

What is it that Cognitive Abilities are Abilities to Do?

Alan Millar

Received: 16 July 2009 / Accepted: 1 October 2009 / Published online: 20 October 2009
© Springer Science + Business Media B.V. 2009

Abstract This article outlines a conception of perceptual-recognitional abilities. These include abilities to recognize certain things from their appearance to some sensory modality, as being of some kind, or as possessing some property. An assumption of the article is that these abilities are crucial for an adequate understanding of perceptual knowledge. The specific aim here is to contrast those abilities with abilities or competences as conceived in the virtue-theoretic literature, with particular reference to views of Ernest Sosa and John Greco. In the course of the discussion, it is emphasized that the notion of exercising a perceptual-recognitional ability is a success notion: exercising such an ability is nothing less than acquiring knowledge. Even so, the view can make sense of our fallibility. It can also be defended in the face of an objection stemming from consideration of what are here called success-rate abilities.

Keywords Competence · Greco · Knowledge · Perceptual-recognitional ability · Sosa · Virtue epistemology

1 Perceptual-recognitional Abilities

The background of this discussion is provided by the idea that a specific kind of ability—a perceptual-recognitional ability—is crucial for an adequate understanding of perceptual knowledge (Millar 2007a, 2007b, 2008a, 2008b).¹ My aim here is to

¹The epistemological framework with which I am working is more fully set out in my contribution to Pritchard et al. [Forthcoming](#).

A. Millar (✉)

Department of Philosophy, University of Stirling, Stirling FK9 4LA, UK
e-mail: alan.millar@stir.ac.uk

contrast those abilities with abilities or competences as conceived in the virtue-theoretic literature. On the virtue-theoretic side, I have in mind especially Ernest Sosa (2007) and John Greco (2007, 2009). I share with virtue theorists the view that an adequate epistemology must account for the acquisition of perceptual knowledge in terms of abilities or competences. I differ from them over how these abilities or competences should be conceived.

Here are some key points about perceptual-recognitional abilities.

- (a) Perceptual-recognitional abilities include abilities to recognize certain things from their appearance to some sensory modality, as being of some kind, or as possessing some property. They also include abilities to recognize certain individuals from their appearance to some sensory modality, as being individuals specified in some way (e.g., as your brother or sister, or as that painting you bought last year). Appearances of things are the ways they look, or sound, or feel, or smell, or taste. You might be able to recognize that something is a rose from the way it looks, or that a piece of fabric is silk from the way it feels, or that a noise is that of a fire alarm from the way it sounds. To have such an ability is to be in command of a way of telling, and therefore coming to know, that something of a certain sort is so. Abilities of this sort are exercised (manifested) when and only when one acquires the relevant knowledge. If I have the ability to recognize orchestral performances as works of Mozart from the way they sound then I manifest that ability if and only if I come to know of some performance, from its sound, that it is a work by Mozart. In a similar way, if I have the ability to ride a bicycle, I manifest that ability when and only when I ride a bicycle. The upshot is that the notion of the manifestation or exercise of an ability is a success notion. This is compatible with its being the case that there are abilities to do something only a small proportion of times on which one tries. Suppose that I can throw a basketball into the basket from 10 m about 10% of the times on which I try. I exercise that ability if and only over a series of trials, each trial comprising a sequence of throws, my success rate is roughly 10%. A single success is not a manifestation of my ability. A near miss is not a manifestation of my ability either. Manifestation is doing what the ability is an ability to do and in this case that is to achieve a success rate of 10% in attempts at doing the thing in question. (I return to such abilities in Sect. 2 below.)
- (b) We are fallible with respect to our recognitional abilities. This is not to say that these abilities are sometimes manifested in making false judgments. That would be incompatible with (a). We are fallible in that we do not always exercise our abilities when we make a judgment of the sort that is implicated by exercises of the ability. I might be good at recognizing roses by sight yet falsely judge of something that it is a rose from the way it looks. This might be due to something affecting me—distraction, inattention, confusion—or it might be due to some unusual environmental circumstance, for instance, that the thing at which I am looking is a silk artifact that is a dead ringer for a rose. Sometimes our fallibility is indicated by our making false denials. I might be able to recognize certain stereotypical orchids as orchids, but falsely take myself to be able to recognize *whether or not* something is an orchid, and falsely judge certain flowers not to be orchids when they do not conform to the stereotype.

- (c) A necessary condition of having an ability to recognize things as being of some kind or as possessing some property is, to a first approximation, that things of that kind or possessing that property should have an appearance that is distinctive of being of that kind or having that property; similarly for recognition of individuals as being individuals specified in some way. An appearance is in the relevant way distinctive if and only if with a high degree of reliability a thing's having that appearance indicates that it belongs to the kind or has the property or is the specified individual. If, for instance, Mary has a twin who in appearance is just like her, then Mary's appearance is obviously not distinctive of being Mary.

