

On the Explanation of Intentionality¹

Tim Crane, Central European University

Contemporary philosophy of mind tends to assume that to explain intentionality we must give an answer to what I will call the ‘question of aboutness’: what makes it the case that any intentional mental state is about something? Two further assumptions usually accompany this question. The first is that a satisfactory answer must be given without using any intentional notions. So, for example, it would not a satisfactory answer to the question to say that a thought is about an object O when it represents O, since ‘represents’ is an intentional notion. The second is that the answer should be, in a certain sense, general: it should apply to all kinds of intentional phenomena. So it would not be an adequate answer to say that ‘perceptual aboutness’ works in one way, and the ‘aboutness of thought’ in another, without any ultimate explanation of why these are both manifestations of intentionality.

My primary aim in this paper is to dispute the assumption that anything that could count as a plausible explanation of intentionality must answer the question of aboutness, so conceived. I call this the ‘aboutness assumption’. My secondary aim is to sketch an alternative way of thinking about intentionality and its place in the natural world.

I will first outline the basic ideas behind the question of aboutness. I will then look at one influential argument for the aboutness assumption, from Hilary Putnam’s, *Reason, Truth and History* (1981). I claim that this argument requires some hidden premises if it is to be convincing, but if you already accept those premises, then you don’t need Putnam’s argument. This suggests the aboutness assumption does not emerge out of quite general philosophical curiosity, but relies on much more specific and controversial ideas. Of course disputing one argument for the aboutness assumption does not refute it; but insofar as its premises are shared by other

¹ Lectures based on this material have been given at the Salzburg conference on Descriptive Psychology in 2018, at the University of Tampere, Eötvös Lóránd University in Budapest, the University of Antwerp and the University of Nottingham (the 2017 Lumsden Lecture), University of Hamburg and the University of Vienna. Thanks to participants on those occasions and especially to Max Kölbel for forcing to me to make myself clearer, to David Papineau for many enlightening and invigorating discussions of this topic, and to Kati Farkas and Craig French for their insights. The ideas for this paper were born in the ‘New Directions’ project in Cambridge, from 2015-2017.

arguments for the assumption, my argument should be of wider interest. In the final section I will outline my alternative approach to the explanation of intentionality.

1. The question of aboutness

The problem of intentionality is often posed in terms of the question: what makes something an object of our thought? Imogen Dickie puts it in this way:

When I look at my dog and think a thought I would express by saying 'he is asleep', there is an intuitive sense in which my thought is 'about' the dog rather than any other individual... But what does this intuitive aboutness consist in, and how is it secured? What makes a thought or utterance about a particular ordinary thing? (Dickie 2015: 1)

Dickie claims that when she has a certain thought occasioned by looking at her dog, this thought is 'intuitively' about her dog. It's easy to understand the situation: Dickie looks at her dog, and thinks *he is asleep* (or, as she more fastidiously puts it, presumably to avoid the implication that all thinking must be in words) she thinks a thought that she would express by saying 'he is asleep'. Obviously, her thought is about her dog. It is her dog that she is thinking about — she herself says as much. What is the problem supposed to be?

Dickie asks what the 'aboutness' of her thought consists in. Once again this question seems to have a straightforward answer. She is looking at her dog, she recognizes the dog and when she thinks *he is asleep*, she intends to be thinking about the dog. Why isn't that enough of a description of what the aboutness consists in? Of course, we can also ask what seeing consists in, and this is a largely empirical question — Dickie will presumably agree, which indicates that this is probably not the sort of answer she is looking for.

But she then broadens the focus with a more general question, not just about this specific thought which she had when looking at her dog, but about thoughts and utterances in general: what makes it the case that any thoughts or utterances are about particular things? Note how much more ambitious this question is, even if it is restricted to particular things. Two features of this question stand out: first, that it is about both thoughts *and* utterances about particular things (singular thought, singular reference); and second, that it seems to require a quite general answer.

Presumably, an answer of the form, 'sometimes I am looking at the thing I am thinking about; sometimes I am imagining it; sometimes it is because I have learned about it from other people, etc.' would not be the sort of thing that Dickie wants. This is, I think, partly because these answers already employ the idea of aboutness, and partly because it is just a list of ways in which we can be thinking about things, and such a list does not tell us what aboutness 'consists in'.

So there are two constraints on an answer to this more ambitious question: the answer should not employ any intentional notions, and the answer should be general, applying to all kinds of intentionality. These assumptions are widely shared in today's philosophy of mind. Nicholas Shea begins his book on representation in cognitive science as follows:

So, the question remains: how do mental states manage to be about things in the external world? That mental representations are about things in the world, although utterly commonplace, is deeply puzzling. How do they get their aboutness? ... This is an undoubted lacuna in our understanding, a void hidden away in the foundations of the cognitive sciences. (Shea 2018: 5).

And in his recent book on the metaphysics of representation, Robert Williams begins by saying representation is 'spooky' and poses the question of his book in entirely general terms: 'What is representation? How do the more primitive aspects of our world come together to generate it?' (Williams 2020:1). Again, the question is not about perception, or about thought, or about imagination, or memory, or decision ... but about representation in general.

