

# New Essays on Singular Thought

Robin Jeshion

Print publication date: 2010

Print ISBN-13: 9780199567881

Published to Oxford Scholarship Online: Sep-10

DOI: 10.1093/acprof:oso/9780199567881.001.0001

We are Acquainted with Ordinary Things\*

Imogen Dickie

DOI: 10.1093/acprof:oso/9780199567881.003.0008

## Abstract and Keywords

To be 'acquainted' with a thing is to be in a position to think about it in virtue of a perceptual link, and without the use of any conceptual or descriptive way of identifying it. There are old arguments for the claim that we cannot be acquainted with ordinary material things. This chapter uses recent empirical results about perception to show how these arguments can be overturned. The chapter has three parts. The first summarizes the relevant results, the second uses these results to construct an account of acquaintance with ordinary things, and the third turns this account against objections to such proposals from claims about the possibility of perception-based error.

Keywords: acquaintance, demonstrative reference, attention, perception

Let us define 'acquaintance' as follows:

Definition: A subject, *S*, is 'acquainted' with an object, *o*, iff *S* is in a position to think about *o* in virtue of a perceptual link with *o* and without the use of any conceptual or descriptive intermediary.

There is an old and intuitive argument for the conclusion that some of our thoughts must be, as I shall say, 'acquaintance-based thoughts': thoughts we are in positions to have because we are acquainted with the objects they are about. The argument goes roughly like this. All other thoughts about particulars depend for their representational content on acquaintance-based

thoughts. So there must be acquaintance-based thoughts. For otherwise none of our thoughts would be about anything at all.

But there are also old arguments for the conclusion that we cannot be acquainted with ordinary middle-sized objects: we can be acquainted only with our own sense data, or some other range of 'simple' things. So, given these two lines of argument, we have

1 Acquaintance provides the basic subject matter of thought.

and

2 We can be acquainted only with our own sense data.

(p.214) But **1** and **2** entail the undesirable

3 Our own sense data are what we are really thinking about.

I think it is fair to say that the mainstream response to this situation has been to concede **2**, and avoid **3** by rejecting **1**. My aim in this chapter is to show how recent empirical results about perception can be used to provide an account of acquaintance with ordinary objects. Given this account, we will be able to reject **2**, so keep **1** without getting stuck with **3**.

The relevant empirical results are results in psychology which suggest that a traditional view of the kind of information we get from perception is wrong. According to this traditional view, which I shall call the 'Old Empiricist View' (hereafter 'OEV'), visual experience presents us with an array of features laid out around us and developing over time.<sup>1</sup> As long as the OEV is assumed, the claim that we can be acquainted only with our own sense data, or with some other range of 'simple' things, is unavoidable. For to assume the OEV is to assume that conceptually unaided perception delivers only a shifting mosaic of features, which you will call 'colour (or texture, or shape) patches' or 'sense data', depending on whether you are prepared to allow that they exist independently of our experience of them.

But the OEV has only ever been an assumption—an assumption that the empiricist tradition has supposed would be obvious to each of us in our own case. And more recent and more systematic empirical testing has suggested that it is a false assumption. There is strong evidence that conceptually unaided perceptual processing does not deliver an array of features. It delivers a world already divided into objects (in senses of 'divided' and 'object' to be explained).

The chapter has three parts. §1 summarizes the empirical evidence against the OEV and in favour of the view that our pre-conceptual processing delivers a world already divided (in some sense of 'divided') into objects (p.215) (in some sense of 'object'). §2 uses the new view of what pre-conceptual processing delivers to build an account of acquaintance-based thought about ordinary objects. §3 uses this account to overturn two arguments which have combined with the OEV to form the traditional case against the claim that we can have acquaintance-based thoughts about ordinary objects—the arguments from analogues for the case of singular thought of Frege's Puzzle, and of the problem of empty names.

I should say at the outset that I am not aiming for empirical comprehensiveness. I shall summarize only some of the evidence that points away from the OEV and towards the claim that perception delivers scenes already divided into objects. I shall not attempt an exhaustive survey of this evidence. Nor shall I consider whatever counter-evidence there might be in the psychological literature. So the spirit of my discussion is as follows. There is an old, hard, and as yet unclosed question about whether it is possible to uphold an intuitive picture of how our thought gets its subject matter while avoiding an unpalatable account of what this subject matter is. The OEV forces a negative answer to this question. But some recent empirical findings suggest an alternative to the OEV. Whether this new, alternative, view of what experience delivers is right is itself an open question. The empirical results that support it might be overturned. Or they might be overwhelmed by counterevidence. But if the OEV is false, the history of the empiricist attempt to give an account of how our thought about the world works is a history of battles fought on terrain defined by a false empirical assumption. And it is at least worthwhile to see what happens to the old question if we discard the OEV and suppose the new account of what experience delivers in its place.

## §1. The Empirical Evidence Against the Old Empiricist View of What Experience Delivers

This section summarizes some of the empirical evidence against the OEV and in favour of an alternative view of what our pre-conceptual processing delivers. According to this alternative view, our pre-conceptual processing does not deliver an array of features. It delivers a world already divided (in some sense of 'divided') into objects (in some sense of 'object').<sup>2</sup>

**(p.216)** The evidence that I am going to present is all evidence for the claim that we sometimes attend to objects (again, in some sense of ‘object’) rather than just to regions in the visual field. So let me start by saying why evidence about what we are attending to will support a conclusion about what is delivered by our pre-conceptual processing.

Consider your visual field as it is laid out in front of you when you look at this page. Your visual field contains much more information than you can access for conceptual thought. For example, you can see a large number of words on the page. But at each moment you will be able to tell what only a few of them are. ‘Selective attention’ is the process by which a specific part of the visual field is highlighted as a part accessible to conceptual thought and verbal reporting. Given this definition of ‘attention’, it follows that pre-attentive processing is ‘encapsulated’ from conceptual thought:<sup>3</sup> if what you are attending to determines what is available for conceptual thought, the processes involved in preparing the field over which attention roams are pre-conceptual. So if attention is sometimes directed to objects (which is to say, if the processes involved in preparing the field over which attention roams deliver up objects) then experience is delivering a world already divided into objects prior to whatever further dividing and packaging might happen when we conceptualize it.

Here are three pieces of evidence for the claim that attention is sometimes directed to objects.