Some clarifications and refinements are called for.

- (i) Every visually perceptible object presents appearances that vary with, for instance, its orientation, the conditions of light, and the points of view to which the appearances are presented. Things that produce a sound can sound different in different circumstances. A person's voice in a narrow tunnel sounds different from the same voice in the open air. The way things feel can vary. If you have cold hands then a warm surface feels warmer than it would if your hands had not been cold. I shall often speak simply of recognizing something from its appearance to some modality, taking it for granted that there will be some range of appearances that facilitate recognition of the thing as being of some kind, or having some property or being some specified individual.
- (ii) Telling from his visual appearance that it is my colleague who has just entered the room does not require that I take in at the level of judgment that he has all of the features that make up his appearance. Echoing J. L. Austin (1946), the red patch close to the beak of a goldfinch is no doubt salient when I spot a goldfinch but this patch aids recognition only because it is an element in the distinctive overall look of a goldfinch. Our visual-recognitional abilities are triggered by a perceived *Gestalt* presented by the thing we recognize, not by prior judgments as to the features of the thing recognized. (It is sometimes possible to identify things as being of some kind from a checklist of features but that is not perceptual recognition as conceived here.) Similar points apply to other modalities. It is possible to recognize the distinctive flavor of some Islay malt whisky while lacking the resources to judge that it has this or that feature that goes to make up the flavor.
- (iii) Things of a certain kind might have a distinctive appearance in one region though the same appearance is not distinctive of things of the same kind in some other region. In fake barn county (Goldman 1976, following Carl Ginet), the appearance that barns present to a roadside point of view is not distinctive of barns since too many structures in that territory that are not barns have that appearance. So it is not surprising that our hapless subject cannot tell of structures in that territory that they are barns from the appearance they present to the roadside. Back in his home county the same appearance might be distinctive of barns. There he might be able to recognize barns as barns from the way they look. So having the ability to recognize Fs from their appearance in one place does not entail having the ability to recognize Fs from the same appearance in another place. There might be regions in between when it is indeterminate whether one has the ability or not and thus indeterminate whether one has the knowledge that would be the exercise of

the ability. For a vast range of things that we can know perceptually such considerations do not matter. For there are vastly many things that have appearances that are to all intents and purposes universally distinctive—their distinctiveness is not domain-specific).

- (iv) Though it is convenient to speak simply of recognizing things from their appearance, the abilities we exercise when we recognize are often sensitive to the surrounding context. In someone's bathroom a dead ringer for a lemon would probably not throw us off. We know it is likely to be, or could well be, soap, so we would not without further inspection judge it to be a lemon. You might be good at recognizing a colleague from the way he walks in the places in which you routinely encounter him, even though you would not recognize him from his walk in some unfamiliar setting. That might be because you think of the walk as being distinctive of your colleague relative to the familiar setting but you are not confident that it marks him out in a larger population and so would withhold judgment.

By reference to a relevant perceptual-recognitional ability we can explain why a person knows something perceptually. Sally might recognize a voice as that of her husband. Her possession of the relevant recognitional ability explains how she is *in a position* to know on this occasion that the voice is her husband's. Her exercise of that ability explains how she comes to acquire the knowledge. It also provides the basis for an explanation of what makes the state that she comes to be in knowledge rather than anything else. It is knowledge just because the requisite recognitional ability was exercised and its exercise is nothing less than the acquisition of knowledge. That is a natural way to explain what her knowing amounts to. And it is informative. The explanation is not merely that she knows that the voice is her husband's because she has exercised some suitable recognitional ability, whatever that might be, but more specifically that she comes to know through this specific recognitional ability—the ability to tell, from its sound, that a voice is that of her husband. Unusual though it would be, she might have come to know the same thing by other means, for instance, by being told that the person whose voice is salient is her husband.

My description of this imagined case of knowledge seems to me to accord well with our ordinary common sense ways of thinking about knowledge. Noticeably, it says nothing about belief. It is not that belief has no place in the fuller story. Just because she knows that the voice is her husband's she believes that it is. Her knowing in the way she does is one mode of believing. We need not think of her believing as a mental state that is an ingredient of her knowledge. We account for her believing in terms of its being entailed by her knowing. Because we explain the knowing in terms of her perception of the voice and the exercise of the recognitional ability, rather than in terms of justified belief and further conditions, we are free to account for the justification of her belief in terms of something the knowledge makes available. More specifically, we can account for the justification of her belief in terms of her possession of a reason to believe that is constituted by the fact that she hears and recognizes her husband's voice.² Her reason is

² Though the matter is controversial in epistemology I take it that in general reasons are constituted by truths or facts that you know. (For present purposes, there is no reason to distinguish between truths and facts.) As I approach these matters, it speaks in favor of a philosophical account if it respects this assumption.

that knowledge-implicating fact by reference to which she would herself account for her taking the voice to be her husband's. It is a fact that is available to her because she, like us, has a reflective ability such that when she hears a voice and recognizes whose voice it is, she can tell that she has done so. On this account, her knowing that the voice is her husband's, in virtue of hearing and recognizing his voice, is not explanatorily posterior to her possession of a reason to believe that the voice is her husband's and to her being justified in believing this in virtue of having this reason. As I understand it, the perceptual-recognitional abilities account of perceptual knowledge is a package that incorporates the conception just outlined of how perceptual knowledge relates to justified belief.