This is, of course, the standard approach to the problem of intentionality in analytic philosophy. (Full disclosure: I put it in something like this way myself (Crane 2016: 7-8)). But it is worth asking about the assumptions that lie behind this approach. Why should we expect the same answer to be given to explain the aboutness of thoughts, experiences and utterances? Why should we expect our account of aboutness to be perfectly general? Why should we expect that an adequate answer can only be given in non-intentional terms? What are the reasons for taking these starting points? In the next section I will examine a famous argument of Hilary Putnam's, which seems to me to encapsulate a particularly influential way of thinking about intentionality, and to offer answers to these questions.

2. Putnam's ant

Putnam's *Reason, Truth and History* starts with a striking and memorable image:

An ant is crawling on a patch of sand. As it crawls, it traces a line in the sand. By pure chance the line that it traces curves and recrosses itself in such a way that it ends up looking like a recognizable caricature of Winston Churchill. Has the ant traced a picture of Winston Churchill, a picture that depicts Churchill? (Putnam 1981: 1)

The obvious answer is to Putnam's question, as he himself insists, is no. The lines on the sand just happen to resemble Churchill; they don't depict Churchill. The lines in the sand are no more a picture of Churchill than a slice of bread resembling the outline of Australia is a picture of Australia. 'Sometimes we see a cloud that's dragonish' says Shakespeare's Mark Antony (*Antony and Cleopatra* IV: 15), and he is right. But the fact that the cloud looks like a dragon does not make it a picture of a dragon, or indeed a dragon-representation of any kind.

The lines on the sand look like Churchill, but they do not represent him. It is a familiar point that resemblance or similarity is not sufficient for representation. It is often claimed that resemblance is not necessary either. As a point about representation in general, this is easily demonstrated by the phenomenon of linguistic representation, since words do not resemble what they represent. But this is only relevant because we are talking about representation in general. After all, it might be relevant to a theory of *pictorial* representation that if a certain kind of picture depicts someone, then it usually resembles them in some way; so the fact that philosophers move so quickly to conclude that resemblance is not necessary for representation indicates that they are looking for a *general* account of representation, not just an account of pictorial representation.

Having established that resemblance is neither sufficient nor necessary for representation, Putnam asks:

If similarity is not necessary or sufficient to make something represent something else, how can anything be necessary or sufficient for this purpose? How on earth can one thing represent (or 'stand for', etc.) a different thing? (Putnam 1981: 2)

Once again there is a natural response. The lines come about by pure chance, since the ant has no intention to draw a picture of, or to represent, Churchill. The ant has no idea who Churchill is — if ants have any ideas about anything at all, they surely do not have ideas of Churchill, and so they cannot intend to represent Churchill by drawing a picture of him. Drawing a picture must at the very least involve the intention to draw a picture. Ants probably have no intentions, but in any case they certainly have no picture-drawing intentions.

Putnam's response is that his original question can be asked again about intention and thought:

To have the intention that anything ... should represent Churchill, I must have been able to think about Churchill in the first place. If lines in the sand, noises, etc., cannot 'in themselves' represent anything, then how is it that thought forms can 'in themselves' represent any thing? Or can they? How can thought reach out and 'grasp' what is external?
(Putnam 1981: 2)

The general point is clear enough: if lines in the sand cannot represent anything, then how can thoughts represent anything? Yet this question also sounds strange. After all, it might be said that the parallel is not quite right, since unlike lines in the sand, the whole nature of thoughts is to represent things — that's the point of thought. So why is it that we can ask the same question about thought that we can ask about lines in the sand?

If we look more closely at the passage we see that Putnam does not just say 'thought', but 'thought *forms*'. What are thought forms? The introduction of the word 'form' is not an accidental slip of the pen on Putnam's part; it embodies an assumption that is needed to make the parallel with lines on the sand intelligible. For the only way to make sense of the parallel is if there is something that relates to thoughts as mere shapes relate to meaningful words, or as lines in the sand relate to a picture in the sand. Putnam himself does not defend the hypothesis that there are such things as thought forms, but without it the argument cannot move forward.

I will return to 'thought forms' shortly. But for the moment, let's go along with the assumption that the same question can be asked about thoughts (e.g. intentions) which Putnam asked about the lines on the sand: what makes thoughts represent anything at all? Given this assumption, Putnam can then claim that 'what goes for physical pictures also goes for mental

images, and for mental representations in general' (Putnam 1981: 3). And from this he concludes that there is nothing about any mental or physical representation, considered in themselves — or intrinsically — that makes them represent what they do.

Putnam's overall argument can therefore be represented as follows:

- (1) No physical representation represents intrinsically
- (2) What goes for physical representations also goes for mental representations
- (3) Therefore no mental representation represents intrinsically
- (4) Therefore mental representation must represent via some non-intrinsic connection

Causation is the leading candidate for a non-intrinsic connection which is proposed to underpin representation. And of course, one of Putnam's aims in the beginning of *Reason, Truth & History* is to use the causal theory of representation to argue against the coherence of the idea that we might be brains in vats. I will not discuss this additional argument here, since my interest is only in his argument which begins with the thought-experiment about the ant and ends with the conclusion labelled (4) above.