### §1.1. First Piece of Evidence—the Automatic Spread of Attention

Figures 7.1 and 7.2




Figure 7.1




Figure 7.2

(Both from Scholl 2001; 'C' stands for 'cue'; 'S' for 'same rectangle'; 'D' for 'different rectangle')

**(p.217)** In a basic 'automatic spread of attention' experiment, people are presented with either two (fig. 7.1) or three (fig. 7.2) rectangles arranged as shown. A spotlight is used to make a spot where the 'C' is; the spot disappears; immediately afterwards there is a 'luminance decrement' (a gradual darkening) either at the point marked 'S' or at the point marked 'D' (so either on the same rectangle as the initial cue or on a different rectangle). It turns out that we are much faster at detecting the decrement when it occurs on the same rectangle as the initial cue than when it occurs on a different rectangle. The effect cannot be caused by a difference in distance from the initial cue, because the 'D' and the 'S' are equidistant from the 'C'. Rather, it seems that attention is attracted by the initial cue, then spreads out to the boundaries of the shape the cue is on. And attention is not just spreading until it hits a solid line on all sides. For the effect remains in fig. 7.2 cases, where the rectangles where the action is taking place are partially occluded. Rather, the boundaries to which attention is spreading seem to be apparent 'object' boundaries in at least the following sense: they are boundaries around a thing which is seen as lying behind, and partially occluded by, something else.<sup>4</sup>

## §1.2. Second Piece of Evidence: Amodal Completion



Figure 7.3.

(From Driver et al. 2001)

Fig. 7.3 seems to contain a square and a complete circle, with the circle lying behind the square, and so partially occluded by it. So when you look at fig. 7.3 your visual system is completing the notched circle into an **(p.218)** apparent object (a circular disc) which you see as lying behind another object (something square). Psychologists call this process of filling in the missing part of an apparent object 'amodal completion'. Now consider fig. 7.4 and fig. 7.5. And suppose that in each case you are asked to find the notched circle.



Figure 7.4.

(From Driver et al. 2001)



Figure 7.5.

(From Driver et al. 2001)

The notched circle is much harder to find in fig. 7.5 than in fig. 7.4. And it has been found that this gap between fig. 7.5 type cases and fig. 7.4 type cases increases with the number of distracters in the diagram. The more extra shapes are put into the frame, the harder it is to find the target shape (**p.219**) in a fig. 7.5 type case. In fig. 7.4 type cases the frame has to get really big (so that you cannot take it in at a glance) or the shapes have to get really small before the notched circle ceases to jump out at you from among the shapes surrounding it.

Psychologists have concluded that the reason the notched circle is so much harder to find in fig. 7.5 is that amodal completion is a 'pre-attentive' process. When you look at fig. 7.4, your pre-attentive processing delivers up a field which contains a notched circle. This notched circle can then grab your attentional spotlight as it goes by. In contrast, when you look at fig. 7.5, your pre-attentive processing fills in the missing part of the notched circle. So the field over which your attention roams does not contain a notched circle. It contains a partially occluded full circle instead. To find the notched circle you have to attend to each shape in turn, looking to see if it contains a 'hidden' notched circle. That is why the target shape is so much harder to find in fig. 7.5 than in fig. 7.4, and why the gap increases with increasing numbers of distracters.<sup>5</sup>

But amodal completion is completion into objects, in some sense of 'object': the completed circle is seen as a bounded figure lying behind the occluding corner of the square. So if attention roams over a field in which amodal completion has already taken place, our pre-conceptual processing is delivering a world already divided into objects of some kind.

### §1.3. Third Piece of Evidence: Multiple Object Tracking

The final piece of empirical evidence I shall present concerns the relative efficiency of attentive and pre-attentive processing. It is generally accepted that pre-attentive processing is much more efficient than post-attentive processing.<sup>6</sup> Sub-personal, pre-attentive, information processing is fast and has enormous capacities. In contrast, our dealings with information after attention has selected it as available for conceptual thought are slow and have relatively small capacities. The more conceptual level tasks you have to perform at once, the slower you will get. But parsing the visual field into objects seems to share the efficiency of pre-attentive processing. So parsing into objects seems to be pre-attentive.

**(p.220)** The main experiments which have been taken to establish this result involve ‘multiple object tracking’. In multiple object tracking experiments subjects are shown a display containing a number of identical dots. Some of the dots are ‘flashed’ to distinguish them from others (a). The dots are then shown moving randomly around the display (b). When the motion stops, the subject is asked to say whether a given dot is one of those which flashed at the outset (c).




Figure 7.6.

(From Scholl 2001)

If parsing into apparent objects were post-attentive, you would expect us to be much better at performing this task when asked to track only one dot (so when only one dot flashes at (a)) than when asked to track two, and worse at tracking three than tracking two, and worse again at tracking four. But it turns out that our success at this task does not fall off in this way. We are as good at keeping track of two, three, or four things as we are at keeping track of one. At five the capacity seems to hit overload and beyond five it falls off suddenly. And this pattern of falling off— uniform success up to some point at which the system's capacity is reached—is the pattern characteristic of pre-attentive, not post-attentive, processing. So it seems that pre-attentive processing is already serving up a scene parsed into objects that we then attend to.<sup>7</sup>



Further experiments along these lines have even shown that there is a link between the apparent objects served up by pre-attentive processing and the ordinary material objects that we would like to be able to say we are thinking about.<sup>8</sup> For it has been found that the capacity to keep track of dots (p.221) in multiple object tracking depends on the way the dots behave. We are able to track dots which trace out continuous spatio-temporal paths. We are not able to track dots which cover the same amount of ground in the same time but by jumpy (non-continuous) motion. Nor can we track through motion that does not preserve the locations of a thing's parts relative to one another (where the dot seems to disassemble and project itself to the new location). Again, we can track through apparent occlusion—where the dot seems to disappear behind a barrier then reappear again. But we cannot track through sudden disappearance and reappearance. So it seems that we can track through motion and change which are in keeping with the basic characteristics of the category of ordinary material objects (spatio-temporal continuity, moving and acting as a whole) but not through motion and change which violate these characteristics. The 'objects' at which our attention is directed are, in some yet to be explained way, closely associated with the ordinary material things that we would like to be able to say are the basic subject matter of our thought.

#### §1.4. The New Psychologists' Findings and the Old Philosophers' Problem

The empirical findings I have summarized all support the conclusion that our attention is sometimes, as psychologists say, 'object-based': our pre-conceptual processing divides the visual field into units (psychologists say 'visual objects') which can attract and hold our attention. Note that the claim is only that attention is sometimes object-based, not that it always is. And it seems clear that sometimes attention is location-based instead: I can attend to one, rather than another, region of a uniformly coloured wall. But for convenience in what follows I shall set cases of location-based attending aside, and say that when your attention is locked to a visual object you are in an 'attentional state'.

