

SOURCES OF SELF-CONSCIOUSNESS¹

By José Luis Bermúdez

1. In thinking about the sources of self-consciousness we need to separate out two dimensions of the problem. There is, first, a genetic dimension. One might wonder about the origins of the capacity to think full-fledged self-conscious thoughts. That is to say, one might be asking about where this capacity comes from. How does it emerge in the normal course of human development? What are the genetic foundations on which it rests? Are there more primitive types of first-person contents from which full-fledged self-conscious thoughts emerge?

The problem of the sources of self-consciousness has an epistemic dimension as well as a genetic one. From the point of view of ontogeny or phylogeny, when we ask about the sources of self-consciousness we are really asking factual questions about where and when self-consciousness emerges in an individual or in a species. When we ask about the sources of self-consciousness from an epistemic point of view, however, the questions that emerge are of a fundamentally different type. Here we are taking the sources of self-consciousness to be the grounds of self-consciousness. What we want to know are the reasons for which self-conscious judgements and utterances are made. The questions that arise here include the following. How are such judgements justified? In virtue of what do they count as knowledge? In what senses are the different types of self-conscious judgements epistemically privileged? What role do self-consciousness judgements play in the cognitive economy?

My interest in this paper, in contrast, is with the role that the genetic dimension of self-consciousness plays in understanding the epistemology of self-consciousness. I will take as my foil a recent account of some key features of the epistemic dimension of a particular type of self-conscious judgement – the account offered by Christopher Peacocke in his recent book *Being Known* (Peacocke 1999). Working through some of the consequences and implications of Peacocke's account will bring out some important ways in which we need to draw upon the sources of self-consciousness in the genetic sense for a proper understanding of the sources of self-consciousness in the epistemic sense.

¹ Earlier versions of this paper were delivered at a conference on Self-knowledge held at the University of Fribourg in November 2000, at the Institut Jean Nicod in Paris and at the Philosophy Department of the University of Eastern Piedmont at Vercelli. I am grateful to audiences at those occasions for very helpful discussion and to Alan Millar and Brie Gertler for written comments

2. Two questions stand at the fore when we think about the epistemic grounds of full-fledged self-conscious judgements. The first question is: what are the grounds on which such judgements are made? The second question is: in virtue of what can such judgements qualify as knowledge? Clearly, answers to the second question will depend crucially on answers to the first. It is difficult, unless one leans towards the most extreme variety of epistemological externalism, to separate a judgement's status as knowledge from the grounds on which it is made.

Peacocke distinguishes two different types of self-conscious judgements according to their grounds. The first category comprises what he terms *representation-dependent* first-person judgements (I am slightly modifying his terminology and will henceforth use "representation-dependent" to abbreviate "representation-dependent first-person"). These judgements involve taking a first-personal representational content at face value and forming a corresponding belief. So, for example, if I am enjoying a visual state with the content that I am in front of the cathedral and, in virtue of that visual state, judge that I am in front of the cathedral, my judgement is representation-dependent.

The second category of full-fledged self-conscious judgements are *representation-independent*. These judgements are not made by taking a first-personal representational content at face value. Here are the examples that Peacocke himself gives:

- I am thinking about Pythagoras's Theorem
- I see a phone on the table
- I remember attending the birthday party
- I remember that Russell was born in 1872
- I am beginning to dream
- I fear that the motion will not be carried

The idea that there is a fundamental difference between these two types of judgement is really quite plausible. Judgements in the first category, those which are representation-dependent, are derived from corresponding first-person contents in a way that judgements in the second category (those which are representation-independent) quite simply are not. When I think about Pythagoras's theorem, for example, the content of my thought is the relation that holds between the lengths of the sides of right-angled triangles. The fact that I am thinking about that relation is not in any sense part of the thought.

So, the ensuing judgement that I am thinking about Pythagoras's theorem is *ampliative*. In contrast, when I look at the cathedral in front of me there is a sense in which the fact that I am in front of the cathedral is part of what I see. What is seen is not just the cathedral, but rather the state of affairs of the cathedral standing in certain specifiable spatial relations to *me*. So, the judgement is in some sense an endorsement of the content of my perception.

Both representation-dependent and representation-independent thoughts involve ascribing certain properties to oneself, but the properties are different in the two cases. Judgements that are representation-independent are typically psychological self-ascriptions, whereas the properties attributed to oneself in representation-dependent judgements will usually not be psychological. One would expect, therefore, that when these two thoughts qualify as knowledge there will be different account of what makes this the case.