2 A Contrasting View of Cognitive Abilities

It is crucial to my account of perceptual-recognitional abilities that such abilities are of their very nature knowledge-acquisition abilities. The exercise of such abilities is nothing less than the acquisition of knowledge. Successful performance is getting it right and what getting it right consists in is determined by what the ability is an ability to do. In the case of perceptual-recognitional abilities the ability is an ability to tell, that is, come to know, that something is such-and-such from its appearance to some sensory modality. Success consists in doing precisely that. On this way of thinking there is no gap between manifestation and success. The manifestation of an ability is the subject's doing what the ability is an ability to do. There is neither manifestation without success nor success without manifestation.

Virtue epistemologists tend to think of the abilities that account for knowledge-acquisition as belief-forming abilities and to think of success in believing in terms of true belief however it is formed. Under virtue theory we can have true belief—thus success—the formation of which is not due to the right kind of ability, that is, the kind that accounts for knowledge when suitable conditions are satisfied. There can be a manifestation of an ability of the right kind in the formation of a false belief, thus without success, and in the formation of true beliefs in Gettier cases. So there can be manifestation without success (conceived as true belief) and success without manifestation, and both success and manifestation when there is no knowledge. That is why virtue theorists insist that for knowledge the true belief must be sufficiently due to the subject's manifesting the right kind of ability. John Greco (2007, 2009) sums up the view by saying that knowledge is achievement, that is, success due to ability. Ernest Sosa puts the point by saying the knowledge is apt belief, that is, true belief that is due to competence (adroitness) (Sosa 2007: Lecture 2).

On this virtue-theoretic view what is it that the ability that explains knowledge-acquisition is an ability to do? One view is that it is a belief-forming disposition such that, in appropriately normal conditions, manifestations of this disposition would be, or would be likely to be, true beliefs.³ That suggests that the ability is an ability to form true beliefs. So, for instance, the cognitive competence implicated in competent

³ Here is Sosa's formulation: a cognitive competence is 'a disposition resident in the agent, one that would in appropriately normal conditions ensure (or make highly likely) the success of any relevant performance issued by it' (2007: 28).

applications of the concept of being red might be a disposition to believe of perceived surfaces that they are red when they are seen to appear red and one has no reason to suppose that conditions are abnormal. In appropriately normal conditions manifestations of this disposition—beliefs concerning some perceived thing to the effect that it is red—would be, or would be likely to be, true. Still, the disposition could be manifested in forming a false belief in circumstances in which, unknown to the agent, the conditions of light are not suitable for recognizing color. But now consider Sosa's kaleidoscope example in which a subject is said to manifest competence in correctly believing a surface to be red even though the setup is such that a trickster is capable of manipulating both which surfaces are visible and the conditions of ambient light. Let us suppose that though occasionally red surfaces appear red in conditions of light suitable for judging color, and similarly for other colors, the trickster's manipulations are such that the color-appearance presented by surfaces is not a reliable indicator of their color. Sosa holds that despite these conditions it can be that the subject exercises the relevant competence when correctly believing a surface to be red from its visual appearance. This is because it can be that the conditions for viewing and discerning the color of a surface are normal (Sosa 2007: 35–36). For at least two reasons this strikes me as puzzling.

- (a) On the occasions of correct belief the conditions of light are suitable for perceiving the colors of surfaces (cf. Sosa 2007: bottom of p. 35 to the top of p. 36) and at some level of description the subject judges in the same way as he would judge in a normal environment. He follows the procedure: judge *x* to be red if it appears red and you have no reason to think conditions are abnormal. Nonetheless, the subject is no good at forming true beliefs with respect to the colors of surfaces within the setup. That tells against regarding any true beliefs as to the colors of surfaces in this setup as manifestations of competence at forming true beliefs as to the colors of surfaces.
- (b) Why suppose that the conditions prevailing in the setup are appropriately normal for the exercise of the relevant competence? For Sosa the answer is that the subject exercises his normal color vision in conditions that are in fact suitable for discerning color (2007: 35–36; cf. 32–33). But while there is no denying that the subject sees the color and judges correctly in direct response to what he sees, the question is whether the conditions are such that the subject can be competent at forming true beliefs as to the colors of surfaces in the setup. It seems plausible that they are not because the trickster ensures that apparent color does not reliably indicate color (so that we have a scenario structurally analogous to the fake-barn scenario). Even if we think of the relevant competence as a competence at forming true beliefs of a certain sort, rather than as a knowledge-acquisition competence, the conditions stand in the way not just of the subject's manifesting competence but of his possessing competence with respect to the colors of surfaces in the setup.⁴