There is an obvious apparent obstacle to using the empirical results I have set out as the basis for an account of acquaintance with ordinary things. For psychologists' 'visual objects' are not philosophers' ordinary objects. The perceptual system treats all of the following as visual objects: dots on screens; vertices of geometrical shapes; shadows; parts of physical objects; ripples on water; patches of reflected light. All of these things can grab and

hold the attentional spotlight. So anyone who wants to build an account of acquaintance-based thought around the empirical results about attention must deal with a gap between the class of potential objects of **(p.222)** attention (which includes parts of ordinary objects, patches of light, and so on) and the class of objects that most of our thoughts about the external world concern (ordinary material things).

In the next section I show how I think this obstacle can be overcome.

## §2. Acquaintance as Attentional Link

### §2.1. The Modal Containment Principle

I shall assume the following widely held background view of how a system of beliefs about particular things is structured,<sup>9</sup> and the role that singular terms (for example, proper names and demonstratives) play in such a system. A system of beliefs is structured like a filing system. Each file is a cluster of information that the speaker takes to be information about a single thing. Taking the information in the file to be about a single thing consists in attempting (within limits imposed by our laziness and lack of logical omniscience) to keep the file free from contradiction, and to keep the file's contents consistent with the general beliefs about particulars that the system contains. So the subject attempts to keep the file free from 'is F' and 'is not F' pairs. And the subject will, in general, let 'is F' into a file only if willing to admit 'Something is F' as a general belief about particulars. A singular term acts as a label on a file in that it plays a label-like role in determining how information entering the system is processed. For example, the term determines where in the system incoming information expressed using it will be sent: all the information I receive in the form of sentences containing the name 'George W. Bush' is sent to my 'George W. Bush' file. And when I discover that two files are actually about the same thing I collapse them together, and may use the term associated with either to label the new file formed as a result. A singular term's referent (if it has one) is the object (if there is one) that the file of beliefs which the term labels is about.

Against the background of this view, questions about the kinds of relations to objects that enable us to think about them become questions about the relations between files of beliefs and the objects they are about. **(p.223)** The account of acquaintance-based thought I am going to propose is built around a principle which I want to suggest any answer to this kind of question must respect. The principle is a theoretical counterpart of an intuitive claim about

the nature of representation: the claim that if an entity (for example, a singular term or a thought component) stands for an object then the way the representing entity can be deployed depends on the ways it is possible for the represented object to behave.<sup>10</sup>

This intuition about representation is most clear-cut in toy cases. For example, suppose that a teacher is showing a class how to use a molecular modelling kit. The kit contains little plastic balls of various colours, and sticks which can be used to join balls together. The teacher's instructions to the class will go something like this. 'We know that oxygen has valency two and carbon has valency four. So an oxygen atom can form two bonds with other atoms, and a carbon atom four. Now use the black balls to represent carbon atoms and the red ones to represent oxygen atoms, and make some models of molecules.' The instructions to use the black balls as representatives of carbon atoms and the red ones as representatives of oxygen atoms could be expanded upon as follows: 'Combine black and red balls only in ways which match the ways you have been told that carbon and oxygen atoms can combine.' A student who combines the balls in ways the atoms cannot combine, building 'impossible' molecules, is not following the teacher's instructions, and is not using the balls as 'representatives' of the atoms in the sense of 'representation' that use of a molecular modelling kit involves.

So in this toy case, a representing relation (*this* kind of ball stands for *this* kind of atom) is established by a stipulation about how the representing entity should be deployed: the stipulation that the representing entity should be deployed only in ways which match the ways it is possible for the represented entity to behave.

To bring out how this intuitive model of representation might apply outside this kind of toy example it is necessary to add a little detail to the (p.224) skeleton 'files' view of how a system of thought about particular things is structured.

Consider the way a file of beliefs develops over time. At  $t_0$  the file has a given informational content. For convenience, let us identify this content with the set  $\{I_{01}, I_{02}, I_{03}, \dots\}$ , where each  $I_{0i}$  is the informational content of a belief in the file at  $t_0$ . At  $t_1$ , the file has content  $\{I_{11}, I_{12}, I_{13}, \dots\}$ ; at  $t_2$  it has content  $\{I_{21}, I_{22}, I_{23}, \dots\}$ ; and so on. So the file consists in a succession of sets of items of information.<sup>11</sup>

The differences between successive sets (so between the file's informational content at successive times) depend on two factors. The first factor is which information is delivered to the file (for example, how my 'George W. Bush' file develops depends partly on what information expressed using this name I am receiving). The second factor is which operations on the information delivered to the file and the information already there the subject performs. These operations divide into two kinds. On the one hand, there are operations which form part of the subject's rational strategy for maintaining the file. On the other, there are operations which are not rationally driven. The deletion of information through forgetfulness is one variety of non-rational operation. Another is the unintended distortion or garbling of information, like the distortion which accidentally replaces my belief that a colleague is from New Brunswick with the belief that the colleague is from Nova Scotia.

I am going to suggest that it is in terms of the first, rationally driven, kind of operation that the notion of deploying a representation in a way which matches the ways the represented object can behave should be understood outside the toy molecular model case. So I shall take a little time to say how these rationally driven operations work.

Let us say that, for each file, the cluster of dispositions that determines which rationally driven operations the subject will perform in maintaining the file is the file's 'Governing Conception'. The minimum requirement on a Governing Conception is that it tend towards keeping the file consistent. This is a requirement on Governing Conceptions because **(p.225)** objects cannot have inconsistent properties: if a cluster of information is allowed to develop in a way that is indifferent to its consistency it is a mere list, rather than a file of information which is potentially about a particular thing. We are not logically omniscient. So we might not succeed in keeping our files of beliefs consistent. But no cluster of information constitutes a file of information potentially about a particular thing if the subject's procedure for incorporating new information into the file is indifferent to whether what is being added is consistent with what is already there.

We can maintain files using only minimal Governing Conceptions. For example, suppose that a subject, S, is told to treat an expression, 'a', as a name, told not to worry about what kind of thing it is a name of, and told to treat a story expressed using 'a' as a reliable source of beliefs about a to the extent that this is compatible with the assumption that 'a' is a name for a particular. Conforming to this instruction, S opens an 'a' file, and accepts,

rejects, or withholds assent from sentences contained in the story in such a way as to avoid the inclusion of contradictory 'is F' and 'is not F' pairs in the file.

