because I think we do not yet understand well enough the kinds of norms that apply to beliefs. However, I will sketch a framework that is motivated by the aspectual data and capable of avoiding the extremes of Cartesianism, reliabilism, and virtue epistemology.

The key to this alternative framework is to recognize the difference between *action-norms* and *state-norms*.²⁹ We are familiar from ordinary ethical life with the sorts of norms that govern our actions: e.g., one should not steal, one should pay back debts, etc. And ethical theorists often seek to articulate more fundamental normative principles governing action in general: e.g., one ought never to act except in such a way that one could will one's maxim as universal law, or one ought to act in such a way that one maximizes overall utility. However, it takes only a moment's reflection to observe that there are also ethical norms pertaining to states: e.g. playgrounds should be free of drugs, hospitals should be accessible to the infirm. These are norms pertaining to states outside of the agents who are assumed to be (at least partially) responsible for the states. Importantly, there are also ethical norms pertaining to states of the agents who are assumed to be (at least partially) responsible for the states: e.g. one should feel guilt for

²⁸ I do not mean to suggest that it could not make sense for certain purposes to change the topic. For example, maybe certain theoretical purposes in cognitive science or empirical psychology would be advanced by abandoning talk of belief and knowledge in favor of less quotidian notions. I take it, however, this is neither what Sosa is suggesting nor a move that many epistemologists would find useful for *their* theoretical purposes.

²⁹ Compare Wilfrid Sellars's distinction between "rules of action" and "rules of criticism" in "Language as Thought and as Communication," *Philosophy and Phenomenological Research* 29, no. 4 (1969), as well as I. L. Humberstone's "Two Sorts of 'Ought's,'" *Analysis* 32, no. 1 (1971). My idea of applying this distinction to belief and knowledge here is similar to David Owens's idea in writing "Some norms are not there to guide action, to govern the exercise of control: their function is to assess what we are," *Reason without Freedom* (London: Routledge, 2000), p. 126.

wronging another, one should want one's children to be happy, etc.

What is illustrative about both of these kinds of state-norms is that, unlike action-norms, their validity does not rest on the assumption that their explicit subject can choose whether to conform to the norm. To take the most extreme example, playgrounds and hospitals do not choose anything, because they are inanimate objects. So it cannot be the case that playgrounds can choose whether they are free of drugs or hospitals can choose how accessible they are. Likewise, however, it seems to me that individual people do not choose, in particular cases, whether to feel guilt for wronging another or choose to want their children to be happy. The validity of state norms would, I think, come into question if there were no discernable connection between things that people can choose to do and conformity to state-norms; however, the connection can be indirect, complicated, and perhaps even uncodifiable.

For example, when we say that a hospital should be accessible to the infirm, we imply (defeasibly) that the hospital's architect should design it that way and/or the builders should use techniques ensuring accessibility and/or the manager should correct any problems with the accessibility, etc. These are things that agents can choose to do, in order to put the hospital in the state it should be in. Because of this, we can say that the connection between this state-norm and action-norms is *other-regarding*. However, when the relevant state is the state of an agent, it seems that we can also have a *self-regarding* connection. When we say, for example, that a person should feel guilt for wronging another, we often imply (defeasibly) that she should perform certain actions, such as reflecting on the pain she has caused, considering the plight of the victim, discussing her wrongdoing with others, seeking psychological help if she does not feel guilty, etc. In this way, state-norms pertaining to the states of an agent can be different from the state-norms with only an other-regarding connection to action norms, since they can imply things about what the explicit subject of the state should do.

However, I think we miss an important part of the picture if we ignore the fact that state-norms pertaining to an agent can also imply things about what others in the agent's community should do. For example, when we say that a person should feel guilt for wronging another, we can also imply that the person's friends should help her to see the plight of the victim and/or the person's parents should have inculcated this emotional capacity, etc. In this way, state-norms with a self-regarding connection to action-norms can also have an other-regarding connection.

