

Draft Only. Comments Welcomed.

Knowledge, Recognition, and Evidence

Alan Millar

1. *Introduction*

There is a curious phenomenon in current epistemology—a clash between widely held theoretical assumptions and common ways in which we think of knowledge. The aim of this discussion is to canvas accounts of specific kinds of knowledge that enable us to do more justice to the common ways of thinking.

The concept of knowledge is fundamental for our ordinary ways of thinking about enquiry in all its forms, systematic and otherwise. The most natural way to describe what interests us when we are interested in whether or not something is so is by saying that we wish to find out whether or not it is so. Finding out is nothing less than coming to know. Because of limitations of time and other resources enquirers may have to settle for less. They may have to settle for a body of evidence that favours one position but does not rule out others. They may not even achieve that. But their enquiries will nonetheless have been conducted under the aspect of finding out and so coming to know.

What is this knowledge that is the goal of enquiry? Certainly something different from true belief grounded in a greater or lesser degree of evidential support. We mark such a distinction in everyday situations. Suppose, for instance, that I have left my house without my keys. I may justifiably be fairly confident, and right, in thinking that they will be on the telephone table in the hall at home, on the grounds that I usually leave them there when at home. My state is different from that of knowing that the keys are on the telephone table because my son has just told me by phone that they are there. There is a difference between knowing that there is a train for Edinburgh scheduled to depart at 7pm because I have just checked the timetable, and being, justifiably, fairly confident, and right, in thinking that this so on the grounds that my friend has told me that he is fairly confident that it is so. There is a difference between knowing that the mains switch is off, because I, or someone else, has checked and being, justifiably, fairly

confident, and right, in thinking that it is off because my working partner would in the usual course of events have switched it off by this time.

That knowledge is different from justified true belief is not news. The point on which I wish to focus lies elsewhere. It is that when we know that something is so the matter is usually settled and, when this is so, an important factor in its being settled is our knowing how we know. Seeing that my computer is switched on establishes in me a relatively stable belief that it is on. My awareness that I see or have recently seen that it is on is apt to render the belief more secure—less vulnerable to being undermined. That is because challenges to the belief have to reckon with the fact that I see or have seen that the computer is on. Let it be granted that knowing that *p* does not entail knowing that one knows that *p*. Still, we often know how we know and typically do in perceptual cases and in many cases of testimony in which we acquire knowledge from others. Consideration of how we know serves to assure us, as well as others, that we know. Assurance on that score can be of practical import. The upshot of my being so assured would be that I do not engage in further enquiry, would discount alternatives to what I know, and would be in a position to be an unhesitant informant to others on what I know. On the matter of the train schedule, for instance, I would be in a position to tell an enquirer that there is a train scheduled to depart for Edinburgh at 7pm. In the cases in which I merely have evidence on the basis of which I am pretty sure, the matter would not be settled for me and I would not take it to be settled. It would be open to me, if I needed to know, to enquire further if practical considerations permitted. If I don't enquire further I would be prepared for disappointment in case I'm wrong. I might make back-up plans. If asked by others about whether a train for Edinburgh is scheduled to depart at 7pm I would not assert that it does. I'd give a qualified response like, 'I think it does' or 'I'm assuming so' or maybe, 'I'm pretty sure that it does'. We all have a sense of when it is appropriate to give such responses, as opposed to unqualified assertion or confirmation. We also understand the significance of these responses when they are given to us. Whether the matter is important or trivial we are liable to feel let down or at least disappointed if others give it to be understood that they know something when the evidence available to them warranted, at best, some fairly high degree of confidence. Of course, one can have assurance that one knows when one does not know. The kind of assurance we want is assurance that is relatively secure because grounded in what we know. If it were not so grounded it would be vulnerable to discovery this is

so. Finding out that you only thought you saw that Mary was at the party, in the absence of another reason to think she was there, might easily leave you still wondering whether she was.

There are strands in epistemological tradition that are apt to obscure the difference between being in a position such that one knows, and the matter is settled, and being in some lesser state implicating true belief in which the matter is not settled. In the next section I consider two such strands. In subsequent sections I outline a conception of knowledge that helps to make sense of the difference.

2. Obscuring the difference

The difference can be obscured by the adoption of a weak conception of evidence adequate for knowledge. Weak conceptions are widespread. In recent years they have been especially visible in contextualist literature. Typical examples include knowing that your bank opens on Saturday on the grounds that a few weeks ago you were in that bank on a Saturday, and knowing that your car is in the car park on the basis the fact that you left it there earlier.¹ There could be knowledge in these cases, depending on how the scenarios are filled out. In the bank case the environment might be one in which banks are permitted to change their opening hours only on a specified date each year and the two occasions that figure in the case relate to Saturdays that are within a period in which opening hours are constant. In the car park case, the car park may be very secure. These sorts of factors are not the ones that matter most in contextualist discussions. In those discussions the factors that are held to determine a sense in which it is true that the subject knows are factors in the context of the attributer of knowledge that bear on whether anything is at stake for the attributer as to whether or not the subject knows. It is taken to be intuitive that an ascription of knowledge could be true in a low-stakes case given the evidence on which the subject relies, even though the evidence as specified clearly does not rule out alternatives that, for all we have been told, seem to be relevant (e.g., that the bank has changed its hours; that the car has been stolen since it was left). This view underestimates the strength of evidence required for knowledge and accordingly blurs the distinction between knowing on the basis of evidence and having a true belief based on evidence that confers on the belief a high degree of probability. Evidential support admits of degrees. It is widely assumed that evidence adequate for knowledge

is just evidence that is strong enough. How strong is strong enough remains obscure. And there is a serious question as to why we should take a matter to be settled when our evidence merely confers a fairly high degree of probability on the proposition in question.