⁴ Does it alter the situation if, over a series of occasions for judging, the true colours appear but could easily not have given the whims of the trickster? I think it does if the series of occasions is extensive enough for testing the subject's success-rate. The subject's competence at judging correctly would in that case be fragile because it could so easily be undermined without the subject's knowledge but he would be good at judging correctly over the occasions in question and so could be said to manifest the relevant competence.

One might respond to (a) by suggesting that the relevant competence *is* manifested because, (i) the subject has the relevant competence, and (ii) the subject judges in the same way as he does in normal environments, doing what would indisputably be a manifestation of competence in those environments. But it is not straightforward that these add up to the subject's manifesting competence. There is an evaluative dimension to competence: being competent at ϕ ing is being good enough at ϕ ing and being good enough at ϕ ing with respect to some environment is being good enough at ϕ ing there. The victim of the trickster is not good enough at correctly judging colors within that setup. Over a series of trials he would too often judge incorrectly. That is a good reason to resist the attribution to him of manifestation of competence when he correctly believes the surface to be red. The same considerations tell against the supposition that the prevailing conditions are appropriately normal. They are normal in the sense that they are such that apparent color is true color and the subject is well placed to see true color. But we are looking for normality with respect to competence at correctly judging color by sight. There is a clear sense in which with respect to that competence the conditions are not appropriately normal.

A somewhat different view has it that competences are relative to suitable environments in a sense that entails that they can be manifested only with respect to such environments.⁵ This is the view taken by Greco (2007; 2009). In its terms we deny that the victim of the kaleidoscope trickster manifests competence when correctly judging the color of surfaces within the setup.⁶ This is much closer to my own view. Still, Greco's view differs from mine in an important respect. He thinks of the abilities responsible for knowledge-acquisition as abilities reliably to form true beliefs. We are to take it that such abilities can be manifested in occasionally forming a false belief. It can easily seem that we are forced to adopt such a view if only to accommodate our fallibility. On any plausible view the ability I manifest when I come to know from its appearance that a flower I am looking at is a rose is such that possession of the ability is compatible with occasionally incorrectly judging things to be roses on the basis of their visual appearance. It might seem, then, that to account for this we have to think of the ability in question as an ability that, by-and-large, yet not always, is manifested in forming true beliefs and sometimes is manifested in forming false beliefs. (It would be manifested in forming a false belief if one were to form it in the same way, at some appropriate level of characterization, as one does when the manifestation is formation of a true belief).⁷ But consideration of fallibility does not compel us to take such a view. Suppose that, in keeping with the perceptual-recognitional abilities account, we think of the ability in question as an ability to tell, that is, come to know, that something is a rose from its visual appearance and, as is entirely natural, take this ability to be manifested

⁵ One might hold that competences can be relative to environments in a weaker sense. Suppose you think that cognitive competences are belief-forming dispositions. You might then think that such a disposition is a competence only in virtue of the fact that there is some type of environment in which the disposition would reliably yield true beliefs, but you might not demand that the competence be manifested only in such an environment.

⁶ I infer this from Greco's treatment of the fake-barn scenario (e.g., in 2009: Sect. 2).

⁷ Sosa's condition C (2007: 33) presupposes that exercises of a competence relevant to knowledge-acquisition can be the formation of a false belief.

when and only when one tells of something, from its visual appearance that it is a rose. We can still make sense of our fallibility as explained under (b) of Sect. 1 above. Our fallibility with respect to this ability is exhibited by our not always exercising it when we aspire to (when from its visual appearance we judge something to be a rose) but not in our sometimes manifesting it successfully, in forming a true belief, and at other times manifesting it unsuccessfully. When I judge the silk rose to be a rose I thought I was telling that it is a rose. I was wrong. Since I did not tell that it was a rose I did not manifest an ability to do so. This is in keeping with what strikes me as a common sense thought: an ability to do something is exercised or manifested if and only if one does what that ability is an ability to do.⁸

