

The State of Knowing

Alan Millar

1. *The plan*

It has been widely assumed over the last fifty years or so that epistemology, so far as it addresses the nature of knowledge, should issue in a conceptual analysis of *knowledge-that*—a statement of necessary and sufficient conditions for the truth of instances of the schema *X knows that p* in terms that do not implicate the concept of knowledge. The traditional way of going about this—the *traditional project*—is to look for a reductive analysis—a set of conditions, including truth, that need to be satisfied if belief is to be turned into knowledge. Some these days are, in my view rightly, sceptical that this is the right way to go about epistemology.¹ Attempts to spell out the further conditions have fallen prey to Gettier counter-examples, and there is an issue as to why we should even expect the concept of knowledge to yield to a traditional analysis. There seem to be relatively few concepts that do. Why suppose that *knowledge* is among them? But what if knowledge is not belief plus the satisfaction of further conditions and what if it does not yield to any reductive analysis? What in that case can we say in general terms about the nature of knowledge? On the view outlined here knowledge that something is so is essentially a state that results from the deployment of a way of telling that such things are so. A way of telling just is a way of acquiring knowledge. So if this approach is to be illuminating it will be not be through the provision of a reductive analysis. Rather it will aim to shed light on the nature of knowledge by describing the variety of ways of telling and exploring the implications for justified belief, the role of

¹ See, especially, Williamson 2000. My debt to his work will be clear. I do not suggest that he would endorse my conception of ways of telling and their role in illuminating what knowledge is. My thinking in this area has also been significantly shaped by reflection on John McDowell's work in epistemology. See especially his 1982.

reliable belief-forming processes, sceptical arguments, and so on. In this essay, especially in section 5, I aim to make some headway on the key issue of how we should conceive of ways of telling.

A theme of the essay is that an adequate epistemology, and therefore an adequate account of ways of telling, should emerge from conceptual enquiry. In sections 2 and 3 I address a challenge to this view that directs us to look to the phenomenon of knowledge rather than the concept. I argue that it is a mistake to draw a sharp contrast between enquiry into concepts and enquiry into phenomena. One upshot of the discussion is that conceptual enquiry should not be identified with conceptual analysis in the style that yields reductions of key concepts.

In section 6 I consider why knowledge matters against the background of the conception sketched in section 5. Knowledge is something that we can often readily acquire, often want, and often transmit to others. It makes a difference to how people act whether they are acting from knowledge or from some lesser state. Traditional approaches to epistemology make this hard to see because they make it hard to understand how facts about knowledge can be as readily available as they often are. I think that the conception of ways of telling sketched here is better able to make sense of the importance we place upon knowledge in our social interactions and engagements with the world.

2. Epistemology, natural kinds and conceptual enquiry

Recently Hilary Kornblith has played down the role of conceptual enquiry in epistemology. Kornblith holds that ‘the subject-matter of epistemology is knowledge itself, not our concept of knowledge’ (2002: 1). Knowledge, he says, is a natural phenomenon; indeed, it is a natural kind. As such it is to be investigated like other natural phenomena. This view naturally gives rise to a number of questions.

- (1) What is the import of the claim that knowledge is a natural kind?
- (2) If knowledge is a natural kind does this lend support the claim that the subject matter of epistemology is not the concept of knowledge?

According to a standard way of thinking, natural kinds include substances like water and gold. Our concept of water is of any liquid that is same liquid as the familiar that we pick out as being water. Sameness of liquid is determined by sameness of nature, and the nature of water is to be H₂O. There are two important implications of such a view that are germane to our present concerns. The first is that the nature of water is not wholly captured by the concept. We can have a perfectly adequate grasp of the concept without, just on that account, being committed to thinking of water as H₂O. It follows that reflection on the concept of water will not reveal the nature of water. The second point is that to discover the nature of water we have to conduct an empirical investigation. It was just such investigation that revealed that water is H₂O.

Against this background at least part of the answer to question (1) above is that the concept of knowledge does not capture the nature of knowledge and so reflection on the concept of knowledge will not reveal the nature of knowledge. This in turn yields an answer to question (2): the subject-matter of epistemology is knowledge itself and that is revealed by empirical enquiry, not by reflection on the concept.

Even supposing that the story just told about water is correct it does not follow that a similar story can be told for all concepts. Consider for instance the concept of furniture. Our concept of furniture is of things that include tables, chairs, chests of drawers, bookcases, and similar things. It seems that there is no more to being furniture than is captured by the concept—to be an item of furniture is just to be suitably similar to things that include tables, chairs, ... , and so on. Similar considerations apply to role-concepts like the concept of a

teacher. The concept of a teacher is just the concept of one who teaches. Arguably, there is no more to the state of being a teacher than what a judgement to the effect that someone is a teacher would specify in virtue of its implicating the concept of a teacher.

Is there any reason to think that the concept and the state of knowledge should be understood as being analogous to the concept of water and the substance water? It is not obvious that we should think of the concept in this way. Perhaps our concept of knowledge is akin to the concept of furniture. Perhaps knowledge comprises states like seeing or hearing that something is so and remembering that something is so, and states that are significantly similar to these. Or perhaps we have no single determinate concept of knowledge.