The significance to us of our knowing is also liable to be obscured in discussions of perceptual knowledge. This is despite the fact that in cases of perceptual knowledge it is especially clear that the matter is settled. I may have perceptual knowledge that Bill has entered the room because I have just seen that Bill has entered the room. In cases like this we have more than just evidence, or even strong evidence, for thinking that something is so. J. L. Austin illustrates the general point forcefully in the following passage.

The situation in which I would properly be said to have *evidence* for the statement that some animal is a pig is that, for example, in which the beast is not actually on view, but I can see plenty of pig-like marks on the ground outside its retreat. ... But if the animal then emerges and stands there plainly in view, there is no longer any question of collecting evidence; its coming into view doesn't provide me with more *evidence* that it's a pig, I can now just *see* that it is, the question is settled.²

Epistemology ought to be able to make sense of such data. Yet theoretical treatments of perceptual knowledge often flout Austin's point. The kinds of treatment I have in mind are sometimes motivated by the conceptual-analytical accounts of knowledge, in terms of justified true belief plus further conditions. This leads theorists to delve below the level of seeing-that and to treat knowing in virtue of seeing-that as involving justification by evidence. Consider how zebra-cases are discussed. Looking into an enclosure in a zoo I see zebras. It is quite commonly thought that in such a situation I have evidence for thinking that the animals I am looking at are zebras. This would be provided by the fact that animals of such-and-such an appearance are before me or perhaps by the fact that it visually appears to me that there are zebras before me.³ In either case, the evidence, so conceived, could obtain even if no zebras were there. From this point it is a short step to thinking that in this situation, because of what constitutes my evidence, I

¹ The bank case is from DeRose 1992.

² Austin 1962: 115.

do not discriminate zebras from cleverly disguised mules. And because that is so it is thought that my justification for thinking that what is before me are zebras does not transmit to the claim that they are not cleverly disguised mules.⁴ But on the view in question it is just as problematic that I know that the animals are zebras. That is because it is far from transparent how it can be that I have evidence adequate for knowledge in the good case, in which I am indeed looking at zebras, when I would have the same evidence if I were looking at cleverly disguised mules. It might seem that there is no problem here. One can be justified in believing what is false, so it might seem that we should expect that there can be pairs of good and bad cases that are on a par with respect to the possession of grounds. The trouble with this is that it obscures the distinction between what is settled and what is not.⁵ What needs explaining is how the same evidence settles the matter in one case and not in the other. It cannot just be that in the good case it is true that *p* and in the bad case it is false that *p*. That gives us that in the good case there is evidence that *p* and it's true that *p*. It does not give us what we would need if this kind of knowledge were evidence-based—that the evidence in the good case clinches the matter.

In the sections that follow I outline conceptions of certain modes of knowledge that make sense of the idea that one can know in those modes in such a way that the matter is settled. In section 3 I deal with perceptual knowledge. In section 4 I deal with knowledge from indicator facts, for instance, knowing that one's car's tank is half full from the fact that the fuel gauge displays a reading on which it is half full. It emerges that knowledge from indicator facts is more like perceptual knowledge than one might think.

3. Perceptual knowledge and recognitional abilities

The fact that a zebra is there is a 'rich' fact in that it concerns the presence of an animal of a certain kind, not just the presence of a thing with certain superficial features that make up it

³ For the latter version, see Wright 2005: 246. The kind of thinking outlined here is familiar in discussion of zebra cases taking a cue from Dretske 1970.

⁴ For this line of thought, see Wright 2002 and Dretske 2005.

⁵ Essentially the same problem arises if one preserves the distinction between perceptual knowledge and evidence-based knowledge but thinks of perceptual knowledge as implicating belief that is justified by one's experiences, rather than by facts to the effect that one's experience has such-and-such a character. I defended such a view of justified belief in Millar 1991, with a view to addressing the scepticism in Davidson 1983. For recent work in this vein, see Pryor 2005.

appearance to some sense modality. From the perspective adopted here the problems arising from zebra cases stem from treating perceptual knowledge involving ‘rich’ facts as if it was implicitly evidenced-based. As I observed in the previous section, there is pressure to do so stemming from conceptual-analytical theories of knowledge that take knowledge to be true belief plus the satisfaction of other conditions. There are other motivations rooted in empiricist instincts to treat perceptually acquired knowledge of rich facts as evidence-based on the grounds that perceptual knowledge does not extend to rich facts but takes in only facts concerning superficial, perceptually detectable features. Yet there is an entirely natural way to account for how there can be perceptual knowledge of rich facts that is not evidence-based.

Looking into my garden I sometimes see that, and thereby know that, a *robin* is on the bird table, not just that something with a certain colour, shape, and way of moving, is there on the table. In such a case my judgement, made just from looking towards the table, is phenomenologically immediate. I do not infer that a robin is there from some prior assumption to the effect that something with certain superficial features is there. Rather, I simply recognise that the bird is a robin. To do so I need a specific recognitional ability—to tell that something is a robin from the way it looks. In exercising this ability I do, of course, respond to the features that make up the visual appearance of the bird. Those features explain why I judge it to be a robin. They coalesce in a manner that triggers recognition, so that the bird is seen as, and seen to be, a robin, but the recognition need not be effected by acceptance of some prior description of superficial features. Recognition is a response to a perceived appearance, not necessarily to a judgement about the appearance perceived or the features that make up the appearance. There is no need to treat knowledge acquired by such recognition as implicitly evidence-based.