Resistance to the preceding claim might derive from reflection on performances like throwing a basketball into the basket or throwing a dart into a dartboard within the 25 ring. Is it not just obvious (i) that these are things that can be done, (ii) that doing them can manifest ability, and (iii) that ability to do them is more or less reliable in that exercises of it are sometimes acts of doing what the ability is an ability to do and sometimes not? In terms of my eminently fallible darts performance let's consider each of these points in turn. First, since I do sometimes hit the board within the 25 ring I can. Here, the sense in which I can do this is simply that it is possible that I do it. That it is possible obviously does not suffice for my having an ability to perform this feat. Second, when I do hit within the 25 ring it is often not entirely accidental that I do. I am trying to do so and, fairly often, make arm movements that significantly raise the chance that I'll hit within, or close to, the 25 ring. But my ability is such that to represent it, without qualification, as an ability to hit the board with the 25 ring is at best misleading. Asked if I can read English I can appropriately answer by saying, without qualification, that I can. Asked if I can hit the dartboard within the 25 ring it would be a mere display of bravado to say without qualification that I can. If one does say this it can mean little more than *This is among the things I sometimes manage to do*. An appropriate answer would be, 'I sometimes do this. I'll try'. That answer would be bizarre in the response to the question whether I can read English. It would not be bizarre from someone learning English but then, accordingly, it would be inaccurate to represent such a person as having an ability to read English. The upshot is that my ability is an ability to hit the board with the 25 ring within some range of proportions, say (optimistically), between 10 and 15% of the times on which I try. Third, the ability so understood is

⁸ The manifestation principle under consideration is that the ability to ϕ is manifested if and only if one ϕ s. One might think that an ability to jump over a bar 50 cm high is manifested if one jumps over a bar 80 cm high, and therefore that one can manifest an ability when doing something other than that which the ability is an ability to do. (Something like this was suggested to me at the Bled Epistemology conference in 2009 by Adam Carter.) Aside from the example, the principle admits of a plausible refinement. Let us say that an ability to ϕ implicates an ability to ψ if and only if necessarily, if one has the ability to ϕ then one has the ability to ψ . Then the right principle might be that the ability to ϕ is manifested if and only if one either ϕ s or does something the ability to do which is implicated by the ability to ϕ . It is not plausible that the example in question forces this modification since it is not plausible that the ability to jump over an 80-cm bar (ability A) implicates the ability to jump over a 50-cm bar (ability B). Some weird contingency might prevent one who has ability A from having ability B. Where there is no such contingency one might still feel there is some sense in which a manifestation of ability A could be a manifestation of ability B. But this sense might just be that some manifestations of ability A as a matter of fact indicate that the subject has ability B.

manifested when and only when over a series of throws, say 100, I hit within the 25 ring within the specified percentages. Doing that, and nothing less than that, is what exercise or manifestation of the ability consists in. So while it is true that when I try to hit within the 25 ring I shall sometimes succeed and often will not, it is not true that single successes are manifestations of my ability or that any failures are. Nor is it true that the ability is more or less reliable. If it really is an ability as specified I shall be just as reliable with respect to it as I am with respect to my ability to read English in that, barring interfering factors, it will be manifested always or nearly always when I undertake a run of attempts. (Note that hitting within the 25 ring is success only in that it is what I was attempting. That does not guarantee that it was success in the sense of an exercise of an ability to do such a thing).

Residual doubts about the foregoing might derive from the thought that it will often be the case that when I hit the 25 ring, or even when I don't but strike close to it, my shot manifests ability. It can certainly be the case that a single hit within the ring, or a near miss, is not entirely accidental because it derives from an ability to throw in a manner that significantly raises the chance that the dart will hit within or close to the ring. That ability can be manifested by a single shot—one such that the arm movement does significantly raise the chance that the dart will hit within the ring. When the arm movement is of that sort it will not be entirely accident if the dart hits within or close to the ring. But the ability in play here is not an ability to hit the dartboard within the 25 ring.

An ability to ϕ can be a success-rate ability—an ability to do something else, say ψ , a proportion of the times on which one tries. The problem we have been addressing arises from thinking of success-rate abilities as if they were abilities to do the other thing, for instance, thinking of the ability to hit the dartboard within the 25 ring a certain proportion of the time as if it were an ability to hit the board within the 25 ring. The abilities to read English, to speak French, to ski, to ride a bicycle, to find one's way from home to the campus, are not success-rate abilities. Nor are perceptual-recognitional abilities.

3 The Grain of Ability Specifications

Fake barn scenarios are ones in which a person aims to exercise an ability to tell of structures that they are barns from the way they look but does not succeed in doing so. By my account the reason is perfectly straightforward. You cannot tell of the structures in the environment in question that they are barns on the basis of appearances presented to the relevant points of view. This might lead one to suggest that when we specify an ability simply as an ability to tell of certain structures that they are barns from the way they look our specification stands proxy for a more specific specification that makes reference to a particular environment. The matter is not quite so straightforward and requires some sorting out.

