

Two-dimensionalism and the epistemology of recognition

Markos Valaris

© Springer Science+Business Media B.V. 2008

Abstract There is reason to expect a reasonable account of a priori knowledge to be linked with an account of the nature of conceptual thought. Recent “two-dimensionalist” accounts of conceptual thought propose an extremely direct connection between the two: on such views, being in a position to know a priori a large number of non-trivial propositions is a necessary condition of concept-possession. In this paper I criticize this view, by arguing that it requires an implausibly internalist and intellectualist conception of capacities we bring to bear in applying concepts in experience. Empirical concept-application depends on the exercise of a variety of capacities, many of which can be grouped together under the general label “recognitional”. As I argue, two-dimensionalism cannot accommodate a plausible account of such capacities. This suggests that the link between a priori knowledge and the nature of conceptual thought is not as direct as twodimensionalists take it to be. I close by briefly sketching a different way to think of that link.

Keywords Two-dimensionalism · Recognition · Epistemic intension · A priori · Externalism · Internalism · David Chalmers · Stephen Yablo · Frank Jackson

1

One prominent strand in the epistemology of a priori knowledge is based on the idea that our capacity to know a priori is to be explained in terms of the nature of conceptual thought. I am inclined to think that something along these lines must, in

M. Valaris (✉)
Philosophy, University of Pittsburgh, 1001 Cathedral of Learning, Pittsburgh, PA 15260, USA
e-mail: mcv8@pitt.edu

fact, be right, as the alternative seems quite unattractive. On the alternative picture, our capacity for a priori knowledge would have to be an optional extra, not consequent upon anything already required by our capacity for conceptual thought. It is easy to see why such a view is generally considered unappealing: for what light could be shed by postulating such a special-purpose capacity? In this paper I will not attempt to explore the relation between a priori knowledge and the nature of conceptual thought with any kind of generality. I will rather focus on one, recent account of that relation, which is to be found in the writings of David Chalmers and Frank Jackson on two-dimensionalism.

These authors propose an account of conceptual thought according to which possessing a concept is a matter of a subject's knowing, or being in a position to know through reflection alone, a large number of conditional propositions concerning what the concept applies to, depending on what the world turns out to be like.¹ Such knowledge, it is claimed, is a priori precisely because it depends only on the subject's possession of the relevant concepts and her capacity for reflection. In particular, it does not depend on experience.²

In this paper I will criticize this view by criticizing the account of concept-possession, and especially concept-application, entailed by it. As I already said, I believe the epistemology of a priori knowledge is indeed to be grounded on an account of conceptual thought. Nevertheless, the two-dimensionalist account of this connection cannot be correct, because it entails an implausibly internalist and intellectualist picture of capacities we bring to bear in empirically applying concepts. Concept-application in experience depends on the exercise of a variety of capacities, many of which can be grouped together under the general label "recognition": we recognize particulars as instantiating kinds, for instance, or as being identical to things we remember from past experience. But, as I will argue, two-dimensionalism is committed to an implausible account of recognition.³

I begin by giving a brief overview of two-dimensionalism, in Sect. 2. In Sect. 3, I introduce in a general way the sort of worry I will be pressing later on in the paper, and which I develop in detail in the sections that follow.

2

The two-dimensionalist account is supposed to apply, with appropriate modifications, both at the level of thought and at the level of language. I will only focus on

¹ The formulation used here is not accidental, but rather defended as a useful heuristic in Chalmers (1996, p. 57).

² Of course, our possession of some – and perhaps even all—of our concepts depends on experience. But this experience is not part of what *justifies* our knowledge of the corresponding conditionals, and justificatory dependence is just what is at issue here.

³ Much work on two-dimensionalism has been published recently. What follows has points of contact with some of it, particularly with Yablo (2002). I will elaborate on this later on in the paper.

the case of thought here.⁴ According to two-dimensionalism, to possess a concept C is to know, or to be in a position to know by reflection alone, what I will call its set of application-conditionals—that is, a set of material conditionals of the following form:

$$D \rightarrow E(C)$$

Here D is supposed to be a variable ranging over exhaustive descriptions of what—given all we know independently of experience—the actual world may turn out to be like, and $E(C)$ a specification of C 's extension, on the hypothesis that D is true of the actual world. Given the set of world-descriptions W , to every concept corresponds a set of application-conditionals, each of which has as its antecedent a member of W and as its consequent a specification of what, if anything, the concept applies to under the supposition that the actual world turns out to conform to the description in the antecedent.

Continuing in the same spirit, we can define for every concept its “epistemic” (or “primary”) intension.⁵ The epistemic intension of a concept C is a partially defined function from W , such that for any D —a member of W —on which it is defined, its value is the extension determined by C 's application-conditional with D as its antecedent. The epistemic intension of a concept should be distinguished from its “counterfactual” or “secondary” intension. The counterfactual intension is defined in the way familiar from ordinary (one-dimensional) modal semantics, and determines what the concept applies to in counterfactual situations. Thus, to take a standard example, the counterfactual or secondary intension of *water* is a constant function, whose value is the same stuff— H_2O —in every world where it is defined (throughout, I use italics to refer to concepts). By contrast, since (for example) we cannot a priori exclude the possibility that it may turn out that there is actually no such stuff as H_2O , the value of the epistemic or primary intension of *water* is not always H_2O . There are ways the world (the actual world) may turn out to be on which the stuff the concept *water* applies to is not H_2O .

Grasping a concept, knowing or being in a position to know its application-conditionals, and grasping its epistemic intension are supposed to be the same state, according to the two-dimensionalist picture.