On the view adopted here the notion of a recognitional ability is the key notion for an understanding of what perceptual knowledge is. The natural way to think of such abilities is in strongly epistemic terms, as knowledge-acquisition abilities. The ability I exercise when I tell of the bird that it is a robin is an ability to do no less than recognise the bird to be a robin from its visual appearance. That is to say, it is no less than an ability to tell, and thus come to know, of something that it is a robin from the way it looks.⁶ It can clearly be explanatory of a particular

⁶ I invoke recognitional abilities as part of a wider project concerning how to pursue the theory of knowledge if one is sympathetic, as I am, to criticisms of traditional epistemology of the sort found in, for instance, McDowell 1993 and 1995, and Williamson 2000. I introduce recognitional abilities in Millar 2007 and forthcoming, in the course of addressing issues raised by disjunctivism.

case of knowledge to describe it as acquired via the exercise of this ability. What I know in such a case can be known in other ways, for instance, from the bird's chirp or by being told that it's a robin by someone looking into the garden. So it is informative to know that a given case is a case of telling from the look of the bird. Note too that this does more than give us an account of how the knowledge came about. It is constitutive of the knowledge that it is knowledge through the exercise of an ability to recognise robins as robins by sight.⁷ Here we account for the knowledge not by saying how an implicated belief is justified, but by describing how the subject knows. (I shall return to the issue of justified belief below.)

Can such an approach illuminate knowledge in general? One might suppose that it cannot do much on this score because it relies on an unanalysed notion of knowing by sight. But there is further illumination in prospect, through enquiry into the character of perceptual-recognitional abilities. For instance, it is important to stress that my ability to tell that the bird is a robin from its look is not simply a matter of my being prepared to judge of things that they are robins in response to having visual experiences of a certain sort. There are possible worlds in which I am so prepared and yet lack the ability. For instance, I would lack the ability in a world that might be imagined in science fiction, in which there are skilfully made robots looking and moving just like robins. Environments in which I count as having the ability to tell that something is a robin from its look must be ones in which I could not easily judge falsely when I judge that something is a robin from its look. Similarly, environments in which I count as being able to tell from their visual appearance that the animals in the enclosure are zebras must be ones in which I am not liable to come across cleverly disguised mules in such places. I count as seeing that the animals are zebras in virtue of having exercised what counts as an ability to recognise zebras as zebras: knowing that the animals are zebras through seeing that they are *is* the exercise of this ability. But I count as having the ability only if, in the context in question, having the appearance of a zebra is a highly reliable indicator that it is a zebra.

When I exercise the ability to recognise zebras as zebras I need have taken no special steps to test for the possibility that things looking just like zebras are not cleverly disguised mules. This might be thought to suggest that I am blind to whether I have exercised the ability. The worry here would be that since exercise of the ability in this context depends on the obtaining of a

⁷ The view is thus in keeping with the WCK thesis defended by Quassim Cassam (2007) to the effect that when one knows that *p* there is some way in which one knows that *p*.

condition to which I need not have paid any attention, I am in no position to know that I have exercised the ability. This would conflict with the obvious fact that generally when we know by looking we are in a position to know that we know by looking and do know if we but attend to the matter. A theory that could not account for my knowing that I see that the animals are zebras when I do see that they are zebras would be a poor theory. In addressing this problem it should not be supposed that, as I look into the enclosure, all I have to go on at the second-order level at which I know that I know is the fact that I have a visual experience such that it is just as if (things looking just like) zebras are before me. It is plausible that there are second-order recognitional abilities by means of which we can tell whether or not we have in some manner perceived something to be so. This may seem startling. In the case of the ability to tell that things are zebras by sight, the recognition is a response to the relevant visual appearance. But how am I supposed to recognise that *I have seen* that the things before me are zebras? What is it by means of which I recognise this to be so? It would be quite off the mark to answer, ‘By introspection of my experience’, if introspection is conceived as an exercise of some inner sense. That we have such an inner sense is, to say the least doubtful and the phenomenology suggests that the only looking taking place as I look at the zebras is the looking that is directed outward, towards them. There is, fortunately, a plausible explanation of how we tell what we see to be so that does not draw on the assumption that we have an inner sense.

It would be arch, but accurate enough, to say that I tell that *I see* that the animals are zebras from their visual appearance: the very appearance from which I tell that the animals are zebras is an appearance from which I tell that I see that they are zebras. It is not especially puzzling how this can be. Just as I have acquired a way of recognising zebras to be zebras that is triggered by episodes in which zebras are noticed and in clear view, so I have acquired a second-order ability—a way of recognising that I know in virtue of seeing that the animals are zebras—that is triggered *by those same episodes*. Both abilities involve responding to the way things look. It is just that the second-order ability is a matter of responding with a judgement that is both about the zebras and how I stand to them. To count as having such an ability, and as exercising it, the context has to be favourable, just as it has to be favourable if I am to count as having the corresponding first-order ability and as exercising it.⁸ The general message, then, is that

⁸ There is a general tendency in epistemological theorising to assume that we have, as they say, reflective access to ‘inner’ states like having an experience such that it looks just as if a zebra is there but not to factive states like seeing

perceptual recognitions of something as being identical to X, or as being an F, or as being G, are not blind, or not typically so, because at least in normal cases the subject will have recognised, from the appearance that is in view, the perceptual mode by which the perceptual recognition has been effected.