Premise (1) is supposed to be illustrated by the thought experiment about the ant. But a couple of things need to be clarified. First, 'intrinsically' is not explained — for the moment we can take it to mean, 'in itself', or 'in isolation from anything else' or 'not in relation to anything else'. Second, strictly speaking the thought experiment is only about pictures; but it is easy to draw the same conclusion if the ant had made the shapes WINSTON CHURCHILL on the sand, and maybe we can generalise from that to all physical representations. Or maybe not; this is not my focus here. Let's agree for the sake of argument that premise (1) is true.

Premise (3) follows from (1) and (2) obviously, and (4) follows from (3) on the assumption that mental representations represent via some kind of connection. The crucial premise of the argument, then, is (2). To assess the validity of the argument, we have to make precise what 'goes for' means; and to assess its soundness, we have to figure out whether premise (2) is true. So what does (2) mean?

The most plausible reading of (2) is something like this: what is true about physical representations *with respect to their representational powers* is also true of mental representations. After all, it cannot be true in general that everything that is true of physical

representations is true of the mental; the issue here is rather the representational character of these representations, and where it comes from. So what is the reason for believing in (2), so understood?

Given that a physical representation is a representation in which we can distinguish representational properties from non-representational properties, the truth of (2) depends on whether there are 'thought forms' as described above: something that, by definition, represents but has properties other than its representational properties. So our question becomes, why believe that there are such things as thought forms?

In the tradition in which Putnam is writing, I think it is plausible to say that there are two broad reasons for believing in thought forms. The first is physicalism; the second is what I call a semantic picture of all representation. The first reason is familiar; the second perhaps less so. I will take them in turn.

3. Thought forms and physicalism

Physicalism says that everything is physical, or that everything is determined by the physical, or supervenes upon the physical, or constituted by the physical, or grounded in the physical. These formulations are perhaps progressively less clear, but not in a way that matters for my purposes here. Nor does it really matter here how exactly we define the physical, though that is important in other contexts. What matters is that the subject-matter of physics, whatever that precisely is, has ontological priority — everything else has to be understood or explained relative to this subject-matter, to the physical. So physicalism requires that anything which is not obviously part of the fundamental physical world must be explained in physical terms.

This reductionist doctrine has dominated analytic philosophy of mind in the last half century or more. Intentionality and consciousness are typically taken as the central mental phenomena which do not seem to fit easily into the world as conceived by physics. Jerry Fodor, in a well-known passage, writes:

I suppose that sooner or later, the physicists will complete the catalogue they've been compiling of the ultimate and irreducible properties of things. When they do, the likes of spin, charm and charge will perhaps appear upon their list. But aboutness surely won't;

intentionality simply doesn't go that deep... If aboutness is real, it must really be something else. (Fodor 1987: 97)

If we start with this metaphysical assumption, then it is natural to see how the introduction of thought forms is required for an understanding of intentionality. The basic idea is that you postulate something in the brain of the thinker, call this the 'vehicle' of representation, and then explain its representational content in terms which are acceptable to physicalism (e.g. in causal or informational terms). In fact this is what Fodor (1987) himself does; as does Hartry Field in a famous paper (1978). Fodor's view is that these vehicles of representation are sentences in a 'language of thought' (1987, appendix). Physicalism gives you a reason to believe in thought forms, whether conceived as sentences in a language of thought, or in some other way.

Understood in this way, it's clear that what motivates the postulation of thought forms is physicalism, not some general concern about the possibility of representation. (Cf. here Shea: 'The physical and biological sciences offer no model of how naturalistically respectable properties could be like that' (Shea 2018: 5)). Putnam's argument, however, was supposed to be a quite general argument — the argument does not anywhere state that it is assuming physicalism. So although appealing to physicalism does give a reason to believe in thought forms, it makes the Putnam argument redundant. If you are a physicalist you will have already rejected 'intrinsic intentionality' (perhaps for Fodor's reasons) and maybe you will postulate something like thought forms instead. So you do not need Putnam's argument.

The starting point of Putnam's argument — as expressed too in the quotes from Dickie, Shea and Williams — is that we will find intentionality philosophically problematic ('puzzling', 'spooky'). As I have reconstructed Putnam's actual argument, the thing that is immediately problematic is the relationship between a *picture* and what it represents. We can then extend this sense of a problem to symbols (e.g. words) too. The natural response is that pictures and words get their representational powers from the states of mind of those who use them or interpret them. To ask then, 'but where do those states of mind get *their* representational powers from?' requires the assumption that 'what goes for physical representations goes for mental representations'. But independently of a belief in physicalism, it really isn't obvious that what goes for physical representation goes for mental representation.