Kornblith suggests that regarding knowledge as a natural kind best accommodates the fruitful practice among cognitive ethologists of treating knowledge as ‘a well-behaved category, a category that features prominently in causal explanations, and thus in successful inductive predications’ (2002: 62). A category is well-behaved in the relevant sense when that to which it applies instantiates clusters of properties that are stable, feature in causal laws, and can form the basis for reliable predictions about the behaviour of what has those properties. Amplifying on this Kornblith says that ‘[i]f we wish to explain why it is that members of a species have survived, we need to appeal to the causal role of the animals’ knowledge of their environment in producing behaviour which allows them to succeed in fulfilling their biological needs’ (2002: 62).

I have no problem with the idea that knowledge plays a causal-explanatory role. But what is it for a *state* like knowing to be a natural kind? The central cases of natural kinds considered in the literature are substances (understood as stuff) and individuals. Samples of the same substance-kind (e.g., gold) and individuals of the same individual-kind (e.g. apples) share certain properties possession of which makes these samples and individuals belong to the same kind. In the case of substance-kinds instances of the kind are samples of the substance. Instances of individual-kinds are obviously individuals belonging to the kind. If

knowledge is a natural kind what are its instances? Kornblith takes knowledge to be a kind to which beliefs belong. On his way of thinking beliefs belong to the kind only if produced by processes that reliably yield true beliefs (2002: 61-63). The properties that go to make up the nature of knowledge would thus be properties, including relational properties, of beliefs. Pursuing the analogy with individuals and substances, the upshot would be that there are metaphysically necessary truths that sum up the nature of knowledge *and are not conceptual truths*. If this is right then the concept of knowledge will not capture the nature of knowledge and reflection on the concept of knowledge cannot yield all there is to know about the nature of knowledge.

It is far from obvious that knowledge is a natural kind because it is not obvious that the concept of knowledge more closely resembles natural kind concepts, like those of water and gold, than it does family resemblance concepts like that of furniture. Nor is it obvious that we conceive of knowledge as having an underlying nature that is not represented by our concept. Whether or not knowledge is a natural kind, the point of most importance for present purposes is that there is no sharp contrast between enquiry into the phenomenon of knowledge and enquiry into the concept of knowledge. We can be sure of that while holding fire on the issue of whether knowledge is a natural kind. The methodological reflections that immediately follow are designed to reinforce this. But I shall in due course also question, as others have done, the claim that knowledge should be conceived as belief that satisfies certain further conditions.

3. In defence of conceptual enquiry

Kornblith's general approach draws upon considerations about how the category of knowledge is deployed among cognitive ethologists. These considerations look very like the sort of considerations one might advance as part of an enquiry into the concept of knowledge.

It is a matter of looking at how a concept is deployed with a view to learning about what, if anything, it picks out. I must acknowledge, however, that Kornblith's scepticism about the importance of conceptual enquiry is not explained just by his acceptance of the claim that knowledge is a natural kind. He targets a particular conception of conceptual enquiry that places a great deal of weight on intuition conceived as *a priori* knowledge.² I have no quarrel with much of what he says about the role and importance of intuition. My concern is that we should not identify conceptual enquiry with conceptual analysis, conceived as aiming to arrive at conceptual reductions by purely *a priori* methods. It is precisely this conception of philosophical method that informs the traditional analytical project in epistemology. But the seminal practitioners of self-consciously conceptual enquiry—proponents of so-called ordinary language or linguistic philosophy—showed little if any interest in conceptual reductions of this sort. Ryle's 'logical geography' (Ryle 1949), Wittgenstein's 'logical grammar' (Wittgenstein 1958) and Austin's 'linguistic phenomenology' (Austin 1956-57/1961) were as much as anything attempts to avoid misleading models of, for instance, thinking, making intelligent moves, meaning, understanding, intending, and so forth. The misleading models were ones that could easily suggest themselves on reflection. They were to be countered by close attention to distinctions we make in practice, and to facts—contingent facts—about what we count as cases of meaning, intending, and the like. The upshot was a conception of philosophy as revealing something of the structure of our understanding of various phenomena. It was about us because it was about our understanding. But it was about the phenomena as well.³

² His principal target is Bealer 1993.

³ Austin counters misconceptions as follows: 'When we examine what we should say when, what words we should use in what situations, we are looking again not *merely* at words (or 'meanings', whatever they may be) but also at the realities we use words to talk about: we are using the sharpened awareness of words to sharpen our perception of, though not as the final arbiter of, the phenomena.' (1956-57/1961: 130). Compare this from A.J. Ayer: 'It is ... indifferent whether, in this manner of philosophizing, we represent ourselves as dealing with