Of course no actual subject is capable of entertaining anything like the exhaustive description of the actual world that D purports to be, let alone evaluating any conditional with D as its antecedent. The two-dimensionalist's claim, therefore, must be not that a subject possessing C actually *knows* any such conditionals (since

⁴ A note about terminology is in order here. According to Chalmers's usage, the term ‘concept’ applies to mental representations, that is, mental particulars with semantic properties analogous to those of words (see Chalmers 2002c, p. 3). This is somewhat surprising, since on Chalmers's Fregean view it would be more natural to take concepts to be non-mental, abstract entities. What could be considered as mental is the state of grasping a concept; but then, that is most naturally construed as general rather than particular, since it is a state many different subjects can be in at many different times. In what follows I will use the term ‘concept’ to refer not to mental particulars, but rather to components of thoughts that can be shared by many subjects.

⁵ My identification of primary and epistemic intensions here follows Chalmers (2006). In earlier work Chalmers took the identity to be something of an open question, to be established by argument (see Chalmers 2002a, pp. 166–167).

knowledge intuitively entails understanding), but rather something like the following counterfactual: just in virtue of possessing a concept *C*, a subject possesses all that would be required in order for her to know the corresponding application-conditionals, if it were not for limitations of memory, computational speed, life-span and the like.

Now, if *D* is allowed to involve any concepts we like, then the two-dimensionalist's claim will be rather uncontroversial. This is because in the absence of any restrictions on the concepts we are allowed to draw upon, we can always employ *C* itself; and in this case, the application-conditionals will be no more informative than a homophonic specification of *C*'s satisfaction-conditions. But it is relatively uncontroversial that a subject who is minimally reflective, and possesses the intentional concepts required to formulate homophonic satisfaction-conditions, will normally be in a position to know such conditions for any concept she possesses. If the two-dimensionalist claim is to be interesting, therefore, restrictions must be imposed on the concepts *D* is allowed to involve. At a minimum, we must not be allowed to use *C* itself, or concepts too closely related to it, in the formulation of *D*. Chalmers imposes much tighter restrictions, claiming that the vocabulary of fundamental physics (or some future refinement thereof), together with purely phenomenal concepts, suffice for the formulation of application-conditionals for all of our concepts.⁶ This is obviously an extremely controversial thesis, and I will not be concerned with it in this paper. I will remain agnostic as to what exactly is allowed in *D*, requiring only that application-conditionals must be more informative than the corresponding homophonic satisfaction-conditions.⁷

For present purposes, at least, what is most interesting about this account of concept-possession are its implications for the epistemology of a priori knowledge. As we already saw, according to this picture possessing a concept *C* entails being in possession of everything required to know its application-conditionals by reflection alone. Of course *we* cannot know these conditionals, since we are not in a position to entertain the exhaustive world-descriptions in their antecedents. Nevertheless, it seems very plausible that this capacity can put *us* in a position to know a priori a lot of more restricted, but still non-trivial propositions: we can, for example, presumably know a priori—just on the basis of our possession of the concept *water*—that the clear, potable stuff that fills lakes and rivers is water.

3

It is clear that two-dimensionalism is committed to a kind of individualism (or internalism) about concept-possession. Possessing a concept, according to

⁶ See, for example, Chalmers and Jackson (2001). Of course two-dimensionalism must offer a different account of what it is to possess the concepts included in the basic vocabulary—perhaps by allowing homophonic application-conditionals for these concepts.

⁷ Another requirement is that application-conditionals must use only “semantically neutral” vocabulary—they must not, in other words, employ any terms whose content depends on which world is treated as actual (see, for example, Chalmers 2002a, p. 166, 2002b, p. 17, fn 5)).

two-dimensionalism, is a matter of capacities which are constitutively independent of one's relations to a physical or social environment. This, as I will suggest, is what is wrong with the view.

Two-dimensionalism of course concedes something to externalist orthodoxy by granting that the counterfactual intension of many of our concepts is constant, with its value fixed by the subject's relations to her environment (since it is, as we saw, fixed by the concept's extension in the actual world). This entails that conditions specifiable at the level of an individual subject will not, in general, be sufficient for concept-possession. A subject who has lived all her life on Putnam's Twin Earth (thus lacking all direct or indirect causal contact with H₂O), for example, cannot express the same concept by her word 'water' as we do, even though all her individualistically specifiable dispositions may match those of a subject on Earth.⁸

This concession, however, does not substantially alter the individualistic character of the account. We should remember that possessing a concept is not a matter that directly involves its counterfactual intension—it is a matter of knowing, or being in a position to know, its application conditionals. Of course, to the extent that a concept is *used*—as opposed to merely mentioned—in its application-conditionals, a subject's being in a position to know a specific set of conditionals implicitly settles which world-description in *W* is true of the world that is actual for her—it settles, in other words, which world is the *actual* actual world for that subject. A subject living on Twin Earth, for example, is not in a position to know any propositions involving the concept *water*; a fortiori, she is not in a position to know its application-conditionals. The crucial point, however, is that the capacity a subject on Earth would manifest in endorsing—and thereby coming to know—the application-conditionals of *water*, is just the same capacity that a subject on Twin Earth would manifest in endorsing the application-conditionals of *her* concept. That capacity is the capacity to pick out, for each possible world-description, the value of the epistemic intension of one's concept on the hypothesis that it is true of the actual world. Since the epistemic intension of *water* is just the same as the epistemic intension of the concept they express on Twin Earth by their word 'water', what it takes for a subject to possess the concept *water* is just the same as what it takes for a subject to possess the Twin Earth concept.

The concession, therefore, does not change the fact that possessing a concept is, according to two-dimensionalism, a matter of dispositions or capacities constitutively independent of one's environment. This is the aspect of the view I will be criticizing here.