But most of our Governing Conceptions are far richer than this. S in the example is not in a position to decide what counts as a legitimate addition with respect to beliefs about what *a* is like at different times. If the file contains the information 'is F at *t*', S can conclude that the file cannot also contain 'is not F at *t*'. But S has no grounds for a decision about whether it could contain 'is not F at *t* + 1'. Similarly, S has no grounds for a decision about pairs of predicates which apply, or could apply, to different parts of a thing at a time: if 'is F at position *p*' is included in the file, S can conclude only that 'is not F at *p*' cannot be included as well; there are no grounds for a decision about 'is not F at *p*\*'. In contrast, for most of the files in our systems of beliefs about particulars, whether the file contains 'is F at *t*' is relevant to whether we will allow 'is not F at *t* + 1' as an addition. Similarly, for many of the files in our systems of beliefs about particulars, whether the file contains 'is F at *p* (at *t*)' is relevant to whether we will happily add 'is not F at *p*\* (at *t*)'. And most of our files are subject to internally driven up-dating: we draw conclusions about what an object which had such and such properties at *t* will be like at *t*\*, and update our files of beliefs accordingly. A file's Governing Conception determines how these updatings and adjustings proceed.

**(p.226)** Let us call the possible ways a file might evolve given its Governing Conception the 'templates' generated by the Governing Conception. So a template lays down a pattern of evolution of properties. And let us say that an object 'fills' a template iff its properties evolve in a way which matches this pattern. Then the condition for a file of beliefs to be about an object that I want to propose can be put like this:

*The Modal Containment Principle* (hereafter 'MCP')—a file of beliefs is 'about' an object only if the templates generated by the file's Governing Conception are templates it is possible for things of the object's category to fill.

The sense of 'possibility' relevant to the MCP is what I shall call 'categorical possibility'. This is possibility relative to a thing's category. A 'category' (in the sense I am concerned with) is a set of things whose members share a kind of life history. 'Ordinary physical object', 'collection of ordinary physical objects', 'place', 'event', and 'natural number' are all category terms. I shall say a little about how our account of category boundaries might be refined in §2.3. But for now let us suppose the following traditional account

of the category of ordinary objects. Ordinary objects are located in space and time. They are causally unified at a time: what one part of an object is like at a time is causally interdependent with what its other parts are like at the time. They are causally unified over time: what an object is like at  $t$  depends causally on what it was like at  $t - 1$ . They trace out continuous spatio-temporal paths. And the relations between their parts are stable: in general, for ordinary object  $o$ , parts of  $o$  which are joined on to one another at  $t$  are still joined on to one another at  $t - 1$ . So for the case of thought about ordinary objects, the MCP says that a file of beliefs is about an ordinary object only if the templates (possible life histories) that the file's Governing Conception generates are templates that ordinary objects could fill: the Governing Conception must rule out possible updatings of the file which would violate spatio-temporal continuity, or causal unity, or the stability of relations between parts.

Note that the MCP is a containment principle, not an exact match principle. The claim is that if I am thinking about  $o$  the patterns in evolution of properties I treat as possible for  $o$  must be patterns to which it is possible for things of  $o$ 's category to conform. There is no suggestion that I must treat as possible all the patterns of evolution in properties that in fact are possible for things of  $o$ 's category: the objects **(p.227)** we think about might have properties and potentialities of which we have no ken. (The parallel claim holds in the molecular model case. The instructions for using the molecular modelling kit stipulate that balls are to be combined only in ways that match the combinatorial possibilities of the corresponding atoms. But an atom's possible behaviour is not fully characterized by its combinatorial possibilities. Elements whose atoms can form the same number of bonds, and even different isotopes of the same element, differ in their possible behaviours—for example, in their stability, and in how stable the compounds they form are. So the stipulation that the balls be deployed in ways that match the atoms' possible behaviours is treating only some of these possible behaviours as representationally relevant.)<sup>12</sup>

Note also that the MCP is only a necessity claim.<sup>13</sup> The suggestion is that if my uses of  $v$  are to refer to  $o$  my  $v$  file must at least (but not at most) have a Governing Conception which keeps faith with  $o$ 's kind. Given that we can assign different names to different members of the same category, there must be additional conditions on reference (conditions which will attach  $v$  to  $o$ , rather than to any other member of  $o$ 's category).

For example, according to one widespread account, the file of beliefs I would express using proper name *v* is about *o* (so *o* is the bearer of my uses of *v*) iff *o* is the 'dominant causal source' of the beliefs in the file. An object, *o*, is the 'causal source' of the belief *S* would express by saying [*v* is *F*] iff *S* was caused to have the belief by standing in an appropriate relation to *o* (for example, by reading about *o* in a newspaper report written by a reliable source who uses *v* as a name for *o*; taking at face value the testimony of someone whose beliefs expressed using *v* are causally derived from *o*, and so on). The 'dominant' causal source of the beliefs in a file (if there is one) is the object (if there is one) which is the causal source of most of the beliefs in the file.<sup>14</sup> If the MCP is true, this account at least needs amendment: *o* is the bearer of my uses of *v* iff *o* is the dominant causal source of the beliefs in my '*v*' file and my '*v*' file's Governing Conception generates templates appropriate to *o*'s category.

**(p.228)** Similarly, if I am using a demonstrative as a label on a file of beliefs formed by taking the information delivered through an attentional channel at face value, if the demonstrative refers to anything it refers to the thing I am attending to. For in this case which beliefs the file contains depends on how things stand with the thing at the end of the attentional channel. But if the MCP is true, a file of beliefs formed by taking the information delivered through an attentional channel at face value is about the object at the end of the channel only if the file's Governing Conception is appropriate to the object's category.<sup>15</sup>

But why should the MCP be accepted? The MCP is a transposition of an element of the account of representation that seems right for the molecular model case to discussion of our ordinary thought and speech about particular things. But what reason is there to think that it is a legitimate transposition? Why think that the match between deployment of representation and nature of thing represented that seems constitutive of representation in the toy case applies to ordinary thought and speech about particular things too?

One reason to think this is just that the MCP combines with the empirical story about attention to yield an account of acquaintance-based thought about ordinary things which saves the intuitive picture of how our thought about the external world works—the picture on which experience provides the subject matter of thought. The second part of this section develops this account.

But if the MCP is true we should expect to come across more direct evidence for it in the form of ordinary speaker intuitions about the necessary and

sufficient conditions for reference and aboutness. And there is some evidence of this kind.