How can conceptual enquiry shed light on knowledge? The beginning of an answer is provided by the consideration that phenomena of which we know have what might be called a *conceptual shape*. Some of what goes into conceptual shape has to do with metaphysical categories. It is part of the conceptual shape of gold that gold is a substance. It is part of the conceptual shape of the property of being fragile that the property is dispositional. It is part of the conceptual shape of a battle that it is an event or related series of events. That X is a substance, or that it is a dispositional property, or that it is an event or series of events, are metaphysical truths that can be seen for what they are only by thinking both about X and about how we think of it. It is far from clear that they are conceptual truths, since, for instance, it would seem that we could have full mastery of the concept of fragility, yet lack altogether the concept of a dispositional property. Nonetheless, it is reflection on our thought and talk about what it is to be fragile that reveals fragility to be dispositional. To take another case, if I just think about battles I shall think of clashing armies, violent death, fear—the sorts of things that go to make up a battle. To home in on the idea of a battle as a series of events I need to reflect on the ways in which events are described by contrast with individuals and states. If *knowledge* admits of a reductive analysis in terms of *belief* then it has the conceptual shape of a composite with belief as one of its ingredients. Whether it is plausible that it has such a shape must be decided by conceptual enquiry.

Granted that conceptual enquiry can tell us something about phenomena, one might wonder why such enquiry should be especially called for in areas of traditional philosophical interest. No one would expect conceptual enquiry into *rust* to yield anything of importance. What is it about key topics in philosophy, like freedom and knowledge, which calls for conceptual enquiry? At least part of the answer is that the phenomena in which philosophers have taken an interest have features that very naturally suggest misleading models of their

words or as dealing with facts. For our enquiry into the use of words can equally be regarded as an enquiry into the nature of the facts which they describe' (Ayer 1956: 29).

character. Some phenomena very easily send us off in the wrong direction. The phenomena that make up our mental life are like this. Given the uncertainties that can attend judgements as to what others think and feel, and a natural inclination we have to think that we know what we ourselves think and feel, it is unsurprising that it should have been thought that looking within is *the* way to investigate mentality in general. When we look within what do we find? In Hume's version, which may stand for many, we find perceptions—experiential items that make up our total current experiential state. Hence, in Hume, the assimilation of beliefs to ideas, conceived as faint copies of impressions. It is a remarkable fact that this picture of our mentality persisted well into the twentieth century.⁴ In philosophical tradition it was dislodged in different ways by Ryle, Wittgenstein, Malcolm, Strawson, Quine, and others, who combined looking outwards at mentality as we find it in the world, with reflection on the kind of understanding we have of mentality. Knowledge also has a feature that is apt to encourage misleading models of its nature.

4. *Why the phenomenon of knowledge is apt to mislead*

Knowledge is apt to mislead because it is intimately linked to belief. If we think of belief in a very broad sense as a matter of thinking something to be so—taking it for true—then knowledge entails belief. Given that knowledge entails belief it is unsurprising that it should have been thought that belief is a constituent of knowledge. With that thought in play it was natural than to wonder what, along with truth, needs to be added to belief to make knowledge. From this perspective knowledge has the conceptual shape of a composite built from ingredients. It is important to appreciate, however, that there are plausible alternative ways to accommodate the assumption that knowledge entails belief. For instance, believing

⁴ See, for instance, the account of belief in Russell 1921.

that *p* might be a determinable of which knowing that *p* is one among many determinates.⁵ There is a kinship between my knowing that my keys have been left at home, because a member of my family tells me that they are there, and thinking that they are there because, given my recent movements, and how I look after my keys when I carry them, it is highly unlikely that they are anywhere else. But there are also important psychological differences. In the case in which I know the matter is settled; I'd be prepared to vouch for the fact that my keys are at home, and would discount the possibility that they are not. In the case in which I merely think they are at home I would be more guarded in what I'd be prepared to say, and my belief would be more vulnerable to doubt.

The assumption that knowledge yields to analysis of the traditional sort gives rise to problems. If our concept of knowledge does yield to such analysis, then we should expect that in applying the concept we should sometimes have regard to whether the conditions made explicit by the analysis are satisfied. It is not required that every time the concept is applied subjects satisfy themselves that the conditions are fulfilled. A child can apply the concept of an uncle to various people that he or she has learned to regard as uncles while having a shaky grasp of what it takes for one person to be the uncle of the other. But given that the uncle of a person is by definition either the sibling of a parent of that person, or the spouse of a woman who is the sibling of one of the parents, one would expect that some people, for instance, lawyers and those tracing family trees, would on occasion have regard to whether those conditions are satisfied. This, I take it, is what actually happens in this case. But in the case of concept of knowledge there seems to be nothing quite comparable. It is not clear that common sense thinking about knowledge explicitly brings into play the concepts that regularly figure in analyses of knowledge, for instance, the concept of a reliable belief-forming process. Everyday ascriptions knowledge, including self-ascriptions, *are* sensitive to considerations about what people saw, or otherwise perceived to be so, and to what they

⁵ Compare Williamson 2000:43-44, and his wider discussion of the relation between belief and knowledge.

remember to be so. And they are sensitive to whether people have evidence that is adequate for knowledge. So they are bound to have regard to whether people are in command of ways of telling that this or that is so, straight off from perception or from recognition of the evidential significance of known facts. It is a further matter whether being in command of ways of telling admits of conceptual analysis in terms of reliable belief-forming processes. It is not even clear that in everyday thinking about knowledge justified belief is routinely considered to be essential for knowledge. Reflection on such thinking in relation to perceptual knowledge might easily suggest that, far from justified belief being an essential component of knowledge, it can derive from what is antecedently known by perceiving.