Now, one may wonder whether *any* non-trivially specifiable cognitive capacities—over and above, that is, than the capacity to think thoughts involving the concept in question, and perhaps to know its homophonic satisfaction-conditions—are entailed by possession of any particular concept. Putnam, for example, famously argues that being able to pick out the extension of 'beech' and 'elm' is not required for understanding those words. More recently, Timothy Williamson argues that having dispositions to infer in accordance with their introduction and elimination rules is not required for possessing the concepts of the logical constants. Possessing

⁸ See Putnam (1975).

these concepts, according to these authors, is rather a matter of participating in the corresponding intellectual practices.⁹

Such arguments certainly appear damaging to the two-dimensionalist view, since on that view possessing a concept requires being in a position to pick out its extension for every way the world may turn out to be. It is not entirely clear how two-dimensionalists should respond to such an objection. I will not here try to sift through the different possibilities, since I do not intend to push this particular objection. On the contrary, I am willing to concede—at least for the sake of the argument—that there is a significant class of concepts possession of which is associated with some distinctive, non-trivially specifiable cognitive capacities.

It is surely the point of many (though by no means of all) of our concepts that they pick out, and thus enable us to classify, things we encounter in our everyday experience. Concepts otherwise very different from each other—such as *yellow*, *round*, *mountain*, *cat* or *chair*—all seem to serve such a function.¹⁰ It is natural to think that, for a certain notion of normality, possessing these concepts is normally—even if not always or necessarily—associated with a capacity to correctly apply them in the course of one's experience. We should expect—and, intuitively, we *do* expect—that a subject who possesses the concept *yellow* will, under favorable conditions, be able to classify yellow things as yellow by sight, and that a subject who possesses the concept *mountain* will be able to classify mountains as mountains, given appropriate information. Developing this natural thought into a plausible account of concept-possession may not be a straightforward matter at all. Nevertheless, for the purposes of this paper, I am going to make the assumption that there are some concepts possession of which normally entails a capacity to pick out their extension. My argument will be that two-dimensionalism needs a very implausible conception of such capacities.

4

Consider a concept of this sort—for example, the concept *water*. Combining the points of the last section with the two-dimensionalist picture, we get the claim that possession of the concept *water* normally entails having all that one would need to be in a position to know its application-conditionals through reflection alone (contingent limitations aside—from now on I will avoid explicitly mentioning this qualification).

Now, in order for this claim to be even remotely plausible, we would need first to exclude cases where possession of the concept intuitively does not involve a reliable disposition to correctly apply the concept. We must, therefore, exclude subjects whose grasp of the concept is deferential, or who adhere to confused but consistent and not easily refutable theories about water, and the like. Assuming this to be

⁹ Putnam (1975) and Williamson (2003, 2006). Similar arguments have been proposed by various other authors as well.

¹⁰ The class of concepts that do not—or do not primarily—serve this function is very diverse, of course. It includes, for example, normative concepts and logical concepts.

possible, the most plausible candidate for a capacity normally entailed by possession of the concept is the capacity to recognize water by its manifest qualities in conditions ordinary for us—that is, the ability to recognize as water liquid water, snow, ice, steam and the like, at least in ordinary environments. Thus the two-dimensionalist claim must be that possessing this recognitional capacity entails being in a position to know the relevant application-conditionals through reflection alone.

How plausible is this claim? Clearly, the assumption that a subject possessing the concept *water* normally has the relevant recognitional capacity entails that, given a suitable form of access to a world or world-description, she will be able to pick out the water in it. This, however, is a conditional claim about what the subject will be in a position to do or know under certain specified conditions. By itself, it as yet tells us nothing about whether the subject is in a position to know the corresponding application-conditional at all, let alone know it by reflection alone. In order to assess this latter claim we need to be more precise on what we take the relevant recognitional capacity to consist in.

In the following couple of sections I will discuss an account of recognition which, if true, would support the two-dimensionalist claim (or at least a reasonable reformulation thereof). I will begin with a rather crude version of the account, which I will subsequently refine. Still, I believe that this whole approach to recognition is very implausible, and in fact evinces a commitment to a very unattractive picture of concept-application. After arguing against it in Sect. 6, I will in Sect. 7 sketch an alternative account of recognition. That account has no tendency at all to support the two-dimensionalist claim.

5

Here is a crude picture of recognition, which nevertheless straightforwardly supports the two-dimensionalist claim. To recognize water in a given world-description is, on this picture, a matter of discursive reasoning.¹¹ Given the relevant world-description *D*, our subject brings to mind a specification stored in memory of what water is like, couched in the same vocabulary as *D*. She then goes through *D* with this specification in mind, figuring out what, if anything, it applies to. Accordingly, to have a capacity to recognize water would consist precisely in having stored in memory a specification of what water is like, which can be recalled

¹¹ I use the term “discursive reasoning” for reasoning that can be represented in words, regardless of whether the subject who engages in it consciously goes through some kind of internal monologue or not. I take it as obvious that not all reasoning is discursive in this sense. When considering how to get a large sofa down a narrow stairwell, for example, I am definitely engaged in a process of reasoning. But it is clear that I do not need to have any knowledge of the geometrical theory that would be required to formulate and solve this problem in a language. My reasoning, therefore, is not discursive in the present sense.

and compared with an appropriately described situation (in the case where the situation is directly experienced we can assume that recognition will work on an appropriately formulated description derived from the experience).¹²

It is clear that on this inferentialist picture the recognitional capacity is also a capacity to know the relevant application-conditionals. In fact, on this picture there is no real difference between recognizing water and reasoning by *modus ponens* from the appropriate application-conditional and world-description. This crude inferentialism about recognition, of course, is not often explicitly defended: it is extremely implausible—just on plain phenomenological grounds—that recognition must be a matter of discursive reasoning. We usually don't know accurate enough discursive specifications of the things we are able to recognize, and, even in the exceptional cases where we do, we certainly do not seem to use them in the way this account supposes. But even if the crude account fails, a more refined version of the same idea might still seem tempting.

Recognition is, obviously, driven by the way something looks (or, more generally, the way it appears to our senses). This might make it seem that recognition *must* be a matter of comparing a specification stored in memory with what is currently before one's senses. If the crude picture is misleading, therefore, one might suggest that the reason is that it assumes that the stored specification must be discursive. But this assumption is not compulsory: the stored specification might be something more like a mental image which the subject cannot fully put into words, for example.