According to Peacocke, representation-dependent self-conscious thoughts qualify as knowledge when the following three things hold

- (a) The thinker is taking a first-person representational state at face value
- (b) That first-person representational state is true
- (c) Certain relevant background informational conditions concerning the proper functioning of the faculty generating the state in question are fulfilled

The basic idea here, put more simply, is that I know that the cathedral is there in front of me if I see that there's a cathedral in front of me with all my perceptual systems functioning as they ought to be and that perception leads to my judgement in a suitably non-inferential way.

Clearly, no account along these lines can be given for self-conscious judgements which are not representation-dependent. There is no way in which taking my occurrent thoughts about Pythagoras's theorem at face value will justify my self-ascriptive judgement that I am thinking about Pythagoras's theorem.² For representation-independent judgements Peacocke offers what he calls the delta account.

² In fact, as Alan Millar pointed out to me, the very idea of taking a thought at face value in the way that one can take an experience at face value is deeply problematic. What makes it possible to take an experience at face value is that one can have the experience and not believe that things are the way that they appear to be, but nothing like this holds for first-order thoughts such as the thought that Pythagoras's theorem is true. If the thought is something one believes, as it clearly is in the case of Pythagoras's theorem, then there is clearly no gap between being aware of the thought and believing it

We can best get a grip both on what it is and on why he calls it the delta account by looking at the explanatory diagram he offers

INSERT FIGURE 6.1 FROM PEACOCKE 1999 P.274

An appropriately formed representation-independent self-ascription will always be a true self-ascription (although not necessarily a true belief).³ This is because the judgement that, for example, I am thinking about Pythagoras's theorem is co-conscious with my thought about Pythagoras's theorem, and hence the first-person element in the self-ascription will necessarily refer to the very individual thinking the first-order thought. But a belief can be true without counting as knowledge. What makes a representation-independent self-ascription knowledgeable? Peacocke's answer here is essentially that it is an *a priori* truth that a representation-independent self-ascription formed in the appropriate way will be true (with respect to its self-ascriptive component – see n.3) because it is an *a priori* truth that any occurrence of the first-person in thought will refer to the thinker of that thought. So, a representation-independent belief formed in the appropriate way will count as knowledge because it is *a priori* that it is true.⁴

3. The key theoretical difference between representation-dependent judgements and representation-independent judgements is that the former but not the latter involve *taking at face value* a representational content in which the first person features. It is natural to ask how we should understand these first person contents. We have already seen one example of a representation-dependent judgement – namely, the judgement that I am in front of a cathedral made on the basis of a perception of the cathedral. It is very clear that my perception of the cathedral does have an important first-person dimension. An integral part of what I see when I see the cathedral is that it stands in

³ The distinction between true self-ascriptions and true beliefs emerges when the subject self-ascribes a factive state. There is a sense in which I might truly form the judgement that I remember that the storming of the Bastille took place in 1790, although of course I can't really remember that the storming of the Bastille took place in 1790, given that it took place on 14 June 1789.

⁴ Again, we need to restrict the truth to the self-ascriptive component. It is not the case that all beliefs which are *a priori* true qualify as knowledge simply in virtue of being *a priori*. Thanks to Brie Gertler on this point

certain spatial relations to me. But what is the relation between the content of the perceptual state and the content of the self-ascription made on the basis of that perceptual state?

Peacocke himself maintains a studied neutrality here:

We can more generally consider examples in which the thinker's reason for making his judgement "I am F" is his being in some state, other than a belief state, which represents a certain content C as correct. In the class of examples I want to consider, the content C may, but need not, be the same as the content "I am F". Some theorists believe that perceptual states have non-conceptual contents which are distinct in kind from the contents which feature in beliefs. What I have to say is orthogonal to that issue. In the class of examples on which I want to focus, it will be the case that the content C, even if it is distinct from the content "I am F", is still one which stands in an implicational relation to the content "I am F". (p.264)

There are, as far as I can see, three possible positions that might be adopted. The first position is that the content C just is the conceptual belief content "I am F" – that is to say, there is a unity of content between the perceptual state and the self-ascription. The second position is that the content C is a first person non-conceptual content. According to this second position the self is explicitly represented in the content of the justifying perception in a way that underwrites, but somehow falls short of, the way in which the self is represented in the content "I am F". The third position is that content C is a non-conceptual content that has no first person dimension. The self is not represented explicitly at all.