It is perhaps not surprising that under the terms of the traditional project knowledge should have turned out to be so elusive. But a theoretical approach under which the nature of knowledge is elusive makes it hard to understand why knowledge can have the practical value for us that it does. It often matters to us whether we know that something is so, as opposed, for instance, to being fairly confident that it is so based on inconclusive evidence. That is because it makes a difference to how we carry out plans and to what we are prepared to say to others. And in social interactions people routinely rely on our discriminating between knowing and lesser states. Enquirers could justifiably feel let down if we were assert that p , giving it to be understood that we know, when the best evidence we have only makes it highly likely that p . Knowledge could not facilitate our social interactions as it does as it does unless facts as to what we, or others, know were often readily available to us. An adequate theory of knowledge should not make it mysterious that we can be easily apprised of such facts.

It might be suggested that I have conflated a pre-philosophical grasp of what knowledge is and a theoretical perspective on the nature of knowledge. The idea would be that at some level we all know what knowledge is, and are sensitive to the conditions that theorists aim to identify. On this picture it is not that knowledge is elusive but that the right theory of

knowledge is elusive. But the picture does not address the heart of the matter. If our concept of knowledge admitted of analysis along traditional lines then the elements of the analysis should inform our thinking, and if they informed our thinking they ought to be more readily recoverable than they have proven to be. It might be suggested that the gap between mastery of a concept and ability to reach a clear view of its character can be explained in terms of an analogy between being able to speak grammatically and representing the rules of grammar explicitly. We can know how to speak grammatically without representing the rules of grammar to ourselves. Might it not be that we can competently apply concepts without having a clear view of their character? The answer is, obviously, 'Yes', since we do apply concepts competently while sometimes lacking a clear view of their character. Nonetheless, the analogy underplays the extent to which, in reflective beings such as us, applications of predicative concepts go hand in hand with some conception of what puts one in a position to apply them. We can accord with the rules of grammar blindly, though even here it helps to be aware of some of the basics. We don't apply concepts blindly according to rules of which we have no conception. We have common sense conceptions of what puts us in a position to apply them. We have conceptions of what puts us in a position to apply the concept of knowledge, but I doubt that these embrace the reducing conditions of traditional analyses.

It might be suggested that if my worry here were well founded it would be hard to make sense of the difficulties encountered in most areas of philosophy. Philosophy exists, it might be said, because of the gap between having mastery of a concept and being in a position to provide a perspicuous representation of its character. Two points need emphasis in response to this challenge. The first is that it can be difficult to achieve a clear view of the character of a concept, not because the facts of the matter are so very elusive, but because prevailing modes of philosophical thinking obstruct a clear view. As I have observed already, many of the phenomena and concepts of interest to philosophy are apt to generate puzzlement. The second point is that in some areas of philosophy the difficulties encountered are not the

difficulties of characterising a determinate concept that we already employ. A great deal of philosophy is conceptual/theoretical development rather than conceptual analysis. The difficulty is in working from the partial and impressionistic concepts we have to something better. I suspect that this is true in areas such as rights, responsibility, and justice.

From the perspective adopted here, facts as to what we, or others, know are often readily available to us. I often know that water in a heating kettle has just boiled or that vegetables that I am cooking are ready to serve. Not only that, when I know these things, I know how I know them—from the sound of boiling water, from the texture of the vegetables. And I often know that others know this or that. For instance, I may know from his reaction that my companion has just seen, and thus knows, that his train is approaching. Knowledge of these kinds is not so very hard to come by. A theory of knowledge should make sense of this.

5. Ways of telling

Knowledge appears less elusive if we think of knowing that something is so as being in a state acquired through the exercise of an ability to tell that something is so. Sometimes we tell just by looking, hearing or through some other sense. Sometimes we tell that something is so by recognising some fact that indicates that it is so, as when a gamekeeper tells from tracks on a path that deer have recently passed. On this view, the state of knowing depends essentially on the deployment of a way of telling—a mode of knowledge-acquisition. This by itself is hardly open to dispute. The bite comes with the detailed characterisations of ways of telling and with, as I see it, the realisation that our conceptions of ways of telling are not reducible to terms that do not implicate the concept of knowledge. This central idea emerges not just from reflection on the difficulties of providing a reductive analysis but also from

reflection on common sense thinking.⁶ We explain particular cases of knowledge in terms of the deployment of specific ways of telling, and we seek to account for knowledge in general by charting, and reflecting on, the character of the various kinds of ways of telling that we have. On any plausible approach along these lines knowledge-states acquired through the exercise of perceptual-recognitional abilities will be central cases of knowledge. These abilities include the ability to tell by sight that a rose, or an empty cup, or one's friend Bill, is there.