How would this admission affect the two-dimensionalist claim that possessing the recognitional capacity puts a subject in a position to know the relevant application-conditionals through reflection alone? Unless the claim is reformulated in some way, it would seem that such an admission would undermine it. As Steven Yablo points out, a lot of our concepts—our “observational concepts”, as he collectively calls them—require for their application an exercise of our sensible capacities. Even in cases where the application is only hypothetical—cases in which we are trying to work out what the concept would apply to, on the hypothesis that such-and-such turns out to be the case—we need an “offline”, imaginative exercise of our sensibility.¹³

¹² This account (as well as the refinement of it I will discuss shortly) suffers from an obvious difficulty, which for reasons of simplicity I will ignore. No specification involving only the qualitative properties of water can capture the required recognitional capacity, because no such specification discriminates between water and twin-water (a substance found only on Twin Earth, a distant but actual planet). But, intuitively, if a person from Earth calls twin-water ‘water’, she has *not* recognized twin-water as water: she has, rather, made a mistake. Her recognitional capacity has misfired in this case. The problem is that on the present account no mistake can be found in her performance, since she has correctly judged that her specification applies to twin-water.

What we need seems to be a requirement that the object picked out be the one the subject has actually interacted with in the past, and from which the stored specification derives. So the specification would end up looking like this (where *S* is supposed to be a specification involving only qualitative properties): “the *x* such that: (*x* is *S* & my past interactions with *x* are the source of *S*)”. These complications will be ignored in the text, since there are more fundamental problems facing the proposed account.

¹³ 2002, pp. 441–492.

One of the cases Yablo discusses in detail is that of a certain family of oval figures, called “Cassinis”.¹⁴ The shapes can be defined algebraically as the sets of points on the Cartesian plane satisfying equations of a certain form. But, as Yablo points out, knowing its algebraic equation is not enough to know that a Cassini is oval: the only way to know *that* is to actually plot the figure, either in one’s imagination or on a piece of paper. But neither of these methods, according to Yablo, qualifies as knowledge through reflection alone. Knowing that Cassinis are oval, as he puts it, requires “peeking”. Thus, even if it is true that a subject who possesses the concept *oval* (and the relevant mathematical sophistication) is in a position to know the conditional “if something is Cassini-shaped, then it is oval”, she is not in a position to know it *through reflection alone*.

This argument is convincing as far as it goes, but it doesn’t go far enough. There is no deep reason, I think, why an inferentialist about recognition cannot build “peeking” into her account.¹⁵ More generally, one could state the inferentialist view in a way that avoids commitment to any particular account of *how* the relevant specification is stored and compared to what is currently before one’s senses.

What the refined inferentialist account should require is just the following. Suppose that, in a specific experiential context, one knows a proposition by recognition (this proposition might be an identity proposition, such as “this is Fred”, or a proposition subsuming something under a kind-concept, such as “this is oval”—where “this” is a perceptual demonstrative). Then, one’s grounds for that knowledge must consist in the cogency of an argument, whose premises include: (i) a formulation in the appropriate phenomenal vocabulary of a specification stored in the subject’s memory, and (ii) a description of her current experience, formulated in the same phenomenal vocabulary.

The crucial point here is that, in order for the grounds of one’s recognitional knowledge to consist in the cogency of this argument, one need not actually arrive at the knowledge by thinking through the argument itself. Of course, in order for the inferentialist account of be true, recognition should, normally at least, involve on the part of the subject *some* sort of reasoning from the stored specification and her current experience to the recognitional proposition. However, the details of any particular process of recognition need not be relevant to the truth of the inferentialist account, as long as an argument of the above form captures its epistemic thrust. If this is correct, then two-dimensionalism should also be allowed to help itself to the same unspecified sort of process as the relevant way of coming to know the application-conditionals of one’s concepts. As Chalmers puts it, peeking may come for free.

¹⁴ 2002, pp. 467–469. As Yablo notes, not all Cassinis turn out to be oval; but, like Yablo, I ignore the point here.

¹⁵ This would seem to be Chalmers’s own response: “given the phenomenal information [...] about appearances (in a pure phenomenal vocabulary) ‘peeking’ comes for free” (2002a, p. 185).

Nevertheless, as I will argue in the next section, we have reason to reject even this refined account: recognition is not to be construed as a matter of inference at all. To recognize an object is to non-inferentially identify an object presented to one with an object one has encountered in the past. To recognize an instance of a kind—such as oval or water—is, analogously, to non-inferentially identify something in one’s environment as an instance of that kind. If this is right, then recognitional capacities offer no support to the two-dimensionalist claim, even in its more relaxed version.¹⁶

6

Consider again the case of the concept *oval*. I take it that *oval* is what we may call a “manifest shape” concept. It is a concept that picks out certain geometrical shapes, but it does so in virtue of a correlation with the way they appear to us. Thus *oval* is not to be identified with any concept formulated in exclusively mathematical terms (and not just because of vagueness). Suppose, now, that possessing the concept *oval* consists in having the capacity to recognize oval-shaped things as such, and that recognition is conceived inferentially—recognizing something, that is, is a matter of matching its current appearance with a stored specification. Now, we can certainly imagine world-descriptions in which, for whatever reason, some things whose contours approximately satisfy circle equations systematically appear oval to us. We may suppose, along with Yablo, that a world-description in which dimes appear oval to us is true of the actual world. Would that be a case in which dimes turn out to be oval, instead of round?

It seems clear that it would not: the existence of such a glitch in our visual system should, intuitively, have no effect on the manifest shape of dimes. Whether dimes are oval or round depends only on their geometrical properties. If we accept the inferentialist account of recognition, however, it seems say that we have to deny that. Since recognition is a matter of the fit between a stored specification and an appearance, on the hypothesis that the way dimes actually appear matches the specification, they are *genuinely recognized* as oval. Given the factiveness of recognition, dimes turn out to be oval.