Maintaining the third option appears to place pressure on the distinction between representation-dependence and representation-independence. If the self is not represented at all in the justifying perceptual state then it is hard to see how the content "I am F" can be reached by taking the content of that state at face value. It would be much more like a case of representation-independence, where the thinker capitalises on the implicational relation between the justifying perception (e.g. the perception that there is a telephone on the table) and the first person belief formed on its basis (e.g. I am looking at a telephone). It looks, therefore, as if we do not have a choice between the second and third options. If we do not go with the first option then the second option is mandatory.

But can we at least be neutral between the first and second understandings of the relation between justifying perception and first-person belief? Here matters are a little more complicated. According to the first understanding, it will be remembered, there is an identity of content between the target perceptual state and the corresponding self-ascription. That is to say, the self is represented the same

way in the original perception and in the self-ascription made on the basis of that perception. If the identity of content thesis is true then the self-ascriptive judgement simply duplicates the content of the original perception. The perceptual content is, roughly, *I am in front of the cathedral*. The self-ascriptive judgement endorses that and is best represented as the propositional attitude of belief being taken to the content *I am in front of the cathedral*. The difference between the self-conscious judgement and the perceptual state on which it is based is simply at the level of force.

The identity of content thesis has had, and continues to have, considerable currency in contemporary philosophy and mind and epistemology. It features most conspicuously in the epistemic theory of perception, according to which perception should be analysed in terms of the acquisition of dispositions to believe (Armstrong 1968), but it is also maintained by philosophers who reject the epistemic theory. John McDowell is a case in point:

Suppose someone is presented with an appearance that it is raining. It seems unproblematic that if his experience is in a suitable way the upshot of the fact that it is raining then the fact itself can make it the case that he knows that it is raining. But that seems unproblematic precisely because the content of the appearance is the content of the knowledge. (McDowell 1982 pp.213-214)

The identity of content thesis is a consequence of McDowell's brand of direct realism – it is because what we believe is what we see that what we know is available to us through perception.

The identity of content thesis raises an obvious question. What has happened to the familiar distinction between the content of belief and the content of perception? If, as many authors (including of course Peacocke himself) have argued, there is a fundamental difference between the content of perception and the content of belief (Peacocke 1992), then how can the content of a belief merely duplicate the content of a perception? We are familiar, for example, with the idea that propositional attitude content is digital, while perceptual content is analogue. Why is this not enough to block the identity of content thesis?

The simple answer is that the analogue/digital distinction does not map precisely onto the perceptual/doxastic distinction. Many theorists, once again including Peacocke, have proposed that propositional attitude contents can contain elements that are fundamentally perceptual in form – what are often termed perceptual-demonstrative modes of presentation (Peacocke 1983, 1992). The example

we are considering seems clearly to fall into this category. The doxastic *I am in front of the cathedral* is best viewed as containing the relevant cathedral under a demonstrative mode of presentation.

But adverting to perceptual-demonstrative contents in this way does not, I think, address the real worry with the identity of content thesis. In essence, the problem is not to do with how the cathedral is represented, but rather with how the self is represented. Whereas there is a certain plausibility in suggesting that there is a single perceptual-demonstrative mode of presentation of the cathedral that can stand in the content both of perception and of belief, it is hard to see how anything comparable holds for the representation of the self in perception and perceptually-based beliefs. The self does not feature in perception in anything like the way in which it does in beliefs – even when the beliefs in question are perceptually-based.

This conclusion follows, I think, from some basic reflections on the nature of belief. I am taking the content of belief to be an abstract object, something that can be believed by many different people and that can be expressed in a declarative sentence. It must be possible for different people, and indeed the same person at different times, to take different attitudes to the same belief-content. Partly this is a matter of it being possible for you to disbelieve something that I believe. Partly it is a matter of it being possible for me to believe something that for you is merely an object of hope – or for me to come to believe that something is the case when previously I simply feared that it might be the case. It's uncontroversial that any plausible general account of the objects of belief must allow this to be the case. And this has certain consequences. Some of them are fairly obvious. It cannot be the case, for example, that I could believe something that it is logically impossible for you to believe. But it also follows, and this is slightly less obvious, that the content of belief must be divorced from the grounds on which one might come to believe that content on any particular occasion. Let's suppose, for example, that I come to the true belief that I am in front of the cathedral because I see it in front of me, while at the same time you come to the erroneous belief that at that very moment I am not in front of the cathedral because you have been told by someone whom you believe to be a good authority on my movements that I am in fact in the museum. There is a single belief-content here which we are evaluating in different ways – namely, the belief-content that I am in front of the cathedral. And if we are to disagree about its truth-value then it cannot be the case that what I believe is tied to, or in any

way reflects, the fact that I arrived at the belief on the basis of perception – because if it was then it wouldn't be in conflict with your belief which you arrived at on the basis of testimony.