For example, consider the case of somebody with synaesthesia, for whom experience of a numeral, or even thought about a number, is accompanied by a visual experience of a colour.<sup>16</sup> Somebody might have this condition but manage to fence off the association with colours from their beliefs about numbers (so that, though thought about the number seven is accompanied by a visual experience of purple, 'is purple' does not end up in the subject's 'seven' file). But usually synaesthetes whose experiences of numerals or **(p.229)** thoughts about numbers are accompanied by colour experiences do believe that numbers are coloured. Ordinary speaker intuition cries out that a synaesthete who says (and means literally) things like 'Seven is purple' can still be using 'seven' as a name for the number seven. So our ordinary notions of reference and aboutness do allow some slack between deployment of a representation and the nature of what is represented. But ordinary speaker intuition also cries out that the synaesthete's claim to be using 'seven' to refer to the number rests on the fact that the belief that seven is purple is, in some sense, 'peripheral' to the file. The required sense of 'peripheral' cannot be captured by saying that categorially deviant beliefs must be outnumbered in the file by categorially legitimate ones. For reference is not undermined even if the purpleness is allowed to invade every belief in the file (so that the file contains 'is a purple prime number'; 'is, purply, the square root of forty nine'; 'is the purple number of things on this table'; and so on, with no colour-free residue). Nor can the required sense of 'peripheral' be captured by saying that the subject must regard purpleness as an unimportant or contingent attribute of the number. Someone's interest in numbers might be primarily aesthetic. And a synaesthete might maintain that numbers are essentially coloured and that each number's colour is essential to it. Rather, the sense of 'peripheral' in which ordinary speaker intuition rules that 'is purple' must be peripheral to a file of beliefs about a number seems to be the sense entailed by the MCP. Ordinary speaker intuition rules that my 'seven' file is not about the number if my beliefs about seven's colour change over time and I believe that a number's arithmetical properties change when its colour does, or if I believe that all and only groups of purple things have seven members. So ordinary speaker intuition seems to be ruling that my 'seven' file may contain 'is purple' as long as the categorially deviant belief is insulated from the file's Governing Conception.



Again, consider this description of the White Spider (an ice-field on the North Face of the Eiger):

The Spider on the Eiger's face is white. Its body consists of ice and eternal snow. Its legs and its predatory arms, all hundreds of feet long, are white too. From that perpetual, fearfully steep field of frozen snow nothing but ice emerges to fill gullies, cracks and crevices. Up and down. To left, to right. In every direction, at every angle of steepness.

And there the Spider waits. (Harrer 1976: 11)

**(p.230)** If we take what is said here at face value, the author of this description is attributing to the ice-field some properties which are appropriate to its kind (a specific location and constitution) and some which are not (the attributes of a patient predator). Ordinary speaker intuition cries out that a subject who says things like 'The Spider waits' and 'The Spider is predatory' (and means them literally) can still be using 'The Spider' as a name for an ice-field. So, again, the example shows that our ordinary notions of reference and aboutness do allow some slack between deployment of a representation and the nature of what is represented. But ordinary speaker intuition also cries out that reference and aboutness will persist only if this kind of mistake is, in some sense, 'peripheral' to the file. And, again, it is plausible that the required sense of 'peripheral' is the sense entailed by the MCP. I can believe that an ice-field has character traits and propositional attitudes. These beliefs can be quite invasive without undermining aboutness: I might think that everybody who has come to grief on the ice-field has been clumsy or unlucky enough to wake it up; that I will get by if I can catch it sleeping; and so on. But my claim to be thinking about the ice-field is undermined if I start believing things like 'The Spider knows where I live and will be hiding in the cupboard when I get home' or 'I can make the Spider go away by the power of positive thinking'. It seems that the 'aboutness' relation between my 'Spider' file and the ice-field can tolerate categorially deviant beliefs only if the match between the file's Governing Conception and ice-field's category remains intact.

I take it that these examples provide some evidence for the MCP for the cases of thought and speech about numbers, and for the case of proper names and the thoughts we use them to express. But my concern in this paper is with the files of beliefs we form by taking what is delivered through attentional channels at face value, and with the perceptual demonstratives we use to express these beliefs. And I think it must be acknowledged that

for this case ordinary speaker intuition does not provide the kind of support for the MCP that it seems to in the cases I have just considered. Suppose that I say, pointing into the shadows, 'That is hungry'. You are in a position to see that there is no animal where I am pointing—my attention has been caught by a thickening in the pattern of light and dark. So this is a parallel to the situations where I use 'seven' as a label on a file whose Governing Conception does not match the category of numbers, or 'The White Spider' as a name on a file whose Governing Conception does (p.231) not match the category of geological features. When I point and say 'That is hungry', if I am referring to anything I am referring to the thing at the end of my attentional channel. But I think I am looking at an animal. So the Governing Conception of the file of beliefs I would express using my demonstrative does not match the category of the demonstrative's potential referent (a shadow). The MCP entails that this kind of case, where a file's Governing Conception does not match the category of its potential referent, is a case of reference failure. In the parallel cases involving 'seven' and 'The Spider', ordinary speaker intuition seems to be in step with this result. But in the perceptual demonstrative case ordinary speaker intuition produces conflicting verdicts. When I say 'That is hungry', you might say either 'It can't be hungry—it's just a shadow' or 'But there's nothing there'. The first response treats my use of 'that' as referring to a thing (a shadow), even though the file of beliefs on which I am using 'that' as a label does not have a Governing Conception which matches this thing's kind. The second treats me as failing to refer.

In §3.1, I shall argue that ordinary speaker intuition does, after all, provide some evidence for the MCP for the case of perceptual demonstratives. For the account of perceptual demonstrative reference which can be built around the MCP provides a reason to favour one set of intuitions about what perceptual demonstrative reference requires (the intuitions which treat cases like the 'That is hungry' case as cases of reference failure) while explaining why the contrary intuitions arise. So, given the MCP, we can explain ordinary speaker intuitions about reference and aboutness, without conceding that our ordinary notions of reference and aboutness are incoherent. But to show how this solution works I need to get the account of perceptual demonstrative reference I want to suggest on the table. So for the moment I shall just assume that the MCP does apply to the files of beliefs we form by taking what is delivered through attentional channels at face value and the perceptual demonstratives we use to express these beliefs. The rest of the section shows how this assumption combines with the empirical results about attention from §1 to yield an account of how this kind of thought and speech work.

## §2.2. Attention, Acquaintance, and Modal Containment

In §1, I told the following empirical story about visual attention. Our pre-attentive processing constructs maps of how features are distributed (p.232) in the space around us. The attentional spotlight is drawn to various kinds of configuration in these maps. The configurations that attract and hold our attention are 'visual objects'. The class of visual objects includes shadows, vertices, ripples on water, reflections on shiny surfaces, and parts of ordinary material things. So many of our attentional states, because they are not activated by ordinary objects, are not states of attending to ordinary things. But attentional states are relatively difficult to sustain over time. They are turned off by sudden disappearance and reappearance; by jumpy, piece-wise, or structure-undermining motion; and by the appearance of qualitative change which cannot be factored in terms of change in viewing conditions for relatively stable objects. So, though many of our attentional states are not states of attending to ordinary things, there is a connection between the conditions for attending and the characteristics that the philosophical tradition has taken to be constitutive of ordinary objecthood: apparent violation of these conditions turns attentional states off.