At this stage it might seem that there is an outstanding problem of making sense of the idea that I can discriminate between zebras and things that aren't zebras. There may remain a temptation to suppose that what is available to me does not enable this discrimination to be made. This thought depends on the model that is here being rejected. It assumes that what is available to me when I look into the enclosure is the same as what would be available to me if the animals were cleverly disguised mules. What we need to restore is the common sense thought that what one can tell from an appearance depends not just on the appearance but on characteristics of the environment in which the judgement is being made. It is possible that in the context of the zoo, the appearance of the zebras discriminates between the zebras and non-zebras: possibly, nothing one might find in that context would have the same appearance.⁹ Note too that on the proposed account knowledge is not severed from the possession of reasons and thus justification. Thanks to second-order abilities of the sort just described, facts as to what one sees can be available to us to serve as reasons for believing this and that. To vary the example, when I spot Bill on the other side of the road I know that I see him, there and then, and at least for a short time I shall retain the memory of having seen him then. At the time, I acquired a reason to believe that he was there then: a reason constituted by the fact that I saw that it was Bill. Subsequently, for a while, I retain that reason. This is the reason I would naturally give by way of providing a justification for my believing that he was there. Similarly, the fact that I have seen that the animals in the enclosure are zebras can constitute a reason to believe that there were zebras in that enclosure. That reason does, of course, discriminate between zebras and cleverly disguised mules since it would not obtain had the animals been mules.

that zebras are before one. The tendency should, I believe, be resisted. There is no reason to think that recognising that one sees that a zebra is there is any more puzzling than seeing that a zebra is there.

⁹ Asked to say whether one might find cleverly painted zebras in the zoo, the zoo director is rightly going to say, simply, 'No'. He is not going to say, 'Well of course, there is always a possibility that something funny is going on'. Were he to contemplate such a possibility he would be bound to do something about it. What might happen is highly context-sensitive. The context can be such that it is just not the case that the animals might be disguised mules. This is a theme of Travis 2005.

The upshot is that, even if we eschew reductive conceptual-analytical accounts of knowledge, there is the prospect of shedding light on perceptual knowledge through consideration of the interplay between perceptual-recognitional abilities, perceiving that something is so, and having justification in the form of reasons for beliefs. We do not give up substantive epistemology by giving up the traditional conceptual-analytical project. That is by way of defence of the general approach to perceptual knowledge. But the interest of this approach for present purposes lies in the contribution it makes to enabling us to make sense of the common sense contrast between knowing that something is so and being in lesser states implicating true belief, even if in forming the lesser state we rely on evidence that warrants a fairly high degree of confidence. The key problem was to understand how in cases of knowing the matter is settled. I have been addressing this for the case of perceptual knowledge. When I know that something is so by means of seeing that it is so, the matter can be settled because One knows in virtue of seeing, via the exercise of a recognitional ability. The exercise of such an ability may well on its own induce assurance that the fact obtains. In any case, such assurance is made secure through the recognition that one knows by seeing and the retention in memory the fact that one knows by seeing. I turn now to consider knowledge that is based on evidence.

4. *Telling from indicator facts*

As we have seen, weak conceptions of evidence obscure the distinction between knowledge and true belief based on evidence that warrants a fairly high degree of confidence. It might easily seem that there is no hope of making sense of a sharp distinction between knowledge and the supposedly lesser state. But there certainly are cases of knowledge based on evidence in which some fact that constitutes evidence that p stands in a very tight relation to its being true that p .¹⁰ It can be that some fact indicates that p , where indication is understood as a factive notion: *the fact that p indicates that q entails q* . Thus the fact that there are such-and-such rings on a tree

¹⁰ There are strands in epistemological tradition in which it is acknowledged that the concept of knowledge is more demanding than accounts of justification incorporating a weak conception of evidence can explain. Fred Dretske (1971) argued that for knowledge there must be reasons for belief that are conclusive. John McDowell has protested against accounts of knowledge of other minds that envisage ‘ascribing knowledge on the strength of something compatible with the falsity of what is supposedly known’ (1982: 372). Charles Travis (2005) has commended a view on which knowledge requires proof—not necessarily mathematical proof but, nonetheless, something far stronger than merely evidence giving what is supposedly known a probability less than 1.

trunk may indicate that the tree was so many years old. The fact that a thermometer reads 97 degrees may indicate that a patient has a temperature of 97 degrees. The fact that the rotor blades of a helicopter in flight have ceased to function indicates that the helicopter will fall to the ground. In such cases I shall call the fact that does the indicating the *indicator fact*. Knowledge based on evidence supplied by an indicator fact is like perceptual knowledge in this respect: in both cases the subject deploys a way of telling that something is so. In the perceptual case the way of telling is the deployment of a perceptual-recognitional ability. Knowledge that p from the indicator fact that q is knowledge resulting from a way of telling this sort of thing from that sort of fact. In both cases whether one has the ability is sensitive to contextual factors. There might be a way of telling that sheep are just over the hill from the fact that there is a certain sort of bleating sound, in a context in which one could not easily judge falsely that sheep are just over the hill from such a sound. That is compatible with the possibility that in other contexts there might be goats that produce the same sound.¹¹ Just as it would be a mistake to suppose that in the zebra scenario I rely on evidence that does not settle the matter, so it would be a mistake to suppose that in this sort of case the evidence comprising the fact that there is this bleating sound would never settle the matter. The key is to acknowledge that what is settled by what is context-sensitive.