But if this is right then it cannot be the case that what I believe is what I see. What I believe, the content of my belief, must be something that is independent of the particular form in which my being in front of the cathedral became manifest to me. As I've stressed, the fact that I am looking at the cathedral is part of what I see – and the self is represented in perception as standing in certain specifiable relations to the cathedral associated with my physical location at the origin of a particular visual array. But then this cannot be the way in which the self is represented in my belief that I am in front of the cathedral. Any such representation of the self would be too specific to feature in a suitably objective belief-content.

4. So how is the self represented in belief? We can remain neutral on the question of whether the first-person concept should be analysed as the sense of the first-person pronoun (as suggested, for example, in Evans 1982) while still noting that there are significant commonalities between the way in which the self is picked out linguistically and the way in which it is represented in belief (and other propositional attitudes). The key to both the first-person pronoun and the first-person concept is the type/token distinction. The first-person pronoun is a linguistic type whose meaning is exhausted by the rule specifying that any appropriately uttered token of that type refers to its utterer. Similarly, the first-person concept is a type which, when it is tokened in thought, plays a familiarly distinctive functional role for the subject who tokens it. Adequate accounts of both the first-person pronoun and the first-person concept, therefore, will have to operate at two levels. One familiar way of doing this for the first person pronoun is to define the linguistic meaning (or *character*, as it is sometimes called) of the first person pronoun as a function from contexts to contents, where contents are functions from possible worlds to the individual of the context (Kaplan 1982, 505-507). One would expect a structurally similar account to hold for the first-person concept, so that the first-person concept is that concept whose tokening generates thoughts with a highly specific and irreducible functional role for the subject who has them. In the case of both pronoun and concept, the account is pitched at the level of the type rather than the token. The meaning of "I" is located at the level of the function from

contexts to individuals, rather than at the level of the individuals picked out by particular tokenings of “I”. So too is the nature of the first person concept fixed in a way that is independent of its tokening in particular individuals.

In light of this the difference between the representation of the self in perception and in thought becomes clear. Frege remarked in ‘The Thought’ that “everyone is presented to himself in a special and primitive way in which he is presented to no-one else” (Frege 1918, 12). Appearances suggest that Frege was intending this peculiar perspective to be reflected at the level of thought, and his comments have usually been interpreted as applying to the sense of the first-person pronoun. What I would like to suggest, however, is that Frege’s pithy phrase applies quite accurately to the way in which the self is presented in perception, but has no application at all to the way in which the self is represented in the content of belief and other propositional attitudes. There is nothing perspectival about the first person concept.

This basic thought can be motivated further by considering the range of representation-dependent thoughts available in other sensory modalities. I can come to the belief that I am near the church because I hear the ringing of the nearby church-bells. My auditory perception has a first-person content, namely, that I am not far from the source of the sound. This is part of the content of perception. It is not something that I infer from what I hear. So, the self must be represented in auditory perception – and yet in a fundamentally different manner from the way in which it is represented in beliefs based upon auditory perception. The same holds with even more force for self-conscious beliefs derived from somatic proprioception. Somatic proprioception is a source of first-person contents. Nonetheless, the self does not feature in the content of somatic proprioception in anything like the way it does in the contents of beliefs based upon somatic proprioception.

5. Representation-dependent judgements are judgements of the form “I am F” that are made by taking a certain justifying perceptual content C at face value. There is nothing wrong with the idea of a content being taken at face value. It is quite clear that when I form the belief that I am in front of the cathedral on the basis of a perception of the cathedral in front of me I am taking my experience at face value. I am not, for example, making any sort of inference from my perception to the presence of the

cathedral. The perception provides an immediate reason for the belief. But the question remains of how best to understand the content that is taken at face value. We have rejected the identity of content thesis, according to which there is no difference in content between the justifying perceptual state and the self-ascriptive belief which is formed on its basis. If the self is explicitly represented in perception then it cannot be represented there in the same way as it is in belief. Since we have also rejected the suggestion that there is no representation of the self in the content of the justifying perceptual state, we are forced to accept that representation-dependent judgements are grounded in a form of representation of the self that is essentially nonconceptual. What this means, of course, is that an account is urgently needed of the grounds of representation-dependent first-person judgements. We need an account of the first-person dimension of the target perceptual states.⁵