If the MCP is assumed, this empirical story transforms surprisingly easily into the foundation for an account of acquaintance-based thought about ordinary things. Here is how.

Recall the definition of 'acquaintance-based' thought that we started with:

S's thought about *o* is an 'acquaintance-based' thought iff S is in a position to have the thought in virtue of a perceptual link with *o* and without the use of any conceptual or descriptive intermediary.

Now suppose that an account of acquaintance-based thought must respect the MCP. This is to suppose that a perceptual link with an object puts you in a position to have an acquaintance-based thought about it only if the perceptual link puts you in a position to maintain a file of beliefs whose Governing Conception keeps faith with the object's category.

And now consider the empirical story about attention. According to this story, not every parcel of information you receive through an attentional link is coming from an ordinary object. When you are in an attentional state you might be attending to a patch of light, or a ripple, or a part of an ordinary object instead. But attentional channels are closed down by

apparent violations of ordinary objecthood. We cannot track 'things' whose spatio-temporal paths are discontinuous; which move in ways which do not maintain stable relations between parts; or which undergo qualitative change (**p.233**) which cannot be factored as change in the properties of causally unified things. And because attentional channels are closed down by apparent violations of ordinary objecthood, the only parcels of features which could be delivered through an attentional channel are parcels of features which could (categorially) be features of the same ordinary object. So, though not every parcel of information received through an attentional channel is information from an ordinary object, every such parcel of information has what I shall call 'ordinary object structure': every parcel of information delivered through an attentional channel consists in a combination of features which could belong to a single ordinary material thing.

A parcel of features being delivered through an attentional channel is not yet a file of beliefs. In talking about what is delivered through attentional channels we are talking about the contents of experiences, not the contents of thoughts. But now consider a file of beliefs which is formed by taking a parcel of features delivered through an attentional channel at face value. The parcel of features will have ordinary object structure. So the file of beliefs will have ordinary object structure too: if I take a parcel of features which could (categorially) all belong to a single ordinary object and form a file which contains just the information corresponding to these features, the resulting file will contain information which could (categorially) all be true of a single ordinary object.

And now suppose that I intend my deployment of the file to be experientially governed: I intend to maintain the file in a way which matches the structure that is already there in my experience. Then conceptual thoughts involving the file (for example, thoughts about how the object might develop across time and how it might relate to other objects) will also be constrained by ordinary object structure. To allow into the file beliefs or combinations of beliefs which stray outside what is categorially possible for ordinary objects is to break faith with the file's founding intention.

So the situation is this. Given the empirical story about attention, an attentional channel can deliver only parcels of features which could be features of a single ordinary thing. So a file of beliefs formed by taking what an attentional channel delivers at face value, and with the intention of keeping faith with the structure the channel carries, will be a file of beliefs

which could all be true of a single ordinary thing, and whose Governing Conception generates templates that ordinary things can fill. Given the MCP, a file of beliefs is about an object only if the file's Governing (p.234) Conception matches the object's category. But if a file of beliefs I form by taking what is delivered through an attentional channel at face value is about an object at all, it is about the object I am attending to—the object which is the source of the information in the file. So, assuming that an attentional link is sufficient to fix reference as long as the MCP is met, we have the following two results:

(1) if the visual object I am attending to is an ordinary object, the file of beliefs I form by taking what is delivered through the attentional channel at face value, and with the intention that my file be experientially governed, is a file of beliefs about this object.

(2) if the visual object I am attending to is not an ordinary object (it is a shadow, or a ripple, or . . . , instead) the file of beliefs I form by taking what is delivered through the attentional channel at face value, and with the intention that the file be experientially governed, is not about anything (because the category of the object the information in the file is from does not match the file's Governing Conception).

So attending to a visual object puts you in a position to have acquaintance-based thoughts iff the visual object is an ordinary object.

And that is the account of acquaintance-based thought about ordinary things that I want to propose.

In terms of the intuitive motivation for the MCP, the proposal can be put like this. The MCP gets its intuitive motivation from the thought that a simple model of representation which seems to work for toy cases (like the molecular model case) might be applicable across the board. In toy cases, we set up representations just by stipulating that one object is to stand for another, and the stipulation commits us to deploying the representing object only in ways which keep faith with the represented object's kind. To extend this model of representation to the cases of linguistic and mental representation we need to find counterparts, for the case of a singular term and a thought component, to the stipulation that *this* kind of sphere is to stand for *this* kind of atom. And it has always been hard to see what these counterparts could be. For to effect parallels to the stipulation you would need to be able to declare '*This* word (or *this* mental token) is going to stand for *that* object'. And to make this declaration you would already require a

way of speaking or thinking about the object. So it seems that to be in a position to make the stipulation which is supposed to be putting **(p.235)** you in a position to think or speak about the object you would need to be thinking or speaking about it already.<sup>17</sup>

But the proposal I have made offers an alternative account of how the relation between a basic representation and what it represents is set up. According to this account, what takes the place of the explicit statement, in the molecular model case, that black spheres stand for carbon atoms, is the intention to let the representation be experientially governed: the intention to form beliefs in a way which does not add to the structure which is already there in experience. Attentional channels deliver only parcels of information which have ordinary object structure. So a file of beliefs formed with the intention to take your experience when you are attending to an object at face value will have ordinary object structure—an ordinary object Governing Conception—too.

### §2.3. The Category of Ordinary Objects

So far I have been assuming what I take to be an uncontroversial account of the category of ordinary objects. Ordinary objects are located in space and time. They are causally unified at a time and over time. They trace out continuous spatio-temporal paths. And the relations between their parts are stable.

But the notion of ‘object’ employed by the account of acquaintance-based thought I have proposed is actually ‘programmatically’ in the following sense. The Governing Conception of a file of beliefs formed by taking what is delivered through an attentional channel at face value is inherited from the limits on what attentional channels can deliver. And being in an attentional state puts you in a position to maintain a file of beliefs about *o* iff *o* has a possible life history which matches this kind of Governing Conception. So the story about acquaintance that I have told will work for any kind of object whose life history matches the Governing Conceptions of files that we form by taking what our attentional channels deliver at face value. The notion of ‘object’ which is at work in the account of acquaintance-based thought I have proposed is just the notion of the kind of thing that matches the Governing Conception of an acquaintance-based file.