What entitles subjects to rely on the indicator fact in the cases under consideration. What entitles the shepherds, for instance, to treat the fact that there is the bleating sound as indicating that sheep are nearby? We may suppose that the shepherds draw a conclusion in keeping with the following inference schema.

- | | | |
|-----|----------------------------------|--|
| (A) | q | (From perception) |
| (B) | The fact that q indicates that p | (Connecting assumption) |
| (C) | p | (From A, B and factivity of indication.) |

For this schema to be part of a plausible explanation of how a subject knows that p, it need not be assumed that anything like a process of reasoning leads to the knowledge that p. What matters

¹¹ The example is from Travis 2005.

is that the subject should have arrived at the judgement that p in view of the indicator fact that q. They are able to do so because they know the significance of the indicator fact. That is to say, they know what it indicates. The question is how to explain this knowledge.

To vary the example, suppose that the indicator fact is that tracks of a certain configuration are on a muddy path and that what is indicated is that deer have recently passed along the path. A common way to approach our problem would be to say that to gain knowledge that deer have recently passed from the presence of the tracks I must have evidence for accepting a generalisation linking the presence of such tracks to the recent passage of deer. My evidence might derive from my personal experience of deer and their behaviour, but if the view is to be plausible it had better allow that justification for accepting the generalisation may derive from the testimony of others. There is much in epistemological tradition that encourages us to think that there must be some plausible story to tell along these lines. I think we should be suspicious of this general approach.

The way of thinking just outlined builds into knowledge from indicator facts an ingredient that looks as if it is, at best, and for many of us, going to rest on inconclusive evidence. That is because many of us do not possess evidence that does more than warrant a fair degree of confidence in the required covering generalisation. There is a question whether support for an empirical generalisation can ever be conclusive. What matters for present purposes is that many of us who undoubtedly have knowledge-acquisition abilities of the sort under consideration do not possess conclusive evidence in support of the required generalisations. We work with a picture of the world, which experience by-and-large does not contradict, on which this or that instrument is very likely to be reliable, this or that type of sign reliably indicates that this or that is so. But we shall have done nothing to test the various elements of this picture. We trust the instruments we use with little if any understanding of how the more sophisticated ones work and having done nothing to test their reliability or the reliability of the sources that tell us how to read them.¹² These considerations raise a problem for the perspective being developed in the present discussion. We want to make sense of the idea that we can obtain knowledge from

¹² Compare Sosa 2006: 117-18: 'Take the gauges that we face as driver of a late-model car. Most of us have a paltry conception of them as little more than screens, displays, that keep us informed about the amount of fuel in our tank, our speed, the rpm of our motor, etc. We take the display to be part of a fuller instrument that reliably delivers its deliverances, But who knows how the display on our dashboard reliably connects with its relevant subject-matter. Our conception hardly extends beyond the distinctive screen or display.'

indicator facts. We seek to do so in keeping with our acknowledging the contrast between knowing and merely having a true belief on the basis of evidence warranting at best a fairly high degree of confidence. But, when we look into what enables us to know that this indicates that, we seem to have to acknowledge that in many cases generalisations for which we do not have strong evidence play a significant role.

One way of responding to the problem would be to deny that in order to know that p from the fact that q we need to know that the relevant connecting assumption is true.¹³ That amounts to denying that the subject need know the significance of the indicator fact. I shall not attempt to refute such a view, though it does seem to make it mysterious how in the absence of knowledge that the fact that q indicates that p one could count as knowing that p . I shall instead propose an alternative account.

When we can tell that our fuel tank is half full from the reading on the fuel gauge we have mastered a simple technique: judging roughly the quantity of fuel in one's tank from the reading on the fuel gauge. We learn that this is what you do to find out roughly how much fuel you have. When we get the hang of it, which in this case is not overly complicated, we have ability to tell roughly the quantity in our fuel tank from the readings on the gauge. The crucial test for whether people have acquired the ability is whether they are good at correctly judging roughly how much fuel they have from the reading on the gauge. The test does not involve a judgement about the subject's evidence for treating readings on fuel gauges as indicators of level of fuel. That is not because we tacitly assume that they have enough evidence on this matter. It is because whether the technique has been mastered, depends on what they can do. In this respect, it is like the ability to add up a column of figures. I suggest that exactly the same points may apply to abilities to tell that deer have passed from the presence of tracks on a path, to tell that a car skidded from tyre marks on a road, to tell that sheep are nearby from the sound of bleating, and to tell the approximate time from a display on one's watch.

On the proposed view the acquisition of knowledge from indicator facts is the exercise of something that is akin to a practical skill or ability. We judge whether or not people have a practical ability to do something by telling whether they can reliably do that thing. Practical abilities rely on the world's being a certain way. For instance, my ability to hit nails into wood

¹³ I take it that this is the approach of Dretske 1971. Compare also his treatment of perceptual knowledge in Dretske 1969, ch. III.

without bending the nails, depends on the nails and the wood having fairly consistent properties that facilitate the manipulations that exercising the ability involves. I presuppose that this is so, though I may do nothing to check on it when I exercise the ability. Yet when I exercise the ability I do not do so blindly. It is not that I go through some routine that reliably leads to my having hammered in nails in a fairly neat way and have no idea what I am doing or how I am doing it. I know what I am doing and there is a constant interaction between my actions with the hammer and the feedback I glean about how the business is progressing. For instance, I have to avoid attempting to drive nails into knots in the wood, and I need to take care over how I strike the nails. Perceptual-recognitional abilities and abilities to recognise the significance of indicator facts are in some respects akin to this practical skill. Subjects who have those abilities presuppose that the world is such as to facilitate possession of the ability. Exercising the abilities does not turn on whether the subject has made checks to determine that the world is that way. Yet typically those who tell some something is such-and-such from its appearance, or tell that something is so from some perceived indicator fact, will know what they are doing. When I judge that it's a robin at the bird table, I know that I am telling that it's a robin from its appearance. If I need to observe more closely to be sure then there will be an interaction between how I observe and the feedback I receive through observing. I shall judge that it is a robin only when I get a good enough look.