Let me approach this at a tangent. There is a general issue here about the extent to which perception is fully conceptual. We are familiar with the idea that perception is concept-laden – that the content of perception falls, as McDowell often puts it, within the space of reasons (McDowell 1994). We do not see objects in a neutral way, attending simply to their geometric, sensory and kinematic properties and then, on the basis of that neutral perception, work out how they fit into our conceptual categorisation of the world. The concepts are already built into the way in which we perceive the world, although nobody has ever managed to make very clear how this works. Nonetheless, to say that perception is concept-laden is not to deny that it has a nonconceptual dimension. In a sense it is completely obvious that perception must have a nonconceptual component because otherwise there would be no qualitative difference between the content of perception and the content of propositional attitudes. The critical question is not whether perception has a nonconceptual content, but whether it has a nonconceptual representational content – and several authors have argued persuasively that the

⁵ This is not a problem that arises simply for self-conscious judgements that are representation-dependent. It arises also for an important class of representation-independent self-ascriptions. Peacocke's list of representation-independent first-person judgements include autobiographical memories (my self-ascribing the memory of attending the birthday party) and self-ascriptions of perceptions (my self-ascribing the perception of the phone on the table). The second of these is quite clearly derived from a first-person content. What makes it true that I see that the phone is on the table is my occurrent perception of the phone on the table – and, as we've seen at some length with the example of my seeing the cathedral, my occurrent seeing that the phone is on the table is a first-person content. The same holds for the third example – my autobiographical memory of attending the birthday party (assuming, of course, that this memory is episodic). My episodic memory of attending the birthday is a memory of *my* having attended the birthday party. In both these cases of representation-independent judgement we need to distinguish, just as we did in the representation-dependent case, between the way in which the self is represented in the original perception or memory and the way in which it is represented in the doxastic states formed on the basis of that target state

nonconceptual content of experience is partially (and indeed perhaps completely) representational. These authors, however, have focused primarily on certain limited aspects of the content of perception (see, e.g. Peacocke 1992, Ch. 3). They have stressed, for example, the way in which the content of perception is analogue rather than digital; the way in which it presents distances in a unit-free manner; and the way in which it supports a fineness of grain that is arguably unavailable to conceptual thought.

However, these three features of the content of perception do not by any means exhaust its nonconceptual representational content. There is in the content of perception, I shall suggest, an ontogenetically primitive form of representation of the self. Full-fledged self-consciousness emerges, as I have stressed elsewhere (Bermúdez 1998), from a foundation of more primitive types of self-awareness – in the propriospecific dimension of exteroceptive perception, in somatic proprioception, in the ability to locate oneself spatially within the world, in the self-representation implicated in primitive forms of social interaction. These primitive forms of self-consciousness are prior to full-fledged self-consciousness *in the order of development* – and indeed are necessary conditions for the emergence of full-fledged self-consciousness. This genetic story by no means entails, of course, that once full-fledged self-consciousness has been attained, these primitive forms of self-consciousness have any further role to play in the cognitive economy. It could simply be that they just drop out of the picture, rather like some of the basic perceptual expectations that very young infants have about the behaviour of distal objects.

Yet, there is a fundamental disanalogy between the perception of distal objects and the perception of the self in this respect. It is plausible to think that young infants parse the perceived array into object-like segments that obey certain basic geometrical and kinematic principles. These have been studied at some length by developmental psychologists. With the acquisition of concepts and language, however, this way of perceiving the world is overridden by the conceptual seeing-as that we have already briefly mentioned. Little if anything remains of the infant perceptual universe in adult perception. But the perception of the self is not at all like that. The fundamental differences in content between first-person representational states and first-person beliefs based on those states point us towards the persistence of a fundamentally nonconceptual representation of the self in ordinary perception. This provides the key to the way in which perceptual states provide grounds for the

representation-dependent first-person judgements that are made by taking them at face value (although not, as we have seen, by duplicating them).

6. The best account of the nonconceptual representation of the self in visual perception comes, I think, from J. J. Gibson's ecological approach to perception and I will simply briefly sketch some of the principal respects that he has drawn to our attention in which the self is represented in visual perception (Gibson 1979).

Gibson stresses certain peculiarities of the phenomenology of the field of vision. Notable among these is the fact that the field of vision is bounded. Vision reveals only a portion of the world to the perceiver at any given time (roughly half in the human case, due to the frontal position of the eyes). The boundedness of the field of vision is part of what is seen, and the field of vision is bounded in a way quite unlike the way in which spaces are bounded within the field of vision. The self appears in perception as the boundary of the visual field – a moveable boundary that is responsive to the will.