The point can be brought out with an analogy. It is sometimes said that Descartes's appeal to mental substance does not answer the question about the metaphysical nature of the mind, since there is as much of a question about how a mental substance can think (or be conscious) as there is about how a material substance can. But this relies on a misunderstanding of the Cartesian notion of substance. For Descartes a substance is characterised in terms of its fundamental attribute. The fundamental attribute of material substance is extension, and the fundamental attribute of mental substance is thought. It's not as if we have some initial idea of what a kind of substance might be, and then can ask 'but how can that substance have that attribute?'. The essence of the kind of substance in question is characterised in terms of that attribute. So for Descartes, it makes no more sense to ask, 'how can a mental substance think, or be conscious?', than it does to ask 'how can a material substance be extended?'.²

What might sustain this misunderstanding is the idea that Cartesian substance is a kind of 'stuff'. It is as if physical or material objects are made out of one kind of stuff, and mental objects (on this understanding) are made out of another. Frank Jackson christens this stuff 'ectoplasm', 'understood as a kind of stuff incompatible with the physicalists' view of what kinds there are—perhaps the stuff out of which thoughts are made according to Descartes)' (Jackson 1998:15).²

But on Descartes's actual view, thoughts are not made of anything at all (Robinson 1993: 163). This is one difference between material substance and minds: material substance has parts, it is made of stuff, whereas minds have none. As Descartes himself says in the 6th *Meditation*, 'when I consider the mind, that is, when I consider myself in so far only as I am a thinking thing, I can distinguish in myself no parts'. If minds were made of some kind of ghostly mental 'stuff' then they would have parts. And if this were so, it would be intelligible how one could ask the same question about mental substance as one does about matter: how can this kind of stuff think? But since no-one has actually held this view of mental substance, the position is a straw one.

This is not to say the Cartesian notion of substance is unproblematic — for one thing, you might think the very idea of substance is misconceived — but only that the problems with are not posed by asking 'how *can* a mental substance think?'. Similarly, I am not saying that there are no

² The word 'ectoplasm' was coined in 1994 by French physiologist Charles Richet, president of the U.K.'s Society for Psychical Research. It is perhaps worth pointing out that this 'real' ectoplasm is a kind of matter: 'white, milky, slimy, smelling of ozone. It was solid, but also sometimes appeared as vaporous, emanating from the mouths, ears, and noses of physical mediums, usually in a darkened room and surrounded by fellow spiritualists' (Blom 2010: 168).

questions to be raised about how intentionality is possible, but only that you should not assume that these questions are raised by thinking that what ‘goes for’ physical representation goes for mental representation too. Physicalists who look for a theory of representation may well make this assumption, but the assumption is not a quite general requirement, independent of metaphysical views like physicalism.

4. Thought forms and the semantic conception of intentionality

The second reason for believing in thought forms is what I call a semantic conception of intentionality. A semantic conception of intentionality explains it in terms of notions like reference, truth, predication, proposition etc. — notions which are drawn from semantic theorising. I use the word ‘semantic’ here as it is used in the discipline of formal semantics, or when we talk about a semantics for a logical calculus. The central concept in semantics is truth; the task of a compositional semantics, for example, is to explain how the truth of sentences is determined by the semantic properties or relations (truth, reference, satisfaction etc.) of their parts (Speaks 2021: §2). So I do not use the word ‘semantic’ to refer (as some do) to any kind of relationship between language (or the mind) and reality, and I do not use ‘semantic’ and ‘intentional’ as synonyms (for me ‘semantic content’ and ‘intentional content’ are not necessarily the same thing). This is partly terminological, but partly substantive. And I think my use of the word ‘semantic’ reflects the dominant disciplinary usage in linguistics and philosophy of language.

The idea that intentionality should be understood in semantic terms is therefore a substantial thesis. It implies that we should understand mental representation as we understand linguistic representation: by relating certain kinds of entities (words) to other kinds (things in the world). Jeff Speaks, in an authoritative survey of recent work in philosophical semantics (or the theory of meaning) talks of the ‘core feature of classical semantics—for example, the assignments of entities as meanings, or the idea that meaning centrally involves word-world relations’ (Speaks 2021: §2.2.4).

In the tradition Speaks is describing, which arguably begins with Frege and Bolzano, and would list Church, Tarski, Montague, Lewis and Davidson as among its most distinguished contributors, different kinds of entities are typically assigned to different kinds of expressions. Names are assigned to objects, sets or functions are assigned to predicates, and truth-values, or

sets of worlds, or states of affairs/situations (etc.) are assigned to sentences, relative to contexts or indices. Assignment is the main business of semantics.

It is one thing, however, to assign meanings to words and sentences, and another to say how any words get their meanings in the first place: i.e. to explain how any assignment is correct. These days the latter enterprise is sometimes called ‘metasemantics’ (see Burgess and Sherman 2014) or ‘foundational theories of meaning’ (Speaks 2021). Seth Yalcin describes metasemantics as seeking ‘a certain sort of “more fundamental” characterization of [semantic] properties. It asks whether and how these semantic properties might admit of some illuminating reduction to, or unification with, nonsemantic properties’ (2014:1). Metasemantics looks for an explanation of semantic properties and relations in non-semantic terms. If this useful division of terminology had been sufficiently established in the 1980s, Fodor’s *Psychosemantics* (1987) might have had the clumsier title, *Psychometasemantics*. Perhaps it was just as well that this terminology was not around in those days.