When the indicator facts are perceived facts, recognition comes in at two levels: first, in the acquisition of knowledge that the indicator fact obtains and second, in the recognition of the significance of the indicator fact. Talk of recognition at the second level seems entirely appropriate. I recognise the significance of the fact that my fuel gauge reads half full. I take in two distinct facts: that the gauge displays a reading of half-full and that this fact indicates that the tank is half full. But there is phenomenological immediacy all the way through. No reasoning is involved even though the judgement about the fuel level depends on recognising the significance of the reading on the fuel gauge. In some cases the indicator fact is conceptualised in such a way that it entails the fact indicated. That there is smoke not only indicates that there is fire; it entails that there is fire. When this is so the only difference from the model of perceptual knowledge considered earlier is that a judgement, the content of which is the content of an indicator fact, is always involved. In perceptual cases there need be no judgement whose content is the content of an indicator fact pertaining only to appearances or to superficial features making up appearances.

Abilities to recognise the significance of perceived indicator facts are to be understood along much the same lines as the perceptual-recognitional abilities considered earlier. Whether one counts as having such an ability, or counts as having exercised it, is context-sensitive. Suppose that some place there are goats that make a bleating sound just like the bleating sounds of the goats around here. In that place our shepherds would not be able to tell sheep from goats just from the bleating, but that does not mean that they can't do so around here. Moreover they can know that they know on occasions on which they exercise their ability. They tell that sheep are nearby from the sound of bleating. From the sound they know how they know; they deploy second-order ability that enables them to recognise how they are telling. They have mastered the technique of telling how they know.

The proposed view yields an account of knowledge of appropriate connecting assumptions. Recall that the schema for connecting assumptions is that the fact that p indicates that q . The view is that what entitles subjects to connecting assumptions is that they recognise the significance of the relevant indicator fact. This is what makes knowledge from indicator facts akin to straightforward perceptual knowledge. In both cases a recognitional ability is exercised. In the indicator cases under consideration knowledge of the connecting assumption is recognitional knowledge of the significance of the indicator fact. This knowledge is borne of the ability to detect that significance—an ability, acquired through appropriate learning, to tell this sort of thing from that sort of fact.

It is crucial for the view that in virtue of having exercised the requisite recognitional ability subjects can be entitled to accept that this fact indicates that such-and-such. We explain the entitlement in terms of the recognitional knowledge of this fact, rather than explaining the recognitional knowledge in terms of prior entitlement to accept the connecting assumption, deriving from evidence-based entitlement to accept some generalisation. The shepherds will have recognitional knowledge of the significance of the indicator fact that there is a bleating sound only if a general assumption to the effect that the occurrence of this sound is a highly reliable indicator of the presence of sheep holds at least *with respect to the environment in which the relevant technique is applied*. No doubt they will have a picture of the world on which something like this general assumption holds. But for them the content of any such assumption need amount to nothing more than what they might convey by saying that you can tell that sheep are nearby from the occurrence of the bleating sound. Two important points emerge at this point.

The first is that it might not be that what the shepherds would convey by saying this need not be that always and everywhere you would be able to tell that sheep are nearby from this bleating sound. That would suggest a far more determinate picture of their world than is necessary for them to be able to tell that sheep are nearby from the bleating in the region in which they exercise this ability. For this reason alone it would be a mistake to suppose that having the ability must depend on prior entitlement to some determinate generalisation relating the occurrence of the sound to the nearby presence of sheep. The second point is that even if the shepherds believed that always and everywhere you would be able to tell that sheep are nearby from the bleating, and even if this belief were strictly false, because somewhere there are goats indistinguishable by their bleating from sheep, this would not undermine the claim that they can tell that sheep are nearby from the bleating.

5. Overestimating the role of evidence in justifying our beliefs

One of the factors that sustains the conception of knowledge from indicator facts that I have been trying to replace is an overestimation of the extent to which our beliefs are justified on the basis of evidence. It is assumed not to be problematic that we have evidence for generalisations covering connecting assumptions bearing on knowledge from indicator facts. Independently of the conception of knowledge from indicator facts that I have been outlining, there is reason to be suspicious of theories of knowledge that too readily posit justification on the basis of evidence.