Here is an example to illustrate this point. Consider three situations. In the first you are attending to a member of the traditional category of **(p.236)**

ordinary objects. In the second, you are attending to a figure on a video screen which you cannot tell apart from an ordinary object. In the third you are attending to a 'stroboscopic' object: an object which can behave in all and only the ways that traditional ordinary objects behave, except that its spatio-temporal path is not really continuous—the path is gappy, but the gaps (points on the path which are not occupied by the object) are too small for us to detect. According to the view I have proposed, in the first situation you are in a position to have acquaintance-based thoughts about the thing you are attending to. And, according to this view, in the second situation you are not in a position to have acquaintance-based thoughts. For an acquaintance-based file of beliefs has a Governing Conception which generates templates that computer-generated figures cannot fill (templates containing possibilities like 'is moving towards me'; 'is touching me'; 'is leaving the room'). But what about the situation in which you are attending to a stroboscopic object? A stroboscopic object matches the Governing Conception of an acquaintance-based file just as well as a traditional ordinary object does. So, according to the view I have proposed, if you are attending to a stroboscopic object you are in a position to have acquaintance-based thoughts about it. This is not to say that ordinary objects and stroboscopic objects belong in the same category. But if there is a difference in category between them, it is a difference to which the account of acquaintance-based thought I have proposed is blind.

### §3. Acquaintance as Attentional Link, Illusions of Singular Thought, and Frege's Puzzle

So far in this chapter I have shown how abandoning the OEV in favour of the account of what perception delivers that some recent evidence suggests brings into view a new model of how acquaintance-based thought about ordinary objects might work. But the OEV forms only part of the traditional case against the claim that we can have acquaintance-based thoughts about ordinary things. I shall close by showing how the account of acquaintance-based thought I have proposed can be used to overturn the two other main elements in this case: the argument against acquaintance based thought about ordinary things from a theory of thought parallel to **(p.237)** the problem of empty names, and the argument against acquaintance-based thought about ordinary things from Frege's Puzzle cases.

#### §3.1. Illusions of Singular Thought

Consider the following examples:

case 1 It seems to you that you are looking at (and keeping track of) something in the shadows. You think 'It's round and it's slow-moving; maybe it's a hedgehog'. In fact there is no ordinary object there—you are looking at a thickening in the general pattern of light and dark.

case 2 It seems to you that you are looking at an object on the horizon. You think 'That is a very old building'. But actually your experience as of a thing on the horizon is caused by a combination of a tree in the middle distance and a speck on your glasses.

It will be convenient to have a term for the kind of thought about an ordinary object that you have when you are looking at it and keeping track of it which is neutral with respect to whether these thoughts are acquaintance-based thoughts. So I shall call these thoughts 'perception-based' thoughts about ordinary objects. In each of cases 1 and 2, you are unable to distinguish your situation from a situation in which you are having a perception-based thought about an ordinary thing. I shall call cases like these cases of 'illusory perception-based thought'.

A case of illusory perception-based thought is indiscriminable by the subject from a case of genuine perception-based thought. So it is, at first sight anyway, plausible that the difference between an illusory perception-based thought and its genuine counterpart is not, as I shall say, a 'normative' difference. Let  $\Sigma_1$  be a situation in which I am attending to an ordinary thing which seems to be both round and slow-moving (so roundness and motion with some slow velocity are among the features being delivered through my attentional channel). Let  $\Sigma_2$  be a situation indiscriminable by me from  $\Sigma_1$ , but in which there is no ordinary thing there. Suppose that, aside from the fact that I am attending to an ordinary object in  $\Sigma_1$ , and there is no object there in  $\Sigma_2$ , everything about my epistemic situation is the same in the two situations (I have the same background beliefs justified in the same ways, all features of context which might influence what it takes for a belief or a transition between beliefs to count as justified are held steady). And let 'That<sub>1</sub>' and 'That<sub>2</sub>' be demonstratives which I use to label (p.238) the files of beliefs formed by taking at face value what is coming through my attentional channels in  $\Sigma_1$  and  $\Sigma_2$  respectively. Then it is, at first sight, plausible that my experience in  $\Sigma_1$  justifies me in thinking 'That<sub>1</sub> is round' iff my experience in  $\Sigma_2$  justifies me in thinking 'That<sub>2</sub> is round'; that I am



justified in  $\Sigma_1$  in reasoning 'That<sub>1</sub> is round; that<sub>1</sub> is slow-moving; so maybe that<sub>1</sub> is a hedgehog' iff the same line of reasoning is justified in  $\Sigma_2$ ; and that if in  $\Sigma_1$  the rational way to try to catch the thing would be to reach out in such-and-such direction with such-and-such rapidity, the analogous movements would constitute the rational way to try to catch what I think I am attending to it  $\Sigma_2$ .

But the claim that there is no normative difference between an illusory perception-based thought and its genuine counterpart generates the following argument for the conclusion that our perception-based thoughts about ordinary objects are not acquaintance-based thoughts:

**1** The features of a subject's situation which generate entitlement to accept a perception-based thought are the same as those which would generate entitlement to accept its illusory counterpart. (From the initially plausible account of cases 1 and 2.)

so

**2** The features of a subject's situation which generate entitlement to accept a perception-based thought can be in place even if there is no ordinary object at the other end of the perceptual link. (From **1**.)

but

**3** An acquaintance-based thought is made available by the subject's standing in a relation to the object it is about (it is a thought you are in a position to have in virtue of a perceptual link with an object). (From the definition of 'acquaintance'.)

so

**4** The features of a subject's situation which generate entitlement to accept an acquaintance-based thought about an ordinary object are in place only if there is an ordinary object at the other end of the perceptual link. (From **3**, given that if the features which generate entitlement to accept a thought are in place the thought must be available.)

(p.239) so

**5** Perception-based thoughts about ordinary objects are not acquaintance-based thoughts. (From **2**, **4**)

The traditional response to this kind of argument was to point to sense data as a class of object for which illusory perception-based thoughts cannot arise, and say that acquaintance-based thoughts are perception-based thoughts about this special kind of object.