It may be difficult to map out in any general way when and how a state norm is other-regarding, self-regarding, or both. However, once we recognize the distinction between action-norms and state-norms, it becomes easy to respect the idea that that belief is a state rather than a performance, without giving up on the idea that beliefs are subject to epistemic norms and so believers are proper objects of deontological evaluation. We must simply conceive of the epistemic norms governing belief as state-norms rather than action-norms. As such, they do not presuppose that conforming to them is something that their explicit subjects can choose to do.³⁰ However, their validity does – it seems to me – rest on indirect, complicated, and perhaps even uncodifiable connections to things that people can choose to do. Insofar as they have a self-regarding connection to action-norms, when we say that someone ought to believe some proposition, we often imply (defeasibly) that this person should perform certain actions, such as looking up the answer to a question and/or reflecting on the consequence of firmly held beliefs

³⁰ See also Matthew Chrisman "Ought to Believe" *Journal of Philosophy*, vol. CV, no. 7, (2008): 346-370, where I use this idea to address the tension between doxastic involuntarists who argue that beliefs are not subject to deontological evaluation because they are not voluntary and epistemic deontologists who argue that beliefs must be free because they are subject to deontological evaluation.

and/or reconsidering the evidence, etc. Broadly speaking, these are the actions that constitute the practice of inquiry. And epistemic norms pertaining to beliefs, considered as state-norms are often connected inferentially to norms pertaining to the actions of inquiry, considered as action-norms.

Again, however, I think we miss an important part of the picture if we ignore the fact that epistemic norms pertaining to beliefs can also imply things about what others in the subject's community should do. Indeed, this must come first in the development of an epistemic subject. There are things that small children ought to believe, but because they are too small to be responsible for these states, it would be wrong to think that these norms imply that there are actions of inquiry that the child ought to perform. But even when someone is old enough to be a mature epistemic subject, and we say that she should believe some proposition, we often imply (defeasibly) that her parents should have taught her the truth of this proposition and/or her friends should correct her if she believes the opposite and/or her colleagues should provide her with evidence of the truth of this proposition, etc. Broadly speaking, these are the actions that constitute the practice of instruction. My suggestion is that epistemic norms pertaining to beliefs, considered as state-norms, bear an other-regarding connection to the norms pertaining to the actions of instruction. This is in addition to whatever self-regarding connections they bear to norms pertaining to the actions of inquiry. (This is a further reason why I think it is wrong to think the activity of maintaining a belief is the sole source of the epistemic norms governing belief.)

So, as I am thinking of them, beliefs are states of mind subject to epistemic state-norms, which have both a self- and an other-regarding connection to action-norms. I think this idea explains why we use a form of words aspectually classified as stative to talk about the object of central epistemic norms: the norms of belief. Moreover, it makes it at least intelligible that we would deontologically evaluate subjects for the belief states they are in. And this is how my framework offers the beginning of progress on the deontological question from above.

Moreover, in my view, when we attribute knowledge to someone, we evaluate this person's belief state as conforming to such norms about what this person ought to believe. These norms are inferentially related to the epistemologically relevant action-norms pertaining to inquiry and/or instruction, but it would be a mistake to conflate epistemic state-norms with those action-norms. For there may not often be an identifiable connection between an epistemic state-norm and a set of specific actions that exhaust its normative force. That is to say that there may be no general or straightforward way to link state-norms pertaining to beliefs to specific action-norms pertaining to inquiry and instruction, but the validity of each specific application may still depend on the validity of a corresponding set of action-norms. It is this connection that I think we should investigate further in seeking a better answer to the axiological question from above. Understanding the specific kind of 'ought to believe', conformity to which tracks with knowledge, will require understanding the various self- and other-regarding connections between epistemological norms of belief, considered as state-norms, and norms of inquiry and instruction, considered as action-norms.

If we think of belief and knowledge in this way, we avoid the pitfalls of Cartesian views, which can now be seen to have incorrectly assumed that belief is itself subject to epistemic action-norms. In my framework, belief is subject only to epistemic state-norms, which bear some inferential connection to the action-norms pertaining to inquiry and instruction. Thus, knowledge is not helpfully viewed as well-chosen belief, although when someone has knowledge his belief conforms to epistemic norms that involve an inferential connection to choices people should make in engaging in the actions of inquiry and instruction.