As I look at a rose in a vase, I know that a rose is before me. I know this just because I see the rose and recognise it to be a rose. This is an informative explanation of my knowing since it picks out the relevant mode of knowing from a range of other possible ways of knowing the same thing. Further,

- (1) my seeing that there is a rose before me is a state that of its very nature is acquired through (a) my visually picking out the rose, and (b) recognising it to be a rose;
- (2) the recognition is the exercise of a certain recognitional ability—an ability to tell by sight that a rose is present;
- (3) the exercise of the ability is to be conceived in success terms—it is exercised only if I tell, and thus come to know, that a rose is there.

There are abilities to do certain things with some degree of reliability and thus a fair proportion of the times on which one tries. An example is a pole-vaulter's ability to clear the

⁶ Compare remarks of Colin McGinn (1984: 541-42). McGinn makes the notion of a discriminative capacity central to his epistemology and doubts that the notion can be analysed in terms that do not bring into play the concept of knowledge. In Millar 2007 and forthcoming (a) and (b), I have outlined the conception of ways of telling in play here, as applied to perceptual knowledge. I believe that the conception can be extended to knowledge of facts from indicating evidence, as illustrated below.

bar at, or close too, the highest height he or she has so far cleared. With respect to abilities of this sort, no special explanation is required of a failure at a single attempt. The ability to read English well is not of this sort. If one can read English well then one reliably does so in the sense that most of the time one succeeds when one attempts to. There may be circumstances (e.g., migraines, tired eyes) when one attempts but fails. And it is important for having the ability that one notices when the circumstances are not propitious and desists from trying when that is so. The difference between this ability and the pole-vaulter's ability is that between (a) abilities such that attempts to exercise them succeed unless something gets in the way, and (b) abilities such that lack of success at an attempt is no great surprise, and does not require explanation in terms of impairment, distraction, or some other interfering factor. Perceptual-recognitional abilities are of the (a)-type. If I can tell by looking whether or not there is a robin on my garden bird table then a failure to judge rightly in conditions that are conducive to the exercise of the ability calls for special explanation.

An important feature of perceptual-recognitional abilities is that they are *indexed to types of environment*. That is to say, they are abilities to recognise that something is so in a certain kind of environment and one counts as having the ability only relative to an environment of the right type. I can recognise roses from the way they look in environments in which things that look just like roses are always, or nearly always, roses. Suppose that in my environment it were to come about that skilfully made artificial roses began to appear unmarked in the sorts of places in which roses usually appear, so that there is a significant chance that anything looking just like a rose is not. Ordinary visual inspection of the kind that yields knowledge would not discriminate the real from the artificial. In that case the environment which I inhabit would have changed in such a way that, even if I am prone to go through the same judgement-forming process as I do now when I judge by looking that something is a rose, I would cease to have the ability to tell that something is a rose by ordinary looking, relative to that environment. I count as having the ability to tell that something is an F from

its having the appearance of an F only relative to environments in which things that have that appearance are nearly always Fs.⁷

It might seem that all that is on offer here is the view that knowledge is the state you get into when you have exercised some ability, whatever that might be, to acquire knowledge. That would be thin beer indeed, but it is not all that is on offer. Recall that a key task for the theorist on this approach is to characterise the varieties of ways of telling that something is so. Fortunately, within the realms of common sense thinking, we have rather specific ideas of what goes into particular perceptual-recognitional abilities. For instance, if you want to be able to tell by sight that a robin is before you then you will need to have learned to recognise robins from their features, including their colouring, shape of beak, movements, and so on. Starting from this common sense standpoint we can move on to a more theoretical level, noting, for instance, that recognition is not the drawing of a conclusion from assumptions descriptive of features. In the case of recognising the robin it is a matter of taking in the distinctive *Gestalt* of the bird.⁸ Though the relevant features are describable, recognitional knowledge that a bird is a robin is not the upshot of an inference from a set of descriptions. Similarly, the particular *Gestalt* that is shown on an anxious friend's face, is something we take in recognitionally, not by making an inference from assumptions that characterise look and demeanour.⁹

Intriguing problems arise when we attempt to achieve a clear view of the character of knowledge based on evidence.¹⁰ On the approach adopted here we need to understand the distinction between the kind of assurance that typically goes with knowing that p and the kind of confidence one has when one is fairly confident that p on the basis of evidence. Evidence

⁷ A reason for doubting that knowledge is a natural kind is that environments that are propitious for the possession of a recognitional ability may shade into environments that are not, so that in some places it may be indeterminate whether one counts as having the ability. Knowledge is as obtainable as it is because the world rarely throws up problems of this sort.

⁸ Compare Austin's remarks (1946/1961: 52-53) on telling a goldfinch from its features.

⁹ Compare McDowell 1982. For related discussion of the scope of perceptual knowledge, see Millar 2000.

can support statements to a greater or less degree. What is required for evidence to be strong enough for knowledge is not generally made clear. Indeed, it is hard to see why we should be concerned to distinguish between evidence-based *knowledge* and true belief that has a high degree of evidential support, but lower than that which is supposed to be adequate for knowledge.