Given this distinction, the familiar contrast between descriptive and causal theories of names (or of reference more generally) might, in some of its applications, involve a conflation. To say that a term refers ‘by description’ is to say something like this: its semantic content includes certain descriptive conditions that something has to meet in order to be the referent of the term; certain predicates have to be true of the thing. If a proper name refers by description, this means that specifying its semantic role — what kind of entity it is associated with the name and how the name contributes to the truth-conditions of sentences in which it occurs — involves specifying what these descriptive conditions are. This is straightforwardly a semantic project.

But these descriptive conditions are themselves about the application conditions for certain words — descriptive phrases — and explaining the meaning of names in terms of the meanings of other words still leaves something to be explained. This explanation is the metasemantic project. A causal metasemantic theory explains the semantic properties of words in terms of causal relations between those words and things in the environment. It is clearly consistent to hold that the meanings of names are descriptive and that they are explained in terms of something like causation (in fact, Lewis (1983) holds a view something like this). So to talk as if the important contrast is between descriptive and causal theories can involve a conflation of the semantic and metaemantic levels.

Our concern here, though, is with thought or intentionality, not reference. Those who think of intentionality in semantic terms should hold that thoughts, and parts of thoughts, can be assigned semantic values. If they also believe in the distinction between semantics and metasemantics, then they should hold that these assignments have to be explained or justified in non-semantic terms. Such an explanation might be in terms of intentional mental states (e.g. in Grice's programme). But if intentional states themselves are understood in semantic terms, then we also need a metasemantic account of intentionality.

Here I will not address the question of which metasemantic proposals should be pursued. Rather I am interested in why should semantic and metasemantic ideas apply to thought at all? Why think that thoughts (episodes of thinking) have semantically significant parts, or parts at all? Why think these parts contribute in a systematic way to the truth-values of whole thought episodes? Before embarking on the 'metasemantic' explanation of intentionality, these questions have to be answered — otherwise we would not have any idea what we are trying to explain in non-semantic terms. I am not saying that there are no good answers to these questions, only that they do need to be answered, and they are rarely addressed explicitly by those who assume a semantic conception of intentionality.

However, it is clear that if you already have a semantic conception of intentionality — i.e. if you think the only systematic way to explain intentionality is in semantic terms — and you believe that metasemantics is a worthwhile project, then you will have a motivation to make Putnam's assumption that what goes for language goes for thought, and you will find Putnam's argument about the ant compelling. Just as the meanings of words must be explained by a metasemantic theory, so must the intentionality of inner words or symbols or representations. If meaning — that which connects words to reality — must be given a general metaphysical explanation, then intentionality — that which connects thoughts to reality — must be given one too.

In other work, I have described the semantic conception of intentionality as 'anti-psychologistic' (Crane 2014). Psychologism in the relevant sense is the view that intentionality should not be understood in purely semantic terms: intentionality is not like reference; not all intentionality involves predication or the determination of truth-value; and not all mental representation need involve anything like a language (whether inner or outer). Psychologism about intentionality understands intentionality in terms of its psychological basis, and does not assume that this basis must be something which requires a semantic treatment in the sense just

described. I will say a little more about what a psychologistic approach to intentionality might involve in the final section of this paper.

5. Giving an account of intentionality

My aim in this paper so far has been to understand the motivation behind the question of aboutness. Answering the question of aboutness, as I am understanding it, requires giving an account of intentionality which is both wholly general (in the sense that it must apply to all kinds of mental representation), and reductive (in the sense that it must explain intentionality without using intentional or mental notions). I took Putnam's argument as representative, and I located its key assumption: that what goes for linguistic or pictorial representation goes for mental representation too. I then argued that physicalism is one reason for believing this assumption, and that a semantic conception of intentionality is another. These are both substantial views, so I conclude that unless other reasons are provided, the requirement to answer the question of aboutness when giving an account of intentionality is not a general philosophical requirement independent of more specific metaphysical views.

It might be objected that I am missing the point. Isn't the point just about the need to give an 'account' of intentionality, or an 'analysis' of it? Isn't the problem of intentionality just one of the central problems of philosophy, like the freedom of the will or the problem of induction? Don't all serious philosophers of mind have to give an account of intentionality? Isn't it just one of the things we have to account for? Isn't saying otherwise being left with something 'spooky' (to use Williams's word) or deeply 'puzzling' (to use Shea's), that is, with an unaccountable mystery?

This was certainly Putnam's view:

Some philosophers have ... what they take to be a proof that the mind is essentially non-physical in nature... Thoughts have the characteristic of intentionality — they can refer to something else; nothing physical has 'intentionality'... This is too quick; just postulating mysterious powers of mind solves nothing. (Putnam 1981: 2)

As we have seen, Putnam thinks that a causal theory of reference (or intentionality) is the only alternative to this mystery. Tim Button, who has made an extensive study of Putnam's various arguments, lists some of Putnam's descriptions of alternatives to a causal theory (under the

heading of a 'magical theory of reference'): a 'direct (and mysterious) grasp of Forms', 'medieval essentialism', 'noetic rays', 'divine intervention' or some other 'occult' or 'spooky' reference-fixing device (Button 2013: 30-31).