Conversely, if we make the supposition that some oval-shaped things—for example, the dining-room table—do not look oval in the actual world, should we say that on that supposition they turn out not to be oval? Again, that seems false: in order for the table not to be oval it has to have different geometrical properties; it’s not enough that the way it looks to us be different. Just as before, however, if

¹⁶ Even though Yablo does not distinguish between inferentialist and non-inferentialist accounts of recognition, some of the things he says suggest he has a non-inferentialist account of recognition in mind. See, primarily, the argument in Yablo (2002, pp. 473–474; also 464–465). If his intention is indeed to argue against the inferentialist conception of recognition or concept-application in general, then his argument about peeking is not sufficient.

recognizing something as *oval* consists in matching a specification with an appearance, we have to wonder whether we can avoid the conclusion that the table turns out not to be oval.¹⁷

This might not be a problem if the subject had other resources, additional to her perceptual appearances and her stored specification, sufficient to rule out the false positive or negatives on the relevant hypotheses. I do not think the resources two-dimensionalism allows are in general sufficient for that, however. Consider how this would be supposed to work in the first of the above examples. In this case, the issue is to rule out the result that our subject recognizes something as oval when it is not, in fact, oval. Our subject has in her possession, we are supposing, a complete—matters of manifest shape aside—description of the world, including both the geometrical properties of things, and how shapes appear to her and to others. Could some of this collateral information tip her off that, despite the match between her appearance and the stored specification, she should not apply the concept *oval* to dimes?

If the distortion affects only a small subset of our subject's appearances (perhaps only approximately dime-shaped round things tend to appear as oval to her) then that very fact might alert her that something is going wrong. The subject, we may assume, knows that whether *oval* applies to an object or not depends only on its geometrical properties; since she knows that the geometry of dimes is similar to that of other things that appear to her round rather than oval, she might reason that her dime-appearances are distorted. Thus she would withhold her application of *oval* to dimes. Even if *all* of our subject's shape-appearances are systematically distorted, she might still be able to reason her way out of the problem as long as we assume that the distortion is statistically abnormal among her peers.

Consider, however, how things would be if we assumed instead that the glitch has a systematic effect on all of a statistically normal subject's visual perceptions. The way dimes appear to our subject in this case is not statistically abnormal at all: this

¹⁷ On this point, I have encountered the objection that, whatever capacity a subject has to recognize manifest shapes (and oval among them), it is retained when she considers other hypothetical scenarios as actual. Thus, it is objected, when the subject considers the specified hypothesis as actual, she can tell that dimes are round and not oval, and accordingly she can discount her appearances on that supposition as misleading. As a point about the evaluation of conditional claims of the relevant form this appears to be all right: after all, we do intuitively judge that dimes are not oval on the relevant hypothesis. Nevertheless, the question here is to *account* for our ability to apply the concept oval, whether hypothetically or not (and thus also for our ability to evaluate the relevant conditionals). This is precisely the feature of two-dimensionalism we are in the process of evaluating. We cannot, therefore, simply *assume* that the ability to correctly apply the concept oval carries over to the hypothetical scenario. Instead, we must examine whether the resources available to two-dimensionalism really do amount to a capacity to correctly apply the concept hypothetically.

The point may be illustrated by considering a simple attempt to capture the epistemic intension of *oval* by a description. Consider the following biconditional: 'x is oval if and only if x looks oval in the actual world (to normal subjects and in normal conditions)'. Such a biconditional seems to be in the spirit of two-dimensionalism (general doubts about the coherence of introducing concepts by means of such biconditionals are not relevant here). This biconditional makes it clear that, on the hypothesis that dimes normally look oval in the actual world, dimes *are* oval. The analogous result might be right for so-called "response dependent" concepts, such as the concept *cool*. But, intuitively, such concepts lack a kind of objectivity that concepts of manifest shape possess.

is just how one would expect dimes to look to our subject, given how everything else looks both to her and to her peers. Now, what could tip our subject off that her dime-experience is not veridical? In order for our subject to figure this out, she must be capable of figuring out how dimes *ought* to appear. But how could she do such a thing? By hypothesis, actual appearances are not going to be of any help: as we have assumed, the way things appear to her is systematically distorted, so that round things appear to her oval instead of round. But neither would her otherwise complete description of the world suffice. For one thing, the description of the world available to her must be free of manifest shape concepts (otherwise the two-dimensionalist claim we are evaluating becomes trivial). But neither can she rely on the appearances of her peers, since they are—as we are assuming—systematically distorted in just the same way as her own. Since neither her appearances nor the description of the world she is considering can be of any help, our subject has no resources that would suffice for her to figure out that her dime-experiences are not veridical.

7

The kind of difficulty discussed in the last section does not depend on the specific features of the examples employed. It rather results from some central features of the two-dimensionalist position. In the next paragraph I will describe where I think the problems stem from. The rest of the section is devoted to presenting a different view, which avoids such difficulties.

The concepts I have been concerned with are associated with recognitional capacities. Given such a concept, it is always an a posteriori matter what (if anything) is in its extension—what, if anything, it applies to in the *actual* actual world. It is not something that can be settled by reflection alone. This, I take it, is uncontroversial. According to two-dimensionalism, however, this further implies that the connection between the recognitional capacity and the concept's extension is not constitutive, but only contingent. This follows from the fact that, when determining the epistemic intension of a concept, the only things we are supposed to hold fixed across every possible world-description are those that can be settled by reflection alone. If, however, the recognitional capacity were constitutively linked to an extension, then that link, even though not knowable by reflection alone, would settle the value of the concept's epistemic intension for every world-description. This, I suggest, is why two-dimensionalism requires an inferentialist account of recognition. (In Sect. 8 I will provide an additional argument, to the effect that two-dimensionalism cannot shed this commitment). On this account, recognition consists in matching a specification stored in one's memory with one's current appearances. Since it is surely contingent which things cause appearances that match the specification, the contingency two-dimensionalism requires is thereby secured. But this is also just what leads to the unacceptable results of the last section.