Our conception of adequate evidence for knowledge, I suggest, should be viewed in the light of the idea that knowledge is acquired through the employment of ways of telling. Specifically, if *e* is evidence that *p* that is adequate for knowledge that *p*, then it must be possible to tell that *p* given *e*. To take a specific case (borrowed from Travis 2005) suppose that in a certain environment the occurrence of a distinctive bleating sound on some occasion is evidence adequate for knowledge that sheep are nearby. This will be so only if the occurrence of the sound on that occasion *indicates* that sheep are nearby. I shall not attempt to provide here a general account of indication. I take it that in the case in hand the indication is grounded in a causal relation: the bleating sound indicates that sheep are nearby because the sound is due to the sheep nearby. Indication is factive: evidence *e* indicates that *p* only if *p*. Evidence *e* may indicate that *p* even if nobody has a way of telling that sort of thing from that sort of evidence. The shepherd in the environment in question will have a way of telling that sheep are nearby from the bleating sound only if he could not easily be wrong in judging sheep to be nearby from the occurrence of this sort of sound. For this to be so it must be that this sort of sound would rarely if ever be heard, in these parts, if no sheep were nearby. All this seems plausible, but it leaves us with a problem of understanding how the shepherd must stand to the fact that this sort of sound would rarely if ever be heard but for the presence of sheep, if he is to be in command of this way of telling that sheep are nearby. A fruitful line of enquiry, I think, is to look into the extent to which evidence-based knowledge is like recognitional knowledge.

¹⁰ The discussion in the remainder of this section was stimulated by Travis 2005. See my critical comments in

The evidence-implicating way of telling of which the shepherd is in command is similar in structure to the way of telling that birdwatchers deploy when recognising that robins are there from sight of their features. Against the background of epistemological tradition it is tempting to represent these ways of telling as resting on acceptance of an inductively based generalisation linking, in the one case, robin-features with robins and, in the other, the bleating sound with sheep. There are two problems with this. One is that inductive support is usually conceived as at best making what it supports merely probable. For knowledge we want more than mere probability. The other is that it is far from clear that those who have such ways of telling are in general justified in accepting the required generalisations on inductive grounds. Exposure to instances may have had a role in inculcating the ability in question and will have had a role in maintaining the ability once it is up and running. It does not follow that to be in command of the relevant way of telling the subject must accept an appropriate connecting generalisation on strong inductive grounds. That the bleating sound would rarely if ever occur but for the nearby presence of sheep will be part of the shepherd's picture of the world—a picture that, as a matter of psychological fact, is sustained by experience. If the role of the picture were just to supply premises for reasoning then the question would arise as to what justifies the shepherd in thinking and acting against the background of the picture. I suggest that this is not the way to think of the role of all of the elements of the picture. The picture ensures that the shepherd does not operate blindly but has some conception of this and other ways of telling. It would put him in a position to appreciate that he has judged wrongly if an occasion were to arise on which the bleating sound was mimicked by a practical joker. It could also enable him to understand that if it should ever come about that the bleating sound ceased to be a reliable indicator of the presence of sheep he would then need more than the sound to be sure that sheep were nearby. The elements of the picture need not be known truths. Indeed, the picture need not be accurate in its

entirety.¹¹ What matters is that elements of it supply a perspective on what is, thanks to the environment, a way of telling that sheep are nearby.¹²

7. *Knowledge matters*

Starting off from a defence of conceptual enquiry in epistemology I suggested that knowledge, like many other phenomena in which philosophers take an interest, has features that are apt to mislead. In particular, because knowledge entails belief it is apt to make us think that it consists in belief that satisfies certain conditions including, obviously, truth. I sketched an alternative approach on which states of knowing essentially depend on ways of telling. Such a view receives support from consideration of the practical value of knowledge.

Knowledge does seem to matter to us. The natural way to represent the goal of enquiry into whether or not something is so is in terms of aiming to find out whether it is so. Finding out is nothing less than coming to know. Since enquiry clearly matters to us, knowledge does as well. This consideration shows that knowledge matters to us, but it does not address the classic value problem deriving from Plato's *Meno*: how to explain why knowledge is better than merely true belief. Adapting Plato's own example, suppose that I have asked someone in a strange town whether a certain road leads to the cathedral. The person asked has no idea whether it does or not, but mischievously says, 'Yes'. As it happens this is correct and I believe what he says. I thus acquire a true belief that the road in question leads to the cathedral, but clearly I do not know that it does. Despite not knowing I shall carry out my intention to visit the cathedral if I take the road I now think leads to it. Why should it matter that I have knowledge rather than true belief that falls short of knowledge? Merely true

¹¹ A theory that is not entirely accurate but is empirically adequate in some domain might put experimenters or observers in command of a way of detecting certain facts from indicator facts within that domain.