This strikes me as a somewhat tendentious way of putting the matter. Obviously, no sensible philosopher will want to postulate mysterious powers of the mind, or to endorse magic or noetic rays. Some sensible philosophers may invoke medieval theories of essence, and others may invoke divine intervention — philosophy is a broad church (so to speak). But surely the choice is not between a causal theory of intentionality on the one hand, and magic or theism on the other. Many naturalistically inclined philosophers will reject both theism and magic — surely they are not obliged merely on the basis of this to embrace a causal theory of intentionality. (For example: they may propose a teleological theory, like Shea.)

Suppose a naturalist — someone who thinks that the mind is part of the natural world, and can be studied as such — is not convinced by the philosophy behind physicalist arguments. Maybe they are not convinced by the causal closure arguments of Papineau (2001) and others, and maybe they are concerned about the category of the physical and what it really involves. Maybe they share Chomsky's scepticism (1990) about the possibility of framing a question of physicalism without 17th century metaphysics. Or maybe they just suspend judgement for broadly naturalistic or empirical reasons. What should they then say about the explanation of intentionality?

I don't myself think that 'naturalism' is a particularly precise term. However, there are some clear connotations of the term. For one thing, the contrast between a naturalistic and a transcendental philosophy has been of historical significance and still survives today and is worth discussing (see Smith and Sullivan 2011). Another thing that naturalists of all kinds will perhaps agree on is that there is no 'First Philosophy': we do not start *ab initio* with our pure metaphysical assumptions and then build things up from there, but we are allowed to appeal to what we already know about the world in forming our theories of parts of it. So, in particular, naturalists about the mind can appeal to what they already know about the mind.

Appealing to what we already know about mental representation is not the same as accepting that it is a mystery. And it does not necessarily involve giving necessary and sufficient conditions for the application of the relational predicate 'x represents y', which do not mention any intentional notions. We can appeal to what we know about different kinds of representation —

e.g. linguistic, pictorial, electronic, neural etc. — without having to say, in non-representational terms, what makes them all representations.

So how should such a philosopher start to give an account of intentionality? In the rest of this paper I will outline one way in which such an account might proceed.

We should start with a proper taxonomy. First there are the metaphysical categories which we need for understanding intentionality. We can distinguish intentional mental states, which involve a person having an intentional mental property over a period of time, from mental events and processes, which occupy time by having temporal parts. Beliefs, for example, are mental states, they involve a subject having a belief property over a period of time. Beliefs do not have temporal parts — there is no such thing as the earlier part of a belief and a later part. Episodes of imagining something, by contrast, are events or processes, entities with temporal parts — singing a tune to yourself in your imagination, for example, is something that has an earlier part and a later part.

Intentional states and events have a certain kind of common structure (from now on, I will talk just about states for simplicity, unless it is important to distinguish events and states). Every intentional state has an object, or more than one — this is the thing or things it is directed on in the world, whether or not that thing exists. The object, or intentional object, of a belief is what the belief is about. A belief may be about many things. It may also be about things that don't exist.

Intentional states fall into, or are classified into, types or kinds — John Searle (1983) called these *modes*, and I followed him in this (Crane 2001). Belief is a mode, so is visual imagination, so is vision, so is desire, and so on. This is my way of making Brentano's point that all mental phenomena involve direction upon an object, 'but they do not all do this in the same way' (1874). Brentano himself thought that there were three modes of direction upon an object — 'presentation', belief, and emotions like desire, love and hate. I have a different view of which modes there are — a crucial question I will address below.

In a given mode, each object is represented in a specific way, under what John Searle (1983) has called an aspect. The way in which the object is represented, in a given mode, is what I call the content of the state. So the planet Venus can be represented as the Morning Star, or as the Evening Star, or as the second planet from the Sun, or as Venus: these are all aspects under which Venus can be represented, and the differences between them are what I call (by stipulation) differences in 'content'.

Every intentional state and event, then, involves an intentional content — the aspectual representation of an intentional object — in an intentional mode. A belief that Austria is old-fashioned, for example, involves the representation of Austria (the object of the belief) as Austria, and a representation of it as old-fashioned. This is the content, and the belief is the mode.

I have presented ideas like these in other places (2001, 2008, 2014) so I will not repeat the details here. But even having given this brief summary of a taxonomy, it is natural to ask: ‘what makes all these states intentional or representational? What is this ‘intentionality’ that they all have in common? You’ve said what intentional modes and contents are, but what is intentionality itself?’ What I want to do now is to try and remove the temptation to keep asking these questions — or at least, asking them as a way of asking the question of aboutness.

My approach to this question is to start not with the question, ‘what is intentionality itself?’, but with the idea of the subject of intentional states. From a naturalistic point of view, the subject is the person (or animal) and a naturalistic approach to the person is to treat the person as a human being, a member of a certain species. There are huge philosophical debates around this (oddly controversial) idea, of course. But although I think the idea can be defended, I won’t do it here. Let’s accept the idea that the human being is the subject of intentional states, and therefore ask: how does intentionality fit into an overall understanding of the human being?

Human beings have many capacities: the capacity for digestion, the capacity for locomotion, for reproduction and so on. Some of these capacities are mental or psychological. Their psychological capacities are, in my view, characterised by their ‘direction upon an object’ — that’s what makes a capacity psychological (Crane 2008). With any capacity, we can distinguish between the capacity itself and its exercises: the capacity to walk and an individual episode of walking, the capacity to reproduce and an individual act of reproduction. We can also distinguish the exercise of a capacity from its results — a particular reproductive act, and the child that comes from it.