Accordingly, the problems are avoided if recognition is not a matter of matching a stored specification with an appearance.¹⁸ Denying the inferentialist model of recognition does not amount to denying the obvious fact that recognition is driven by the way an object appears to us, and therefore that it requires information about the object's appearance to be stored in us in some way. The point is simply that the process by means of which the stored specification is compared to the present experience is not to be construed as an inference. This process may be, for all that the phenomenology of recognition requires, just a sub-personal, neural process, not to be construed as a cognitive act of the subject herself.

So what would the non-inferentialist account of recognition look like? Consider first the case of recognition of a particular object *o*. According to the present account, recognizing *o* in one's perceptual environment is to non-inferentially know a proposition of the form "this is *o*" (where 'this' is an appropriate perceptual demonstrative mode of presentation of *o*).¹⁹ To have a capacity to recognize *o* will, accordingly, be to be able to non-inferentially know that *o* is identical to an object in one's perceptual environment, if it is presented to one under an appropriate appearance.

One thing to immediately note is that, on this view, recognitional knowledge does not depend on attributing any particular appearance to *o*. For all that the present account requires, the subject need not be able to think of the way *o* appears at all, except perhaps by using the very same perceptual demonstrative—as in "*o* looks like *this*". One derives propositions about how *o* appears from one's recognitional knowledge, rather than the other way round. (On the other hand, the specification of a recognitional capacity must make reference to a specific appearance. This reflects the intuition that, for example, two people who have known Fred in two completely different periods of his life—perhaps one knew him as a young child and the other as an old man—will have distinct recognitional capacities with respect to the same person.)

The case of recognition of kinds can be treated in an exactly parallel fashion. Suppose that 'this' expresses a demonstrative mode of presentation of an item in one's present perceptual context. To recognize that item as belonging to a kind *k* will be to non-inferentially know the proposition "this is a *k*". To have a capacity to recognize *k* will be to have the capacity to non-inferentially know of instances of *k*, presented to the subject under an appropriate appearance, that they are instances of *k*. Again, the way the kind looks to the subject is relevant to the individuation of the recognitional capacity, but the knowledge the subject acquires by recognizing

¹⁸ Strictly speaking, denying inferentialism is not the only way to avoid the problems of Sect. 6. There seems to be room for an externally constrained inferentialism, according to which what *oval* applies to is held constant, even though recognition still operates inferentially (the possibility of this hybrid view was suggested to me by John McDowell). This view is still not acceptable to two-dimensionalists for just the reasons explained in the text: it requires us to hold fixed something that is not knowable by reflection alone. I do not consider this view (except briefly in Sect. 9), because I find the combination of constant intension and inferentialism unmotivated.

¹⁹ I am assuming that states of knowing are individuated more finely than reference is, so that the knowledge described in the text does not collapse to knowledge of the proposition that *o* is *o*.

instances of the kind is not derived from prior knowledge that something with that appearance is present in her environment.

If this non-inferentialist account of recognition is correct, then the two-dimensionalist claim must be false. I will argue for this explicitly in the next section, but the reason why should be intuitively clear: on the present account, recognitional capacities are individuated in part in terms of what they pick out in the subject's environment—in the *actual* actual world. If, on the hypothesis that the world turns out to be otherwise, the subject happens to pick out something different, that is just a mistake: on that hypothesis, the subject's recognitional capacity simply misfires. According to two-dimensionalism, on the other hand, the extension of a concept cannot play a constitutive role in individuating the capacity that possessing that concept consists in. This, as we saw, is why two-dimensionalism leads to the unacceptable results of Sect. 6.

8

We have seen an argument against the inferentialist conception of recognition, which I suggested two-dimensionalism is committed to. In this section, I will attempt to argue in a more detailed way that this commitment is not optional. I will do this by arguing that a hybrid view combining the main two-dimensionalist thesis with a non-inferentialist account of recognition does not result in a coherent account of concept-possession.

It might be suggested (although I do not think it actually *has* been suggested) that, even though there is no specification the subject employs in order to pick out what her concept applies to, there still is a specification which describes her non-inferential recognitional capacity. After all, assuming that the subject has such a capacity, one could define a function from ways the world may turn out to be to items of the appropriate category for the concept to apply to, such that for each way the world may turn out to be such an item is the value of the function if and only if it is what the subject would pick out if the world turned out to be that way. (I will call this function the 'recognition-function'.)

It is very questionable whether the disposition captured by the recognition-function can correctly be described as a disposition to *know* the application-conditionals, even assuming that it may capture a disposition to merely *conform* to them. The two-dimensionalist claim is that possessing a concept puts one in a position to know its application-conditionals through reflection alone. But if the capacity to apply the concept through recognition of its instances is non-inferential, then that capacity does not entail that the subject herself has any direct cognitive access to the conditions under which the concept applies. On the non-inferentialist picture, therefore, it is more plausible that a subject is only able to know the application-conditionals of her concepts a posteriori, through a course of experience in which those concepts are employed.

But this is not the objection I want to push here. An even more fundamental problem with the proposal to combine two-dimensionalism with a non-inferentialist account of recognition is that the proposed recognition-function cannot even capture

a disposition to *conform* to the application-conditionals, on pain of rendering the two-dimensionalist account of concept possession incoherent. This argument gives us even more reason to think that two-dimensionalism must be committed to inferentialism about recognition in concept-application.

Now, if there is any plausibility in the idea that the recognition-function captures a disposition to conform to the application-conditionals for a given concept *C*, it must be that for every world-description, the value of the recognition-function for a subject who possesses *C* is the same as the value of *C*'s epistemic or primary intension for that world-description. But the epistemic intension of a concept fixes what it applies to for every world-description, on the supposition that the description is true of the actual world. Thus if the subject's recognition-function for a concept *C* is guaranteed to match its epistemic intension, then the subject is guaranteed to never be mistaken in her application of the concept.