¹² The notion of a perspective has figured prominently in Ernest Sosa's epistemology. See, for instance, several essays in Sosa 1991. It is not clear to me that possession of a perspective on how we know is essential for

belief, it seems, serves just as well for practical purposes. The classic *Meno* problem is to explain why knowledge is more valuable than merely true belief, assuming that it is. But it brings in its train the further problem of why knowledge is more valuable than, for instance, justified true belief. On some ways of filling out the scenario just described I would have a justified true belief as to which road leads to the cathedral. Is there any reason to think that knowledge is more valuable than that?¹³

I shall sketch a conception of why knowledge is of practical value to us, which I think will help us to address these value problems. The key idea is that knowledge on the part of reflective creatures like us *typically* has *instrumental value* in relation to thought and action, additional to that accruing to the implicated belief solely in virtue of its being true. In particular, knowledge that *p* typically enables us to have the practical benefits that go with well-grounded *assurance* that *p* and well-grounded *security* in believing that *p*

Assurance that *p* is a matter of being sure that *p* in the way one is liable to be when one knows or takes oneself to know that *p*.¹⁴ For a further example of the difference that assurance makes, consider the difference between my being sure that I have sent off a letter because I recall doing so, and being fairly confident that I have because so far as I can see the letter, which I had sealed in an envelope, is no longer among my things. In the latter case, I think I must have sent it and have some basis for thinking so, but I neither know, nor take myself to know, that I have sent it. In the former case the matter is settled. I am assured in the way one is liable to be assured when one knows or takes oneself to know. I discount the possibility that the letter is still around and would be prepared to vouch for having sent it.

knowledge, or even for the kind of knowledge we have. It is necessary for the role that knowledge plays in the lives of reflective beings like us. See section 7.

¹³ Duncan Pritchard (forthcoming) describes this as an instance of the secondary value problem: why is knowledge more valuable than some lesser state comprising true belief plus some other condition.

¹⁴ I make no apology for introducing assurance via the concept of knowledge. The naturalness of doing so reflects the fact that concept of knowledge and its cognates are more basic in our thinking than concepts that are often used to pick out the ingredients of knowledge.

When assurance results from knowledge concerning our environment does it come about simply because of the mode of acquisition of the knowledge or does it depend on knowledge that, and how, one has the knowledge in question? It may well be that our simply seeing that p would normally instil assurance that p , irrespective of whether we knew that, or how, we know that p . Whether that is right or not, when we know that p through seeing that p , or other perceptual means, we typically know that, and how, we know that p , at least for a while. Knowing how one knows has a significant role in maintaining assurance. If I know that my keys are in my pocket because I can feel them there, my assurance that they are is grounded in my awareness of the fact that I feel them there, and thus know that they are there. Of course, it is possible to be assured when one falsely takes oneself to know. But that does not make knowledge, and awareness of how one knows, any less important in the production of assurance. Assurance may arise from sources other than knowledge, but it does not arise willy-nilly. Much of which we are assured comes from what we know, and when it does not that is often because it is to us just as if we know.

Knowledge produces assurance that is well-grounded, because grounded in the deployment of some way of telling. It also produces well-ground security of belief. Security is potential to remain stable in a range of circumstances that one might well encounter. Consider again the case of the road to the cathedral. Though my belief happens to be true, my informant does not know, and therefore I do not know. Under the conception of justification that prevails in Gettier cases I may have a justified true belief. As things stand, I am just as sure as I would have been had I known that this is the right road. But now what if someone else, who also does not know, were to overhear the interchange with my informant and then, catching up with me, tell me that this informant does not know which road leads to the cathedral. In that case, all else equal, there is now an issue as to whether I have any basis for thinking that I am on the right road. As a result I might well give up my belief to that effect. Even if as things actually worked out I reached the cathedral by acting on a merely true

belief, my true belief would not be secure because vulnerable to the discovery that it was not well-founded. (Cf. the burglar example in Williamson 2000: 62-63.) For practical purposes security grounded in knowledge is important: it is important not only that we lay hold of true beliefs on things that matter to us but that we hold on to them once we have them. This enables us to act on them without time-consuming re-checking and it enables us to be non-vacillating informants.

Like assurance, security is induced and maintained in large measure because of knowledge of how one knows—of the way of telling deployed. Suppose I become assured that my keys are on my desk because I see that they are, but then I lose track of the fact that I have seen them. I might for a while continue to be prone to act as if they were there and to vouch for their being there if asked, but if it were to occur to me that I have no idea how I know, then my assurance would be liable to be undermined. Our awareness of how we know, or came to know, is thus an important factor in maintaining assurance and thus rendering our beliefs secure. The value of well-grounded security is practical; it has to do with how it enables us to confront a range of possible situations that we might encounter. It does not follow that a person who has a justified true belief, or even a merely true belief, will always, as things actually turn out, be worse off from a practical point of view than someone who knows that *p*. The point is that knowledge that *p*, in beings such as ourselves, who can know how they know, is fit to sustain assurance in the face of challenges that arise from apparently contrary evidence and inducements to doubt. We should go for knowledge, then, not just because by doing so we are likely to have a better chance of avoiding falsehood, but because knowledge once gained is useful for smooth interactions with others and enables us conveniently to carry out our projects.¹⁵

¹⁵ In recent discussions, for instance, Kvanvig (2003), it seems to be assumed that a solution to value problems must show that knowing is always is more valuable than corresponding lesser states would have been. For reflective beings like us, there is a sense in which this will be so, since it is always a good thing to be well prepared for what one might encounter. But as things actually pan out one may not encounter a situation that calls for one to exploit the fact that one knows, and might have been no worse off than if one had had, say, a