This idea also helps us to avoid the impoverishment of reliabilist conceptions of belief and knowledge. Like digestive and immune systems, which can be just as properly subject to state-norms as playgrounds and hospitals, believers are subject to state-norms with other-regarding connections to action-norms about what others should do. However, reliabilist views can now be seen to have assumed that this way of being subject to norms is enough to capture the normative dimensions of epistemology. However, in my view, beliefs are states of a potential inquirer, and, because of this, beliefs are subject to state-norms that may have self-regarding implications about which actions of inquiry their possessors should perform.

The more general claim I want to make is that understanding the distinction between action-norms and state-norms can provide the template for an alternative route between the pitfalls of Cartesianism and the impoverishment of reliabilism, without succumbing to the temptation to treat belief and knowledge as kinds of performances, as Sosa does. Moreover, this sort of account fits well with the observation that belief and knowledge are two of the most important cognitive states for us to keep track of in order to understand and cooperate with others in a common environment. Each of us has only a largely incomplete set of information about the features of our physical and social environment, but luckily each of us also has a different set of information. This means that we can, in principle share information with one another, and broaden the information on the basis of which we make particular decisions about how to behave. But to do that, we need some way of keeping track of who has which pieces of information, and, additionally, of keeping track of whose informational states contain systematic or idiosyncratic errors.³¹

In virtue of being state-descriptions, our ordinary belief- and knowledge-attributions are well suited for this task. When I do not know, e.g., where the food is, it will surely be a step in the right direction if I know that you do know where the food is; then, I only need to convince you to share that information with me. And even when I do know, e.g., that the water is over the hill, if I am looking for you to tell me where the food is and I think you may be fetching water, then I can form a strategy for finding you: look over the hill. Moreover, even if I do not know whether you know, e.g., which plants are poisonous, if I know that your beliefs about such matters are unreliable, then I can form a better strategy for finding out which plants to eat: ask someone else who has more reliable beliefs about such matters.

These are simplistic ecological observations, but I think they illustrate how having concepts of belief and knowledge conceived as states would serve us in surviving and cooperating in a common environment in a condition of less than full information. Moreover, if my discussion above is moving in the right direction, then belief and knowledge are both cognitive states, but, in the environment in which it is useful to keep track of these things, it is plausible to think that they play different roles. As we have already seen, if I am looking for the food, then, if you have knowledge about where the food is, you will prove more useful than someone who merely has a belief about where the food is. To be sure, as Socrates already pointed out in the *Meno*, assuming that both you and the other person have correct beliefs about such matters, the instrumental value of finding out your beliefs may prove identical in the specific case. However, we are often interested in keeping track of each other's cognitive states not only for a present

³¹ This idea is inspired by Edward Craig, *Knowledge and the State of Nature: An Essay in Conceptual Synthesis* (Oxford: Clarendon Press, 1990), who uses a "state of nature" argument to recommend viewing the proto-concept of knowledge as one used to "tag reliable informants". However, it is unclear to me whether the status conferred by being tagged in this way must be considered central to the concept of knowledge – either as it was used in some hypothetical state of nature or as we use it today – opposed to a peripheral upshot of keeping track of who is in which informational states.

purpose, but also for purposes of unspecified future exchanges of information. Hence, even if you and someone else agree in your opinions, e.g., that some specific plant is poisonous, it will prove useful to me to be able to keep track of the fact that you merely believe this and the other person knows this. For in the future, it will probably be a better strategy to seek information about other plants from the other person instead of you. Alternatively, even if I think that both of your opinions are false, if I think one person's opinion amounts to justified belief, whereas the other's does not, then I'll have some idea of who to go to first in the future and also some idea of whose practices of inquiry to aim to emulate when it comes to such matters.

In light of these ideas, the answer to the axiological question in epistemology does not seem to me to lie in any special idea of "performance normativity". Rather the distinctive value of knowledge will depend precisely on the fact that it is not a performance. As a cognitive state, it is the sort of thing we can count on being available to depend upon, even as other things change. Of course, belief falling short of knowledge is also a state, and we depend on such states when we have to. But because the long term functioning of our information economy depends on the difference between knowledge and other sorts of belief, we depend on them in different ways. This should then be reflected in the connections between epistemological norms of belief and the action-norms governing actions constituting the practices inquiry and instruction. Hence, I think it will be something about this difference which explains why we value knowledge more than beliefs falling short of knowledge. And this is how my framework offers the beginning of progress on the axiological question from above.