The boundedness of the visual field is not the only way in which the self becomes manifest in visual perception according to Gibson. The field of vision contains various body-parts which hide, or *occlude*, the environment. The nose is a particularly obvious example, so distinctively present in just about every visual experience. The cheekbones, and perhaps the eyebrows, occupy a slightly less dominant position in the field of vision. And so too, to a still lesser extent, do the bodily extremities, hands, arms, feet and legs. They protrude into the field of vision from below in a way that occludes the environment, and yet which differs from the way in which one non-bodily physical object in the field of vision might occlude another. All objects, bodily and non-bodily, can present a range of solid angles in the field of vision (where by a solid angle is meant an angle with its apex at the eye and its base at some perceived object), and the size of those angles will vary according to the distance of the object from the point of observation. The further away the object is, the smaller the angle will be. But the solid angles subtended by occluding body-parts cannot be reduced below a certain minimum. Perceived body-parts are, according to Gibson, 'subjective objects' in the content of visual perception.

These *self-specifying structural invariants* provide only a fraction of the self-specifying information available in visual perception. There are two more important types of self-specifying information.

The mass of constantly changing visual information generated by the subject's motion poses an immense challenge to the perceptual systems. How can the visual experiences generated by motion be decoded so that subjects perceive that they are moving through the world? Gibson's notion of visual kinesthesia is his answer to this traditional problem. Whereas many theorists have assumed that motion perception can only be explained by the hypothesis of mechanisms which parse cues in the neutral sensations into information about movement and information about static objects, the crucial idea behind visual kinesthesia is that the patterns of flow in the optic array and the relations between variant and invariant features make available information about the movement of the perceiver, as well as about the environment.

As an example of such a visually kinesthetic invariant, consider that the optical flow in any field of vision starts from a centre, that is itself stationary. This stationary centre specifies the point that is being approached, when the perceiver is moving. The aiming point of locomotion is at the vanishing point of optical flow. Striking experiments have brought out the significance of visual kinesthesia. In the so-called 'moving-room' experiments, subjects are placed on the solid floors of rooms whose walls and ceilings can be made to glide over a solid and immovable floor. If experimental subjects are prevented from seeing their feet and the floor is hidden, then moving the walls backwards and forwards on the sagittal plane creates in the subjects the illusion that they are moving back and forth. This provides strong support for the thesis that the movement of the perceiver can be detected purely visually, since visual specification of movement seems to be all that is available. An even more striking illustration emerges when young children are placed in the moving room, because they actually sway and lose their balance.

A further important form of self-specifying information is available to be picked up in the field of vision, according to the theory of ecological optics. This is due to the direct perception of a class of higher-order invariants which Gibson terms affordances. It is in the theory of affordances that we find

the most sustained development of the ecological view that the fundamentals of perceptual experience are dictated by the organism's need to navigate and act in its environment, because the animal and the environment are complementary. According to Gibson, the possibilities which the environment affords are not learnt through experience. Nor are they inferred. They are directly perceived as higher-order invariants. And of course, the perception of affordances is a form of self-perception - or, at least, a way in which self-specifying information is perceived. The whole notion of an affordance is that of environmental information about one's own possibilities for action and reaction.

7. Gibson's characterisation of the phenomenology of visual perception offers important insights into how the self is represented nonconceptually in the content of perception. To return to our earlier discussion of representation-dependent judgements, it is the "ecological self" which features in the content of basic perceptual states and which provides the epistemic grounds for self-conscious judgements made on the basis of those perceptual states. This brings us to the second of the two questions identified at the beginning of section 2. To understand the epistemology of full-fledged self-conscious judgements it is not enough to understand the grounds on which those judgements are based. We need also to understand why, when those judgements count as knowledge, they should do so.

It will be helpful to begin by examining how representation-independent judgements are grounded. Representation-independent judgements are, in essence, self-ascriptions of mental states and it follows from the definition of representation-independence that the self-ascriptive component of the representation-independent judgement is not grounded in the content of the state that the subject self-ascribes. A representation-independent judgement is made when a subject makes the transition from ϕ -ing that p to the self-ascriptive judgement "I am ϕ -ing that p ", but there is nothing in the content p that corresponds to the first-person concept that appears in the self-ascriptive judgement (in obvious contrast to representation-dependent judgements). So how does the subject make the transition? The simple answer is that the transition is available to subjects who have a basic mastery of what might be termed a simple theory of introspection (by analogy with the simple theory of perception and action discussed in Evans 1982 and Campbell 1994). In this case the simple theory of introspection amounts to nothing more than some level of mastery of the *a priori* link between being aware of a thought and

it being the case that one is thinking it – which is, of course, the same *a priori* link that governs the status as knowledge of the resulting representation-independent judgement (according to Peacocke's delta account).⁶ What we see in this case is that the subject's reason for forming the judgement is an appreciation of precisely the factor that qualifies the judgement as knowledge. Let us describe epistemic situations of this type as being *canonically internalist*.