For justification on the basis of evidence there must be evidence that supports a proposition and one must believe that believes that proposition for the reason provided by that evidence. Suppose that I came to know that Ulan Bator was the capital of Mongolia because this was indicated in an encyclopedia I once read and that I knew to be a reliable source. I had good reason to believe that Ulan Bator is the capital of Mongolia, and this reason was the reason for which I came so to believe. I may continue to believe for that reason because the reason continues to sustain my belief. But now suppose that some time after learning about the capital of Mongolia in the manner described I irretrievably forget where I obtained this information, but continued to believe that Ulan Bator is the capital of Mongolia. In that case, I no longer believe for the reason that the reliable source contained the information that Ulan Bator is the capital of Mongolia. Perhaps my appreciation of the character of the source played a causal role in laying

down a firm belief. But even so I no longer believe in view of the fact that the source contained the relevant information. And it is no longer the case that my belief is justified by my having evidence constituted by the fact that the relevant information was contained in a reliable source. I had such evidence but no longer have it. Since I no longer have it, my having it cannot be what justifies my belief.¹⁴

The case I have just been considering assumes that the belief in question was once justified by possession of evidence. Arguably, in many cases of acceptance of information, we were never, properly speaking, justified through possession of evidence. Varying my geographical example again, suppose that once a teacher of mine wrote on the board, ‘Ulan Bator is the capital of Mongolia’ and that this fact sticks in my mind. The teacher was a reliable source but at that stage in my development I was in no position to determine that she was. Though her writing this on the board was evidence that Ulan Bator is the capital of Mongolia I was in no position to assess it as evidence.¹⁵ At this stage I simply imbibed this information that was given to me as I imbibed much else besides. I am not disputing that it can be right or reasonable for people uncritically to take in information that is presented to them through formal or informal education. If they were not to take that information on board they would not become equipped with a perspective on the world that would enable them to function adequately and furnish them with the means to acquire knowledge and form beliefs justified on the basis of evidence. The issue is

¹⁴ The claim here is compatible with its being relevant to whether I *know* that Ulan Bator is the capital of Mongolia, that I gleaned this from evidence supplied by a reliable source. But the evidence supplied by the source is not evidence on the basis of which—in view of which—I now have the belief. So the evidence is not evidence I now possess that supplies me with a reason, and thus justification, for this belief. This would be manifested by the fact that I would not allude to the evidence if called upon to justify my belief and that recollection of the evidence plays no role in sustaining my belief. Some may take it that accepting this view would commit us to thinking that we do not justifiably believe much of what we take to be factual knowledge, since it is no longer based on appropriate evidence. However, accepting the view incurs no such commitment. We can make sense of my having justification for my belief about the capital of Mongolia even though I have forgotten the source of the information. For my justification may derive from my distinctly remembering that Ulan Bator is the capital of Mongolia. Plausibly, it is a factor relevant to my remembering that I once learned this fact from consulting the encyclopedia. But even so my justification now derives from my distinctly remembering the fact, not from the evidence from which I learned the fact. A great deal of our factual knowledge is like this.

¹⁵ This view is compatible with diverse epistemologies of testimony. On any plausible view of testimony someone’s saying that p can be evidence for believing that p. Plausibly, someone’s saying that p can sometimes be evidence that p, and rightly be taken to be such, even though the recipient of the testimony has not taken any special steps to check the credibility of the speaker on the matter in hand. I take it, however, that a plausible theory will demand that to be justified in accepting testimony one should have recognised the speaker as someone who is credible on the matter in hand. That requires an ability to make discriminations between speakers—an ability which I did not possess in the imagined scenario.

whether the rightness of taking that information on board should be understood on the model of being justified in virtue of having evidence that justifies.

It might be suggested that although when we first accept what we are told on various factual matters we are not justified on the basis of evidence, we have a great deal of experience subsequently which supplies us with ample evidence for what we believe. From this it might be inferred that when we are justified in believing what we were once told, our justification derives at least in part from our possessing evidence subsequently encountered.¹⁶ It is certainly true that having once been told which city is the capital of Mongolia I may have read travel publications or heard news broadcasts in which Ulan Bator is referred to as the capital of Mongolia. Experience of this sort might well play a role in sustaining my belief. But again the issue is not whether it sustains my belief but whether, if the belief is justified, it is justified on the basis of evidence gleaned through such experience. We need to distinguish between being induced to believe something, or sustained in believing it, by what is in fact evidence, and being justified in believing it in virtue of having appreciated the significance of the evidence.

Each of us has vastly many beliefs and habits of expectation that, even though induced by experience, cannot plausibly be regarded as being justified on the basis of evidence. We routinely tell that someone wants to see us from the fact that there is knocking at our office door. On particular occasions the connecting assumption is that the fact that there is a knocking sound at the door indicates that there is then someone wishing to see one. As theorists we are apt to explain our entitlement to such an assumption in particular cases by invoking inductive evidence for a generalisation linking knocking sounds on doors with people on the far side of the door wishing to see us. But, notwithstanding that we have lots of experience of knocking sounds succeeded by people entering our offices, it is far from clear that we should posit a justified belief on the basis of evidence provided by such experience. Few of us will have assessed such evidence or submitted any generalisation to test. What we have done is get the hang of a practice—the practice of knocking on a door to a room to indicate that you wish to enter. Experience, including the experience of learning and engaging in the practice, has induced in us a habit of expectation. That we have such a habit is entirely reasonable. I am questioning whether its reasonableness is to be understood in terms of justification by inductive evidence.

¹⁶ A view of this sort is advanced in Adler 1994 by way of countering the view that knowledge of what we were once told depends on trust.

We work with a perspective on the world elements of which are not justified on the basis of evidence, elements of which are false. Our perspectives, though flawed, inform our thinking. Elements that are false or not justified by evidence may enter the thinking that goes into the acquisition or exercise of recognitional abilities. They may figure in the stories we tell about our recognitional abilities. (Recall the possibility that the shepherds falsely believe that the bleating always and everywhere would be from sheep.) We should not try to finesse these difficulties by trying to show that, after all, we always have justification by evidence in cases in which it is clear that we have a recognitional ability. But nor should we suppose that the flaws in our perspectives, when they bear on what we take to be recognitional abilities, undermine the claim that we have those recognitional abilities. We shall be able to make sense of this only if we give up the idea that we have show that we can be entitled, on the basis of evidence, to hold to the perspectives in order to show that we have recognitional abilities.