These distinctions apply to the psychological. The capacity for vision is one thing, the exercise of this capacity is another. When we are asleep, or with our eyes closed, we still have our capacity, but we are not exercising it. The result of the exercises of this capacity may be a belief — a persisting mental state. Exercises of a capacity can last a long time, and they are typically event- or process-like. Lasting a long time is not a sign that something is not an event; some events last longer than others, and there is no *a priori* limit on the length of an event. (For an

important recent application of the idea of a capacity to the case of perception, see Schellenberg 2018.)

Vision or perception is not the only case, though; the distinction between capacities and their exercises also applies to many other intentional phenomena. The capacity to reason, for example, can be exercised in acts of judgement, or in processes of thinking something through. The product of judgement is belief, and this can be the product of processes of thinking too. The capacity to speak a language is exercised in acts of speaking or understanding; changes to this capacity that come from new understanding can result in additions to the subject's stored knowledge of language (a result of the exercise of this capacity). The capacity to feel pain — to be able to sense the presence of noxious stimuli — is exercised in experiences of pain; the results of this capacity may be states of memories of pains, or dispositions to avoid painful stimuli. If there is such a thing as the capacity of the will, then decisions are the exercises of this capacity and intentions can be its results. And similarly for other mental capacities.

This approach does not take it for granted that each psychological concept in our everyday language picks out a genuine psychological kind, which has the capacity-exercise structure just mentioned. In some cases our commonsense psychology makes more distinctions than are plausibly psychologically real. One obvious case is the distinction we make between seeing things and seeing that something is the case. While there is an undeniable conceptual distinction here, this should not be taken to imply that these two concepts pick out distinct kinds of seeing or visual capacities. Distinctions in visual capacities — as for example hypothesised in the two visual systems of Milner and Goodale (2013) — should be made on the basis of empirical discoveries, not merely on the basis of commonsense psychological conceptual distinctions.

In other cases, commonsense does not make sufficient distinctions — we use the word 'memory' for the experience of recalling an event in which we participated ('episodic memory'), and also for the ability to recall facts (so-called 'semantic memory'). Is there a common element in these, apart from the factuality of what is being recalled? Is there a difference between semantic memory and belief? These questions are the subject of empirical research and cannot be answered by reflection on commonsense psychology alone.

It should perhaps go without saying that on a naturalistic, psychologistic approach, the question of which mental capacities there are is an empirical one. But my point here is not just this obvious one, but to draw attention to how this approach to the metaphysics of mind differs

from many standard approaches, and how it more closely harmonises with the approach taken in psychology and cognitive science than many standard approaches. Standard approaches in the philosophy of mind tend to start with the idea of a mental state or a mental event or a mental property, and then asks how these mental states/events/properties are related to physical states/events/properties. My investigation starts a step further back, with an account of the mental in terms of the mental (or intentional) capacities of organisms. The notion of a mental state or event is then understood in terms of the exercise or result of a capacity. The next step then is to explain how an organism can have such a capacity — what the basis of this capacity is in the brain and body of the organism.

Since the starting point is with the organism, not with fundamental physical particles and fields, this approach to the mind-body question will not necessarily take a stand on physicalism, conceived of as a global thesis about the nature of reality as a whole. The relationship of physics to mental reality is one thing; the relationship of an organism's body to its mental capacities is another. Starting with the mental capacities of an organism brings the philosophy of mind closer to the science of the mind than the standard 'mental—physical' starting point. In listing the capacities I just briefly referred to — perception, memory, language, reasoning, sensation — I could be listing the chapter headings of a textbook in cognitive psychology.

Let me return, finally, to the place of the analysis of intentionality in this way of thinking about the mind. Intentionality is what characterises all the mental capacities as mental; and I have claimed that it should be understood rather abstractly in terms of the ideas of object, content and mode (or attitude). But the search for a general explanation of intentionality should stop there. The intentionality of each capacity — that is, the nature of its objects and contents, how it represents its objects — will be given its own specific explanation. The intentionality of vision — what its objects are, how they are represented, what its representational mechanisms are — needs a different explanation from the intentionality of language, reasoning and so on. Once we have given these explanations, then I claim there should be no need to ask any further questions of the form, 'but what makes it the case that any representation is *about* its objects?' in the style of Williams (2020). The explanation of aboutness is undertaken on a capacity-by-capacity basis; and these explanations are allowed to appeal to notions like representation, object, content and so on. What needs empirical explanation are the psychological capacities; but intentionality itself is not a psychological capacity.

I will finish with two analogies from recent philosophical literature, which may help in making this point clear. The first is from the metaphysics of parts and wholes. Peter van Inwagen (1990) argued that any metaphysics of material objects must answer what he called the ‘special composition question’: how do the parts of an object come together to make a whole? I am sceptical that we need to answer this question for objects in general. The answer will be different for different kinds of objects: how the parts of a car come together to make a car requires different account from how the parts of the human body come together to make a body. But once you have told the story, for a specific kind of thing, of how the parts of things of that kind hang together — e.g. for a body: ‘the hip bone’s connected to the thigh bone etc.’ — then there is no further question about composition to answer.