A simple ascription of an ability, like the ability to ski, leaves a great deal unspecified about the ability ascribed. The incorporated specification covers everything from the novice's ability to ski only on gentle slopes to the downhill racer's ability, which enables him or her to negotiate challenging slaloms. Those abilities differ qua abilities. We can specify them more precisely: the novice has an

ability to ski on gentle slopes; the downhill racer can ski pretty well anywhere there are conditions that make skiing practically possible. Their abilities differ in that one has an ability to ski in one limited sort of environment and the other an ability to ski in a very wide range of environments. Suppose it is suggested that we do not need to think of these abilities as relative to types of environment. The alternative would be that there is the ability to ski (and that's it), but some can exercise that ability—ski—in some specified narrow range of conditions and others can exercise it in a specified wider range of conditions. The trouble with this way of thinking is that the novice's ability differs *qua* ability from the downhill racer's ability and that difference just is a difference with respect to the conditions in which the subject can ski. If we represent the abilities to be constitutively independent of the conditions in which one can ski we cannot adequately specify the difference. We need to, and we can easily, make sense of the idea of an ability as being relative to, or with respect to, some environment or type of environment.

It does not follow from these considerations that the ability-specifications with which we routinely work stand proxy for specific descriptions specifying either a particular environment or type of environment. That a person has an ability presupposes that there are conditions suitable for its exercise. Call these *favorable conditions*. In most of the contexts in which we attribute abilities favorable conditions either obtain or are not hard to come by in the agent's sphere of operation. Those attributing the ability will understand that its exercise requires favorable conditions but might have a rather vague idea of what those are. The mere attribution of an ability, under a description like 'the ability to ski' or 'the ability to speak French', entails nothing very specific about which conditions are favorable for its exercise. Evidently to ride a bicycle the conditions have to be such that it is possible to maintain balance and the person riding must be able to maintain balance. But the claim that someone can ride a bicycle is not shorthand for a longer statement to the effect that the person can ride a bicycle provided that this, that or the other condition obtains. That said, some attributions of abilities are fairly specific, as when a person is said to be able to ski on gentle slopes, and attributions that are not specific can sometimes be made more specific.

Our routine attributions of perceptual-recognitional abilities are made against a background in which it is assumed that those to whom we ascribe the abilities, without qualification, operate in environments favorable to the exercise of those abilities. For a given ability the environment of operation—the environment in which the ability is, or is liable to be, exercised—will normally be favorable in that the things one has the ability to recognize have a relevantly distinctive appearance. Examples of the fake barn-type are ones in which the environment of operation is not favorable to the exercise of the relevant ability; hence knowledge is not acquired. But we should not infer from this that our specifications of perceptual-recognitional abilities routinely build in a specific environment parameter, so that variations in that parameter represent different abilities.⁹ Even so, examples of the fake barn type call for specifications of recognitional abilities that are more precise than we normally need.

⁹ In my 2007a and 2007b I spoke of abilities as being indexed to environments or types of environment. This is potentially misleading since it might be taken to imply that ability-specifications routinely incorporate a parameter for environment or environment-type.

If we think along the lines just set out we can meet a challenge that might be put to the account of perceptual-recognitional abilities. The challenge is that we do not in practice work with fine-grained conceptions of abilities and, in any case, are at a loss to specify with any accuracy the environments to which abilities so conceived are supposed to be relative.¹⁰ We meet the challenge by denying that the conceptions of perceptual-recognitional abilities on the proposed account are fine grained. Those abilities, like others, are abilities to do things in circumstances favorable to doing those things. Exactly which circumstances are favorable need not be built into our concepts of those abilities.

4 Reconciliation?¹¹

A possible reaction to my account of perceptual-recognitional abilities would be to concede that the account is plausible but claim that it is not at odds with more mainstream epistemological accounts in terms of true belief plus the satisfaction of other conditions (in short, in terms of true belief+). Most would concede that we have command of vastly many ways of telling things perceptually. These might include the ways exhibited when telling that one's coffee is cold from the feel of the mug containing it, telling that Bob Dylan is singing on the radio from the sound of his voice, telling that there is mint in the salad from its taste, and telling that it has been frosty overnight from the look of the grass in one's garden. It is open to virtue theorists to seek to give an account of such ways of telling in terms of competences in whichever way they conceive of them. This being so it is important to emphasize that the perceptual-recognitional abilities account does not merely highlight that we have such ways of telling. As I indicated at the end of Sect. 1, it is a package that incorporates a conception of the relation between knowledge through perceptual recognition and justified beliefs. This conception is incompatible with accounts of knowledge in terms of true belief+.

On the perceptual-recognitional abilities account what makes it the case that I know that the flower before me is a rose is that I see the rose recognizing it, and in that way knowing it, to be a rose. We do not account for this knowledge in terms of having a justification for a belief but in terms of the exercise of a specified recognitional ability. The exercise of that ability *is* the acquisition of knowledge. Nonetheless, in such situations, as a reflective knower I shall normally be aware—know—as I see the rose that I see that it is a rose and in that way know that it is. When this is so this fact—that I see that it is a rose—provides me with a reason to believe that it is a rose.¹² An advantage of this account is that it enables us to preserve the very natural thought that reasons are constituted by true propositions.