This conception of the value of knowledge imposes an adequacy constraint on accounts of knowledge: that they should make sense of how knowledge can have the practical value for us that it does have. That requires that it should make sense of how we are as good as we are at telling whether we or other people know this or that. While I do not claim to have established here that traditional accounts do not satisfy this constraint I have given reasons for doubting that they do (in section 4). From the standpoint adopted here the conception of knowledge as necessarily deriving from ways of telling has a better prospect of satisfying the constraint because it lends itself to an explanation of how we tell whether we or other people know. This is because it accords a central place to kinds of ways of telling that manifestly feature in our common sense thinking: telling from its visible features that a rose or a robin is present; telling that there is fire from the presence of smoke; telling that my keys are at home from the fact that a member of my family tells me that they are. Of course, it is not enough that these ways of telling should be familiar; it also needs to be plausible that we can often tell when they have been exercised. I think that the prospects on this front are good and that we should invoke the notion of higher-order ways of telling to address the matter. The very experiences that enable me to know that it's a rose before me enable me to know that I see that a rose is before me, thanks to a higher-order recognitional ability whereby I can tell that I see a rose. This ability is one that anyone who can tell by looking that a rose is present is likely to have. The shepherd can tell that he knows that sheep are nearby on hearing the bleating, since the very experiences that enable him to detect the bleating also trigger recognition that he knows the sheep to be nearby from the bleating, thanks to a higher-order ability to tell that he knows such things. And we often know when others know because we can tell what they perceive, or tell that they recognise the significance of some indicator fact. Even this sort of way of telling, I suspect, is closer to recognitional ways of telling than

Gettierised justified true belief. This may be weaker than what some are looking for but I do not see it as an objection to the account.

would generally be supposed, though the matter calls for much more discussion than I have been able to give it here.¹⁶

Department of Philosophy

University of Stirling

Stirling FK9 4LA

U.K.

References

Austin, J. L. (1946) 'Other Minds', *Proceedings of the Aristotelian Society*, Supplementary Volume 20, 148-87. Reprinted in Austin (1961).

Austin, J. L. (1956-57) 'A Plea for Excuses', *Proceedings of the Aristotelian Society*. Reprinted in Austin (1961).

Austin, J. L. (1961) *Philosophical Papers* J. O. Urmson and G. J. Warnock (eds.) (Oxford: Clarendon Press)

Ayer, A. J. (1956) *The Problem of Knowledge* (London: Macmillan).

Bealer, G. (1993) 'The Incoherence of Empiricism' in S. Wagner and R. Warner (eds.) *Naturalism: A Critical Appraisal* (Notre Dame: Notre Dame University Press).

Kornblith, H. (2002) *Knowledge and Its Place in Nature* (Oxford: Clarendon Press).

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

¹⁶ This article contributes to a project on the Value of Knowledge funded by a Research Grant from the UK's Arts and Humanities Research Council. I am grateful to the Council for its generous support and to co-participants in the project, Adrian Haddock and Duncan Pritchard, for discussions and comments on the matters discussed here. I also thank participants at workshops and conferences funded by the Council's Research Grant for the stimulus provided by their contributions and by their comments on some of the ideas presented here.

- McDowell, J. (1982) 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy* 68, 1981, 455-79. Reprinted in McDowell's *Meaning, Knowledge and Reality* (Cambridge, MA.: Harvard University Press, 1998), 369-394.
- McGinn, C. (1984) 'The Concept of Knowledge' *Midwest Studies in Philosophy IX*, 529-54.
- Millar, A. (2000) 'The Scope of Perceptual Knowledge' *Philosophy* 75, 73-88.
- Millar, A. (2005) 'Travis' Sense of Occasion' *The Philosophical Quarterly* 55, 337-42.
- Millar, A. (2007) 'What the Disjunctivist is Right About' *Philosophy and Phenomenological Research*.
- Millar, A. (Forthcoming (a)) 'Perceptual-Recognitional Abilities and Perceptual Knowledge' in A. Haddock and Fiona Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge* (Oxford: Clarendon Press).
- Millar, A. (Forthcoming (b)) 'Disjunctivism and Scepticism' in John Greco (ed.) *The Oxford Handbook to Scepticism* (Oxford: Oxford University Press).
- Pritchard, D. (Forthcoming) 'Recent Work on Epistemic Value', *American Philosophical Quarterly*.
- Russell, B. (1921) *The Analysis of Knowledge* (London: George Allen and Unwin).
- Ryle, G. (1949) *The Concept of Mind* (London: Hutchinson).
- Ryle, G. (1970) 'Autobiographical' in O. P. Wood and G. Pitcher (eds.) *Ryle: A Collection of Critical Essays* (London: Macmillan).
- Sosa, E. (1991) *Knowledge in Perspective* (Cambridge: Cambridge University Press).
- Travis, C. (2005) 'A Sense of Occasion' *The Philosophical Quarterly* 55, 286-314.
- Williamson, T. (2000) *Knowledge and Its Limits* (Oxford: Oxford University Press).
- Wittgenstein, L. (1958) *The Philosophical Investigations* (Oxford: Blackwell).