6. Concluding remarks

This discussion has been framed by a concern to account for the well-grounded assurance that can be gained when people know things as opposed to having a true belief based on evidential support that merely warrants a high degree of confidence. I have addressed this concern by focusing on modes of knowledge that supply us with paradigmatic cases of knowledge. These are cases of perceptual knowledge—cases in which one tells something through some mode of perceiving—and knowledge from indicator facts—cases in which one tells something on the basis of some fact that (factively) indicates that it is so. I have sought to extend the range of recognitional knowledge—knowledge acquired through the exercise of a recognitional ability—to cover knowledge from indicator facts. In both the initial perceptual cases and the cases involving indicator facts, the knowledge is not to be viewed as based on evidence that more or less strongly supports the proposition in question. In the perceptual cases the knowledge not evidence-based at all, in keeping with the observations by Austin noted in section 2. The indicator cases rely on evidence that (factively) indicates that the proposition in question is true. The view emerging seems to me to be entirely in line with common sense. But, importantly for the purposes of this discussion, it also enables us to make sense of the contrast between the sorts of knowledge I have considered and true belief based on evidence providing a high degree of

confidence. In cases of the latter sort the matter is not settled even though it is pretty unlikely that the proposition in question is false. In cases of the former sort the matter is settled: either one perceives that p or one has evidence indicating that p.

The proposals I have canvassed raise the question of how far knowledge extends beyond the kinds of cases I have considered. Can there, for instance, be knowledge of empirical generalisations. These are matters for another occasion.

References

- Adler, Jonathan E. (1994) 'Testimony, Trust, Knowing' *The Journal of Philosophy*, 91, 264-75.
- Austin, J. L. (1962) *Sense and Sensibility* (Oxford: Clarendon Press).
- Cassam, Quassim (2007) 'Ways of Knowing', *Aristotelian Society Proceedings*
- Davidson, Donald (1983) 'A Coherence Theory of Truth and Knowledge' in Dieter Henrich (ed.) *Kant oder Hegel* (Stuttgart: Klett-Cotta). Reprinted in Ernest LePore (ed.) *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson* (Oxford: Blackwell, 1989), 307-19.
- DeRose, Keith (1992) 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52, 913-29.
- Dretske, Fred (1969) *Seeing and Knowing* (London: Routledge and Kegan Paul).
- Dretske, Fred (1970) 'Epistemic Operators' *The Journal of Philosophy*, 67, 1007-1023.
Reprinted in Dretske 2000.
- Dretske, Fred (1971) 'Conclusive Reasons' *Australasian Journal of Philosophy*, 49, 1-22.
Reprinted in Dretske 2000.
- Dretske, Fred (1992) 'Perceptual Knowledge' in Jonathan Dancy and Ernest Sosa (eds.) *A Companion to Epistemology* (Oxford: Blackwell).
- Dretske, Fred (2005) 'Is Knowledge Closed Under Known Entailment' in Steup and Sosa (eds.) (2005), 13-26.
- McDowell, John 'Criteria, Defeasibility, and Knowledge', *Proceedings of the British Academy*, 68, 455-79. Reprinted in McDowell 1998.

- McDowell, John (1993) 'Knowledge By Hearsay' in B. K. Matilal and A. Chakrabarti (eds.) *Knowing from Words: Western and Indian Philosophical Analyses of Understanding and Testimony* (Kluwer: Dordrecht), 195-224. Reprinted in McDowell (1998).
- McDowell, John (1995) 'Knowledge and the Internal', *Philosophy and Phenomenological Research*, 55, 877-93.
- McDowell, John (1998) *Meaning, Knowledge, and Reality* (Cambridge, MA.: Harvard University Press)
- Millar, Alan (1991) *Reasons and Experience* (Oxford: Clarendon Press).
- Millar, Alan (2001) 'The Scope of Perceptual Knowledge' *Philosophy*, 75, 73-88.
- Millar, Alan (2007) 'What the Disjunctivist is Right About', *Philosophy and Phenomenological Research*, 74, 176-98.
- Millar, Alan (2008) 'Perceptual-Recognitional Abilities and Perceptual Knowledge' in Adrian Haddock and Fiona Macpherson (eds.) *Disjunctivism: Perception, Action, Knowledge* (Oxford: Clarendon Press).
- Pryor, Jim (2005) 'There is Immediate Justification' in Steup and Sosa (eds.) (2005), 181-202.
- Sosa, Ernest (2006) 'Knowledge: Instrumental and Testimonial' in Jennifer Lackey and Ernest Sosa (eds.) *The Epistemology of Testimony* (Oxford: Clarendon Press).
- Steup, Matthias and Sosa, Ernest (eds.) (2005) *Contemporary Debates in Epistemology* (Oxford: Blackwell).
- Travis, Charles (2005) 'A Sense of Occasion', *Philosophical Quarterly*, 55, 286-314.
- Williamson, Timothy (2000) *Knowledge and Its Limits* (Oxford: Oxford University Press).
- Wright, Crispin (2002) '(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell' *Philosophy and Phenomenological Research*, 65, 330-48.