¹⁰ For skepticism about whether we routinely work with fine-grained conceptions of abilities, see Duncan Pritchard's contribution to Pritchard et al. [forthcoming](#).

¹¹ This section is prompted by Ernest Sosa's emphasis in discussion at Bled 2009 that the most important respects in which our outlooks differ have to do with our differences over whether a reductive account of knowledge is possible and over the relation between justified belief and knowledge.

¹² This account does not entail that knowing that *p* entails being justified in believing that *p*. Reflective knowers with the conceptual resources to tell when they see that something is so have access facts as to what they see to be so that can serve as reasons for belief.

Regress arguments have been taken to require us to explain justified belief in cases of perceptual knowledge in terms of sensory experiences and accordingly either to sever justified belief from the possession of reasons (Pryor 2002) or posit a special class of reasons that are constituted by experience (as I did in Millar 1991). On the view I am proposing here there is no need for these maneuvers because we are no longer seeking to explain knowledge in terms of justified belief. We explain justified belief that *p* in perceptual cases in terms of a reason that is constituted by its being the case that we see that *p* and that is available to us because we perceptually know that *p* and are aware that we see, or have seen, that *p*. There is no mystery as to how reasons of the sort under consideration—factive reasons pertaining to what we see, or have seen, to be so—can be available to us. For it is plausible that we have higher-order recognitional abilities for telling what we see to be so. My first-order ability to tell that the flower is rose is a learned ability correctly to apply the concept of a rose to things that present a suitable appearance. The higher-order ability by which I tell that I see that *this* flower is a rose is a learned ability correctly to apply the concept *see that ... is a rose* to me and a rose when by seeing I demonstratively pick out that rose, visually recognizing it to be a rose. The higher-order ability is as much dependent on a favorable environment as the corresponding first-order ability. Set me in a fake rose scenario, where visually rose-like things could easily not be roses, and I would be no more able to tell of something I see that *I see that* it is a rose than I am able to tell that it is a rose. This has an impact on what reasons I can have to justify my beliefs. In a fake-rose scenario, there are no facts to the effect that I see that something at which I am looking is a rose, even when I am looking at a rose. Evidently in that case I lack reasons constituted by such facts. Seeing that something before me is a rose entails visually recognizing it to be a rose, and visually recognizing something to be a rose is telling that it's a rose from its visual appearance. In a fake rose scenario I cannot tell that anything is a rose from its visual appearance, since no such appearance is distinctive of roses. Here as elsewhere, whether I have certain reasons, and the kind of justification that these reasons provide, depends on what is true and on what I can know.¹³

Acceptance of the perceptual-recognitional abilities account does not merely reflect a preference for a certain level of description; it is incompatible with any theory that accounts for perceptual knowledge in terms of a conception of justified belief or of manifestation of competence that does not implicate the concept of knowledge. Virtue-theory falls into this camp.

It is striking that our pre-theoretical thinking about perceptual knowledge is so naturally couched in terms of locutions like, 'telling that', 'recognizing that', 'seeing that', and the like, by contrast with talk of belief-formation, of the justification of belief, and indeed of the justification of belief provided by experiences. To take a banal example: suppose that I look out of the window to determine whether it has stopped raining. I see that it has stopped and I tell others accordingly. I can do so responsibly because I have personally checked and understand how I have done so.

¹³ We are not barred from acknowledging that belief based on false assumptions can be reasonable. It can be reasonable because so far as the subject is concerned the situation is indiscriminable from one in which the subject has a reason constituted by truths. But a reasonable belief need not be a belief that is well founded.

The others readily accept that the rain has stopped on my say so, and can do so responsibly. After all, they know me and they know that I can tell this sort of thing in that sort of way. It would be far from natural for me or the others to think of what is going on in this situation at the level of forming justified beliefs. The natural setting for talking about beliefs is where people believe but do not know, even if they have some reason to think what they believe to be true. Where someone is clearly able to tell perceptually, that is, perceptually determine (establish, settle, come to know), whether something is so, it would be odd to characterize the outcome of the telling in terms of true belief satisfying further conditions. Indeed, as in the rain-check case, it is problematic, to say the least, whether those who are perfectly capable of understanding, and making, knowledgeable attributions of perceptual knowledge are bound to have the conceptual and theoretical resources for thinking about the knowledge in terms of beliefs justified by visual experience, if the experience is conceived as an 'internal' component of seeing.