This, however, is not a coherent account of concept possession. To possess a concept is to be subject to a normative standard for one's performances—this seems to be one of the few uncontroversial lessons of Wittgenstein's *Philosophical Investigations*.²⁰ But we cannot make sense of a state's placing its subject under a normative standard if the state is simply *defined* so as to make every performance of the subject count as correct. This, however, is precisely what we have done by stipulating that the epistemic intension must match the subject's recognition-function.

To return to Yablo's example, there are possible world-descriptions in which a visual glitch makes our subject see dimes as oval. Supposing one of these to be true of the actual world, she will pick dimes out as items to which the concept *oval* applies. The intuitive response is that this will be a mistake on her part: dimes do not stop being round simply because our subject takes them to be oval. But on the present account no mistake has been made: the concept our subject uses has an epistemic intension that is guaranteed to match her recognition-function; thus, if on the hypothesis that the above world-description is true of the actual world she is disposed to pick dimes as falling under her concept, then dimes *do* fall under her concept on that hypothesis.

Perhaps one might object here that, in the spirit of two-dimensionalism, the recognition-function should not have been defined in terms of the subject's mere disposition to pick out a certain item, but rather in terms of her disposition to pick out an item *only after ideal reflection*. If the definition of the recognition-function is thus qualified, the response would go, then the identification of the recognition-function with the epistemic intension would be rendered harmless: since no actual subject is capable of ideal reflection, mistakes will be made just to that extent.

This response, however, simply misses the point. The account we have been discussing is supposed to incorporate the non-inferential account of recognition. If this is so, then the distinction between ideally reflective and ordinary concept-application cannot do the work it is supposed to do. If the way a concept is applied in experience is non-inferential, then—intuitively—we must make room for mistakes that are not errors in reasoning of any sort. I take something to be oval,

²⁰ The idea is also at the heart of Kripke (1982).

and I turn out to be wrong. Since my application of the concept is not the product of a process of reasoning, my error is not an error in reasoning either. It may be simply a “brute” error (as opposed to a “procedural” one) in the sense that, no matter how finely we analyze the process leading up to it, the error may not be due to anything that *I* did wrong at any particular step of the way. Since brute error is not due to anything the subject does wrong in the first place, granting her ideal powers of reflection should do nothing to shield her from brute error. It follows that even ideally reflective concept-application cannot be infallible. But since, on the present proposal, two-dimensionalism identifies a concept’s epistemic intension with what the subject applies the concept to under ideal reflection, ideally reflective concept-application *has* to be infallible. The possibility of error, therefore, is still not properly accounted for.

As I already said, I do not think this type of account has been explicitly defended by any proponent of two-dimensionalism. Perhaps Jackson comes closest, in a discussion of Gareth Evans’s non-inferentialist account of recognition. Jackson’s claim in that context is that recognitional capacities can be assimilated to capacities for descriptive identification, even though recognition itself remains non-inferential. According to Jackson, this is because “being disposed to be recognized by a subject as a such-and-such is a descriptive property of something that can and does serve to pick it out from other things”.²¹ Of course this is a descriptive property that can serve to pick out certain things; but it cannot be the property, or the concept, of a “such-and-such”, for in that case the subject would be rendered infallible in her judgments about which things the concept applies to: if she feels disposed to call something a such-and-such, then that very fact *makes* the thing a such-and-such. An attempt to qualify this by appeal to ideal reflection will fail, in the way just discussed.

9

One might perhaps accept the argument of Sects. 6 and 7, and the non-inferentialist account of recognition that is its upshot, but still wonder whether the non-inferential character of recognition might be just an accident, due to contingent limitations of our cognitive apparatus. Perhaps we humans need non-inferential, hard-wired recognitional skills to overcome limitations of computational power: recognition in us works much faster and perhaps more reliably in the non-inferential fashion than it would if it had to rely on reasoning. But we can certainly imagine beings equipped with much more computational power. Perhaps such beings either have no need for non-inferential recognitional capacities, doing instead all the recognitional work through reasoning, or, while still relying in practice on non-inferential recognitional capacities, they are also capable of having personal-level attitudes to the

²¹ Jackson (1998, p. 66). See Evans (1982, pp. 287–289).

information employed by their sub-personal recognitional systems.²² It is arguable that such beings would be able to know the application-conditionals of their concepts through reflection alone.

Perhaps this may be thought not to matter, since it may be claimed that the concepts these beings use are different from the ones *we* use, due precisely to the different cognitive processes used in their application. Suppose, for example, that these beings have a concept—call it *schmoval*—that classifies shapes using the very same appearances our concept *oval* uses. If this concept is used inferentially, then it will differ from our concept *oval* for the reasons outlined in §6 above: on the hypothesis that our beings' visual system has a glitch that makes dimes appear oval to them, dimes will be *schmoval*, but they will still not be oval. Thus *schmoval* and *oval* are not the same concept, since their epistemic intensions differ.

It is not, however, essential to the spirit of the current proposal that the supposed beings apply their concepts in this simple inferential way. Perhaps some externally constrained form of inferentialism is true of these beings: perhaps *schmoval* always tracks the manifest shape *oval* by means of appearance, even in worlds where the associated recognitional capacity fails to pick that shape out. From this it follows that two-dimensionalism is still false of *schmoval*, since its epistemic intension is a constant, with its value fixed by the property it picks out in our beings' *actual* actual world (see also n. 18 above). Nevertheless, it may be worth considering the proposal for its own sake, even though it cannot help save the letter of two-dimensionalism.