The second analogy is with Timothy Williamson’s view of knowledge (2000). On this view, knowledge cannot be reduced or defined, but we can say general informative things about it (that it entails truth, for example). For Williamson, knowledge is a general determinable kind — the ‘most general factive mental state’ — which has various determinates: seeing, remembering etc. These are the ‘ways of knowing’, and an account of knowledge is given by giving accounts of these different ways of knowing. Once these accounts have been given, then there is no further requirement to give an account of knowledge as such.

In both cases, there is a general phenomenon — parthood, knowledge — about which we can say general things of an abstract nature, but which we primarily understand in terms of their particular kinds of manifestations (the parts of a thing of a certain kind, a particular way of knowing). Similarly, I am claiming, with intentionality. We can say things of an abstract nature about intentionality, but when it comes to explaining how any intentional phenomenon gets its objects, the answer will have to be given for each intentional capacity in its own terms. We will therefore get specific explanations of the intentionality of perception, of motivation, of reason, of language and so on. But not of intentionality as such.

6. Conclusions

I have argued that it is not obvious that the explanation of intentionality requires that we give an answer to the question of aboutness. One famous argument for this requirement relies on controversial and implausible assumptions. This assumption can be made by physicalism and antipsychologism; but these views are not obligatory.

In particular, naturalism — understood as a view distinct from, though compatible with, physicalism — does not require an answer to the question of aboutness. Sometimes it looks as if there are only two choices: giving an account of intentionality by answering the question of aboutness in general, reductive terms; and proudly declaring oneself to be a non-reductionist who is not concerned with these scientific questions. And it can seem as if the first approach is the only respectable approach for a naturalist to take. In this paper I have questioned this false dichotomy. The first approach is not obligatory for naturalists, and the second involves putting one's head in the sand (or perhaps leaving it in the clouds). I conclude that it is possible to explain intentionality naturalistically without being required to answer the question of aboutness.

References

- Blom, Jan Dirk (2010) *A Dictionary of Hallucinations* (London: Springer)
- Brentano, Franz (1874) *Psychology from an Empirical Standpoint* (London: Routledge 1993)
- Burgess, Alexis and Brett Sherman (2014) *Metasemantics New Essays on the Foundations of Meaning* (Oxford: Oxford University Press)
- Crane, Tim (2001) *Elements of Mind* (Oxford: Oxford University Press)
- Crane, Tim (2009), 'Intentionalism' in *Oxford Handbook to the Philosophy of Mind* edited by Ansgar Beckermann, Brian McLaughlin and Sven Walter (Oxford: Oxford University Press) 474-93
- Crane, Tim (2014) *Aspects of Psychologism* (Cambridge, MA: Harvard University Press)
- Crane, Tim (2016) *The Mechanical Mind*, 3rd edition (London: Routledge)
- Dickie, Imogen (2015) *Fixing Reference* (Oxford: Oxford University Press)
- Field, Hartry (1978) 'Mental Representation' *Erkenntnis* 13: 9–61
- Fodor, Jerry (1987) *Psychosemantics* (Cambridge, MA: MIT Press)
- Jackson, Frank (1998) *From Metaphysics to Ethics* (Oxford: Oxford University Press)
- Lewis, David (1983) 'New Work for Theory of Universals' *Australasian Journal of Philosophy*
- Milner, David and Melvyn Goodale (2013) *Sight Unseen* (Oxford University Press)
- Chomsky, Noam (1990) 'Language and Nature', *Mind*
- Papineau, David (2001) 'The Rise of Physicalism' in Carl Gillett and Barry Loewer (eds.) *Physicalism and Its Discontents* (Cambridge: Cambridge University Press) 3-36
- Putnam, Hilary (1981) *Reason, Truth and History* (Cambridge: Cambridge University Press)

Howard Robinson (1993) 'The Anti-Materialist Strategy and the "Knowledge Argument' in Robinson (ed.) *Arguments Against Materialism* (Oxford: Oxford University Press)

Schellenberg, Susanna (2018) *The Unity of Perception* (Oxford University Press)

Searle, John R. (1983) *Intentionality* (Cambridge: Cambridge University Press).

Shea, Nicholas (2018) *Representation in Cognitive Science*, (Oxford: Oxford University Press)

Smith, Joel and Peter Sullivan (eds.) (2011) *Transcendental Philosophy and Naturalism* (Oxford: Oxford University Press)

Speaks, Jeff (2021) 'Theories of Meaning' SEP Speaks, Jeff, "Theories of Meaning", *The Stanford Encyclopedia of Philosophy* (Spring 2021 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2021/entries/meaning/>>

Williams, Robert G. (2020) *The Metaphysics of Representation* (Oxford: Oxford University Press)

Williamson, Timothy (2000) *Knowledge and its Limits* (Oxford: Oxford University Press)

Van InWagen, Peter (1990) *Material Beings* (Ithaca: Cornell UP)

Yalcin, Seth (2014) 'Semantics and Metasemantics in the Context of Generative Grammar' in Burgess and Sherman (eds) (2014), 1-46