In order to see how representation-dependent judgements can be just as canonically internalist as representation-independent judgements, we need to explore the relation between representation-dependence, representation-independence and the phenomenon which Sydney Shoemaker has described as immunity to error relative to the first person pronoun. As Shoemaker puts it (with reference to statements rather than judgements):

. . . to say that a statement "a is \emptyset " is subject to error through misidentification relative to the term 'a' means that the following is possible: the speaker knows some particular thing to be \emptyset , but makes the mistake of asserting "a is \emptyset " because, and only because, he mistakenly thinks that the thing he knows to be \emptyset is what 'a' refers to. (Shoemaker 1968, pp.7-8)

Immunity to error through misidentification relative to the first person pronoun (henceforth abbreviated to: immunity to error through) is a property of judgements that are not based upon an act of identification of a particular object as oneself. They are judgements that one has a certain property based upon grounds such that knowing that something has that property is *ipso facto* to know that one has that property oneself.

It is clear that representation-independent self-ascriptions are immune to error.⁷ If I form the judgement that I am thinking about Pythagoras's theorem because I am aware of thinking about Pythagoras's theorem then that judgement is not based upon an identification of myself as the author of the relevant thought. Rather, knowing that there is a thought about Pythagoras's theorem just is knowing that that thought is a thought which I am having. It is equally clear that representation-dependent thoughts share this feature of being immune to error through misidentification. If I form the judgement that I am in front of the cathedral because I see the cathedral in front of me then this is not because I form the belief that someone is in front of the cathedral and then identify that person as

⁶ The simple theory of introspection might also contain such basic principles as that, if one wants to find out whether one believes that *p*, the best way to proceed is to determine whether *p* is in fact the case.

⁷ As Peacocke himself points out (Peacocke 1999 p 270)

myself. Knowing on the basis of perception in this way that someone is in front of the cathedral just is knowing that I myself am in front of the cathedral.

The immunity to error through misidentification of first-person judgements is intimately linked with their status as knowledge, in virtue of the connection between immunity to error through misidentification and the defining property of the first-person concept, namely, that of picking out the thinker of any thought in which it features (in the appropriate way). In a judgement of the form “I am F” which is immune to error through misidentification the only possibility of error is with respect to the predicative component. And hence, provided that the information grounding the predicative component has been acquired in an appropriate manner, one would expect the judgement to count as knowledge. It is natural to conclude, therefore, that an explanation of the immunity to error of first-person judgements will provide the key to explaining their status as knowledge.

Nonetheless, representation-independent and representation-dependent judgements secure their immunity to error through misidentification in fundamentally different ways. In the case of representation-independent judgements it is precisely the feature which secures their status as knowledge that underwrites their immunity to error through misidentification – namely, the *a priori* link between being aware of a thought and it being the case that one is thinking it. In the case of representation-dependent judgements, on the other hand, immunity to error through misidentification has to be a function of the states on which they are based. Let me offer the following conjecture. Representation-dependent judgements are immune to error through misidentification because they involve taking at face value first-person contents which are themselves immune to error through misidentification. It is easy to see how this works in the case of judgements that are based on perceptual states, such as the example we have been considering of my forming the belief that I am in front of the cathedral in virtue of seeing the cathedral in front of me. When I see the cathedral in front of me I do not see that the cathedral is in front of someone and then work out that I am that person. Seeing that someone is in front of the cathedral just is seeing that I am in front of the cathedral. And, given this, it is hardly surprising that the resulting judgement is itself immune to error through misidentification. It is plausible to conjecture that a broadly similar account is true of other types of

representation-dependent judgements (such as those based on the deliverances of somatic proprioception, for example).⁸

In a sense, however, this merely pushes the problem one step further back. Representation-dependent judgements are immune to error because they are grounded in first-person states that are themselves immune to error through misidentification. But in virtue of what are those first-person states immune to error through misidentification? Immunity to error through misidentification is a formal property of representational states. We need to know what is responsible for it. The answer, in brief (and at least as far as visual perception is concerned), lies in the simple fact that the information about the self provided within the content of visual perception is such that it is manifest to the perceiver that it could not be about any creature except the self. To return to the earlier example, the information that I have about my own spatial position relative to the cathedral could not be information about anyone else's spatial position relative to the cathedral. This is because the body is manifest as the originating point of perception in such a way as to ensure that the self-conscious judgement that I am perceiving the cathedral is immune to error through misidentification.