It might be suggested that concepts that have to do with belief-formation and that figure in mainstream epistemology *are* implicit in our ordinary thinking. After all, being in command of a way of telling surely has something to do with reliability. Indeed it does. Being in command of any way of telling whether something of a certain sort is so entails being good at telling whether or not things of this sort are so. Being good at this obviously entails reliably getting it right when one aspires to do so. However, this kind of reliability is not the reliability of a belief-forming disposition; it is the reliability of *a person* with respect to a task—the task of telling whether or not something of a certain sort is so. The idea of this sort of reliability is not hard to grasp and forms part of our understanding of what it is to have a way of telling something. The other brand of reliability marks a further level of theoretical commitment.

Suppose that we grant that when people are reliable in the sort of way described, and they tell that something is so, they form a belief and thereby manifest a reliable belief-forming disposition. What might the disposition be? To have something to work with, let us suppose that the way of telling is telling by looking whether or not it is raining, and that the following holds of those who have command of this way of telling: (i) they have a disposition such that provided that they have visual experiences in a certain range, and do not take there to be reasons not to take their experiences at face value, they will form the belief that it is raining; (ii) they have a disposition such that provided that they have visual experiences in a certain distinct range, and do not take there to be reasons not to take their experience at face value, they will form the belief that it is not raining. Even if something like this holds it is far from obvious that it plays any role in attributions of knowledge. I have already cast doubt (a couple of paragraphs back) on whether we routinely work with concepts of experiences that fit the bill. Even if we do there is the further question concerning our grasp of the relevant ranges of experience. Obviously, those that trigger the manifestation of disposition (i) must be those that facilitate the recognition that it is raining, and those that trigger the manifestation of disposition (ii) must be those that facilitate recognition that it is not raining. But if the enterprise is to delve below the notion of recognition to a conceptual level that would enable us to spell out what recognition and recognitional abilities amount to in terms that do not implicate the concept of knowledge, yet are implicit in our ordinary thinking

about knowledge, then we have a problem: it is far from clear we have to hand a way to characterize the range of experiences independently of their being such as to facilitate recognition. Insofar as we have a grasp of what it is to have an experience such that it looks as if it is raining it appears to be derivative from our grasp of what it is like to see that it is raining. Insofar as we have a grasp of the range of experiences that might trigger visual recognition of rain it appears, again, to be derivative from our grasp of what it is like to see that it is raining. Perhaps it is possible that there should be some independent characterization of the relevant experiences from within the framework of a mature cognitive science. I have no wish to prejudge that matter. But even if there is a way yet to be developed it has no bearing on our present attributions of knowledge. That raises a general question concerning the relation between accounts of the nature of knowledge and accounts of the epistemology of knowledge-attribution. Let us not rule out that the former can draw on conceptual resources that are not part of our pre-theoretical understanding of knowledge. Still, if I am right in thinking that we are good at attributing knowledge to ourselves and to others, and that this reflects an understanding of what knowledge is, then we should attend more closely than is customary to whether accounts of knowledge fit with the understanding of knowledge that we deploy in practice.¹⁴

References

- Austin, J. L. (1946). Other minds. *Proceedings of the Aristotelian Society, supplementary volume 20*, 148–187 and reprinted in J. L. Austin's *Philosophical Papers*. Oxford: Clarendon Press, 1961, 44–84.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *The Journal of Philosophy*, 73(20), 771–791.
- Greco, J. (2007). The nature of ability and the purpose of knowledge. In E. Sosa & E. Villanueva (Eds.), *Philosophical issues*, 17, 2007: *The Metaphysics of Epistemology*, 57–69.
- Greco, J. (2009). Knowledge and success from ability. *Philosophical Studies*, 142(1), 17–26.
- Millar, A. (1991). *Reasons and experience*. Oxford: Clarendon Press.
- Millar, A. (2007a). What the disjunctivist is right about. *Philosophy and Phenomenological Research*, 74(1), 176–198.
- Millar, A. (2007b). The state of knowing. In E. Sosa & E. Villanueva (Eds.), *Philosophical issues*, 17, 2007: *The Metaphysics of Epistemology*, 179–96.
- Millar, A. (2008a). Perceptual-recognitional abilities and perceptual knowledge. In A. Haddock & F. E. Macpherson (Eds.), *Disjunctivism: Perception, action, knowledge*. Oxford: Clarendon Press.
- Millar, A. (2008b). Disjunctivism and skepticism. In J. Greco (Ed.), *The Oxford handbook of skepticism*. Oxford: Oxford University Press.
- Pritchard, D., Millar, A., & Haddock, A. (Forthcoming). *The nature and value of knowledge*. Oxford: Clarendon Press.
- Pryor, J. (2002). The skeptic and the dogmatist. *Nous*, 34(4), 517–549.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge*. Oxford: Oxford University Press.

¹⁴ I am grateful to the audience at the Bled Epistemology Conference in 2009 for discussion of some of the matters explored here and to helpful conversations with my colleagues Adrian Haddock and Peter Sullivan.