Now, let us imagine beings of whom this externally constrained inferentialism is true, and who are moreover able to infer both ways between the algebraic equation a figure satisfies and the way it appears to them. They are able, just by looking at (or visually imagining) a plane figure, and with the help of background knowledge concerning their visual capacities, to infer to the algebraic equation it satisfies. Conversely, given an appropriate algebraic equation, they have no trouble inferring from it and the relevant knowledge about their visual capacities to infer what the corresponding figure would look like to them. Such creatures would not need to rely on non-inferential responses in order to group together the things we call *oval*: they would be able to achieve that classification inferentially, by means of explicit reasoning.²³ Still, it does not seem altogether out of place to say that their concept *schmoval* is just the same concept as the concept *oval* that we use. After all, their concepts group together the very same shapes, and they do so by means of the very same appearance. But now, the argument goes, if such beings are indeed possible, couldn't we idealize away from the non-inferential character of *our* recognitional

²² If concept-application is inferential all the way down, then, plausibly, there must be some basic concepts that are applied on the basis of only trivial inferences. Such inferences will consist in detaching from a homophonic statement of the satisfaction-conditions of a concept, on the basis of an experience described using the concept in question itself (See also footnote 6 above).

²³ As already noted, at the most basic level this reasoning will simply involve applying a concept on the basis of its homophonic satisfaction-conditions and an experience described using that very concept. I am assuming here (as I have been doing throughout) that manifest shape concepts are not basic in this sense—that is, I assume that on the two-dimensionalist view there *will* be non-trivial application-conditionals that one must be in a position to know in order to possess these concepts.

capacities, regarding it as just one more of the contingent limitations we in any case have to idealize away from?²⁴

I think this argument should be resisted. The kind of idealization proposed here is significantly different from the one that was built into the two-dimensionalist account from the start. That idealization involved abstracting from contingent limitations on information storage and processing capacities; crucially, it did not involve a fundamentally *new* source of information. But this is precisely what the present proposal involves. According to the non-inferentialist account of recognition (which is accepted as true of us by the present proposal), the information driving our recognitional capacities is not available to us through reflection; it is available to us only through ordinary empirical investigation. The proposed idealization invites us to lift that limitation. In effect, it invites us to consider what we *would* be in a position to know through reflection, if we suppose that a lot of information that, for us, is available only through empirical investigation, were available to us otherwise. There is no reason to think that the answer to this question is relevant to what *we* can and cannot know a priori.

10

Two-dimensionalism, therefore, does not offer a plausible account of concepts normally associated with recognitional capacities. In closing this paper, I would like to connect this particular, negative point with the more general question concerning the epistemology of a priori knowledge with which I began.

As I said at the very beginning, it seems reasonable to attempt to ground the epistemology of the a priori in an account of conceptual thought. According to the two-dimensionalist version of this idea, the link is especially direct: possessing a concept just becomes a matter of being in a position to know by reflection alone a large number of non-trivial propositions. The failure of two-dimensionalism should, I think, alert us to the possibility that the link we are looking for might be much less direct than this.

According to current externalist views of conceptual thought, concept-possession is a matter of participating in a particular intellectual practice, situated in a particular environment. It is not a matter of being in a position to know any particular non-trivial propositions involving this or that particular concept. Nevertheless, this does not need to imply that concept-possession is easy: even though we may not be able to determine a fixed set of non-trivial propositions one needs to be in a position to know in order to possess any particular concept, membership in a conceptual practice certainly does seem to require a rich array of cognitive skills. In particular, I want to suggest that we cannot make sense of someone's being a member of our conceptual practice unless we can interpret her as, *on the whole*, sufficiently reliable in conforming to the norms fixing the application of our concepts, and, furthermore, reasonably capable of reflective reasoning in this respect.

²⁴ This kind of response was suggested to me by Cian Dorr.

This paper is not the place to elaborate on this idea. But I would like to suggest that something like this holistic constraint on conceptual thought might be a promising starting point for an epistemology of the a priori.

References

- Chalmers, D. (1996). *The conscious mind: In search of a fundamental theory*. New York: Oxford University Press.
- Chalmers, D., & Jackson, F. (2001). Conceptual analysis and reductive explanation. *Philosophical Review*, 110, 315–361.
- Chalmers, D. (2002a). Does conceivability entail possibility? In T. Gendler & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 145–221). New York: Oxford University Press.
- Chalmers, D. (2002b). On sense and intension. *Philosophical Perspectives*, 16, 2002. Retrieved May 9th 2007 from: <http://consc.net/papers/intension.pdf>
- Chalmers, D. (2002c). The components of content. Retrieved May 9th 2007 from: <http://consc.net/papers/content.pdf>
- Chalmers, D. (2006). The foundations of two-dimensional semantics. In M. Garcia-Carpintero & J. Macia (Eds.), *Two-dimensional semantics: foundations and applications*. Oxford University Press, 2006. Retrieved May 9th 2007 from <http://consc.net/papers/foundations.html>
- Evans, G. (1982). *The varieties of reference*. In J. McDowell (Ed.), New York: Clarendon Press-Oxford University Press.
- Jackson, F. (1998). *From metaphysics to ethics: A defense of conceptual analysis*. New York: Oxford University Press.
- Kripke, S. A. (1982). *Wittgenstein on rules and private language*. Cambridge: Harvard University Press.
- Putnam, H. (1975). The meaning of “meaning”. In K. Gudenson (Ed.), *Language, mind, and knowledge* (pp. 131–191). Minneapolis: Minnesota Studies in the Philosophy of Science-University of Minnesota Press.
- Williamson, T. (2003). Blind reasoning II: Understanding and inference. *Aristotelian Society*, 77(Suppl.), 249–293.
- Williamson, T. (2006). Conceptual truth. *Aristotelian Society*, 80, 1–41.
- Yablo, S. (2002). Coulda, woulda, shoulda. In T. Gendler & J. Hawthorne (Eds.), *Conceivability and possibility* (pp. 441–492). New York: Oxford University Press.