In order to get a better understanding of why this is the case we need to revert to the insights into the phenomenology of visual perception gleaned from Gibson. The self is presented in perception as a distinctive and peculiar form of object in virtue of certain features of the content of visual perception. Some of these derive from the fact that the world is presented to the perceiver largely in terms of the affordances it offers. Others arise from the features that lead Gibson to describe perceived body parts as subjective objects in the field of view – the fact, for example, that body parts can vary in perceived size only within a limited range. The way in which the body appears as the frame and boundary of the visual array is also extremely important. The embodied self is the originating point of perception in a much richer sense than the camera lens is the originating point of a photograph. The way in which the embodied self appears in perception, therefore, reflects the immunity to error of the relevant perceptual content and, by extension, that of the resulting belief content.

⁸ For further discussion of the immunity to error of judgements based upon somatic proprioception, see Evans 1982, Cassam 1995 and Bermúdez 1998

None of these features of the content of visual perception are transmitted to the content of self-conscious judgements made on the basis of visual perception, but nonetheless they underwrite the status as knowledge of those self-conscious judgements. They do this because they make it the case that the perceiver's reason for forming the judgement includes an appreciation of precisely those factors that qualifies the judgement as knowledge. The immunity to error of representation-dependent judgements is a function of the fact that the bodily self is the origin of perception. The distinctive way in which the bodily self is presented in visual perception ensures that this is manifest to the perceiver – it is part of what is perceived. This ensures that, to return to the terminology introduced earlier, epistemic situations leading to the formation of representation-dependent judgements are *canonically internalist*.

8. By way of conclusion, then, let me return to the general issues with which I began. As I mentioned at the beginning of the paper, we can understand the sources of self-consciousness in two different ways – in an epistemic way, as the grounds of self-consciousness and in a genetic way, as the origins of self-consciousness. What I've tried to argue is that the sources of self-consciousness in the genetic sense are ultimately also the sources of self-consciousness in the epistemic sense. The primitive foundations from which self-consciousness emerges in the course of cognitive development are also the foundation for the epistemic status of full-fledged self-conscious thoughts. My discussion has been restricted to the relatively straightforward example of visual perception and the Gibsonian account of its first-person content. Obviously a fuller development of this general theme will require extending the analysis to other types of first-person content. I hope that I have done enough to make it seem that this is a task worth undertaking.

Department of Philosophy

University of Stirling

Stirling FK9 4LA

REFERENCES

- Armstrong, D. M. 1968. *A Materialist Theory of Mind*. London: Routledge.
- Bermúdez, J. L. 1998. *The Paradox of Self-Consciousness*, Cambridge MA: MIT Press.
- Campbell, J. 1994. *Past, Space and Self*. Cambridge MA: MIT Press.
- Cassam, Q. 1995. 'Introspection and Bodily Self-Ascription' in J. Bermúdez, A. J. Marcel and N. Eilan (eds.), *The Body and the Self*, Cambridge MA: MIT Press.
- Dancy, J. 1986. *Perceptual Knowledge*. Oxford: Oxford University Press.
- Evans, G. 1982. *Varieties of Reference*. Oxford: Oxford University Press.
- Frege, G. 1918. 'Thoughts' in G. Frege, *Logical Investigations*, edited by P. T. Geach, Oxford: Basil Blackwell, 1977.
- Gibson, J. J. 1979. *The Ecological Approach to Visual Perception*, Boston: Boston Mifflin.
- Kaplan, D. 1989. 'Demonstratives', in J. Almog, J. Perry and H. Wettstein (eds.), *Themes from Kaplan*, Oxford: Oxford University Press.
- McDowell, J. 1982. 'Criteria, Defeasibility and Knowledge', reprinted as Essay XII in Dancy 1986.
- McDowell, J. 1994. *Mind and World*. Harvard: Harvard University Press.
- Peacocke, C. 1992. *A Study of Concepts*. Cambridge MA: MIT Press.
- Peacocke, C. 1999. *Being Known*. Oxford: Oxford University Press.
- Shoemaker, S. 1968. 'Self-reference and self-awareness', in *The Journal of Philosophy* 65, 555-567.
Page references to Shoemaker 1984.
- Shoemaker, S. 1984. *Identity, Cause and Mind*. Cambridge: Cambridge University Press.