An alternative response is to reject the initially plausible account of examples like Case 1 and Case 2 which generates **1**. For example, you might argue as follows. There are cases which seem to the subject to be cases of perception-based thought about an ordinary object but where there is no ordinary object at the end of the perceptual link. But a subject's entitlement to accept a perception-based thought about an ordinary object, and the normative connections between a perception-based thought and a subject's subsequent actions, are both generated by a relation between the subject and the object the thought is about: the relation in which you stand to an object when you are attending to it. When you attend to an object, the fact that you are attending determines what it would be for your perceptual system to be delivering right information about the object: right information will be right information about *this* object—the one singled out by the fact that you are attending to it. And the fact that you are attending determines what it would be for your attempts to act on the object to succeed: successful action will at least be action on *this* object (the object singled out by your attention). If you are having an illusory perception-based thought, there is no object singled out as the object you would have to be right about for the thought to be true. And, in this case, there is no object singled out as the object that you must succeed in acting on if your intention to do something is to be fulfilled. So illusory perception-based thoughts are not norm-governed. And **1** and **2** are false: the features of a subject's situation which generate entitlement to accept a perception-based thought are in place only if there is an object at the other end of the perceptual link.<sup>18</sup>

**(p.240)** Because we cannot tell cases of illusory and genuine perception-based thought apart, to make this move is to reject the claim that normatively relevant differences in thought content are transparent to us: the claim that if my situation entitles me to accept thought  $\tau_1$  but not thought  $\tau_2$ , or licenses some action given  $\tau_2$  but not  $\tau_1$ , there must be a difference in content between  $\tau_1$  and  $\tau_2$  which I am in a position to detect. I take it that having to give up this claim is a cost.

I shall not survey the long history of attempts to offset or downplay this cost here. And I do not even want to suggest that it cannot be offset or downplayed to most people's satisfaction. Rather, I want to show that the

proposal about acquaintance-based thought that I have made in this paper provides a different solution to the problem. It provides a different solution because it provides a coherent way to reject the move from **3** to **4**.

Here is how. The account I have proposed distinguishes between the justification-conferring and aboutness-determining roles of attentional states. According to this account, an attentional state confers justification regardless of whether the subject is actually attending to any ordinary thing. This is because an attentional state confers justification in virtue of the fact that a parcel of features delivered through an attentional channel has ordinary object structure: if you form a file of beliefs by taking what is delivered through an attentional channel at face value, and intend to let your file be governed by the structure that is already there in your experience, your intention constrains you to maintain the file as an ordinary-object file. This account of how attentional states confer justification applies regardless of whether you are attending to an ordinary object or to some other kind of visual object. So illusory perception-based thoughts are norm-governed. And they are norm-governed for just the same reason as genuine cases of perception-based thought are: attentional states confer justification. But, on this account, the beliefs you form by taking what is delivered through an attentional channel at face value are beliefs about an object iff you actually are attending to an ordinary object. For an attentional link provides a way of thinking about a thing iff the MCP is satisfied—iff the Governing Conception of the file of beliefs you form by taking what is coming to you through the attentional channel at face value matches the category of the thing at the end of the channel. Files of beliefs formed by taking what reaches you through an attentional channel at face value always have ordinary object structure. So if you are not attending to an ordinary thing (**p.241**) (if you are attending just to a dot on a screen, or a ripple on water, or a thickening in the shadows) there is nothing that the beliefs you form by taking what comes through the attentional channel at face value are about.

It may seem that in this case the belief you form, though not about anything, still has representational content. For it seems that we can still say what would be required for such a belief to be true. Suppose I have a belief which I would express by saying 'That is F', where I am using the token of 'That' as a label on a file of beliefs that I form by taking what I receive through an attentional channel at face value. In fact, I am not attending to an ordinary object. But if I were, my belief would be true iff this object was F. So it is tempting to think that we have

My belief would be true iff, for some ordinary object  $x$ , the attentional state I am in was in fact a state of attending to  $x$ , and  $Fx$ .

and that once we have this claim nothing can stand in the way of  
My belief represents the world as being such that for some ordinary object  $x$  I am attending to  $x$  and  $Fx$ .

But this temptation should be resisted. For being in an attentional state is standing in a relation to a visual object—the visual object you are attending to. Attending to a different visual object is being in a different attentional state. So, given that I am not attending to an ordinary object, the proposed statement of what it would take for my belief to be true actually says something like (where  $\alpha$  is the visual object I am in fact attending to)

My belief would be true iff, for some  $x$  other than  $\alpha$ , my state of attending to  $\alpha$  was a state of attending to  $x$ , and  $Fx$ .

and this makes no sense. What we have to say instead is something like  
The belief I would form if, in a situation I cannot discriminate from this one, I was attending to an ordinary object,  $x$ , would be true iff  $Fx$ .

But this is no longer an account of what it would take for the belief I in fact form (in the illusory situation) to be true.

So the account of acquaintance-based thought I have proposed entails that **3** is true and **4** is false. You can have an acquaintance-based thought only if the thing you are attending to is an ordinary object. But the features of your situation which entitle you to hold an acquaintance-based belief (**p.242**) would be in place even if there were no ordinary object at the end of the attentional link: the entitlement is generated just by the fact that you are attending, regardless of whether there is any ordinary object that you are attending to. Illusory perception-based thoughts are norm-governed (because they are generated by taking what is delivered through attentional channels at face value), but they lack representational content (because there are no objects that they are about).<sup>19</sup>

Though cases of illusory perception-based thought are commonplace, they are actually the exception among cases involving attention to visual objects which are not ordinary objects. It is much more usual to attend to a visual object which is not an ordinary object (it is a shadow, or a reflection, or a part of an ordinary object instead) without ever forming a file of purported

acquaintance-based beliefs. Because attentional states are turned off by apparent violations of ordinary objecthood, many of these attentional states are so short-lived that you do not get the opportunity to open a corresponding file of beliefs. In other cases you do open a file, but it is not an acquaintance-based file: you provide a mediating description or sortal concept ('the shadow of that bird'; 'that patch of light'; 'his hand'). In this kind of case your file of beliefs is still a file of beliefs you are in a position to have because of a perceptual link. But now the perceptual link is descriptively or conceptually mediated. It is only in a third kind of case, commonplace but fortunately not too common, where you attend to something which is not an ordinary object, and keep your attention fixed, and (therefore) do open a file, but do not provide a mediating description (because the visual object you are attending to looks just like an ordinary object) that the degenerate situation I have suggested—justification without representational content—arises.

### §3.2. Acquaintance-Based Thought about Ordinary Things and Frege's Puzzle

Here is the problem that Frege's Puzzle apparently raises for anybody who wants to say that perception-based thoughts (p.243) about ordinary objects are acquaintance-based thoughts. Let  $\tau_1$  and  $\tau_2$  be perception-based thoughts about the same ordinary object. Then it appears that in some cases  $\tau_1$  and  $\tau_2$  are transparently about the same object while in others they are not. I might see a dog one day and then again the next and wonder whether they are the same or merely similar. And if I am in this position it would be irrational for me to move from 'That dog<sub>(yesterday)</sub> lives up the street' and 'That dog<sub>(today)</sub> has caught a squirrel' to 'A dog which lives up the street has caught a squirrel'. But if, looking at my dog, I think first 'He is hungry' then 'He is tired' the fact that the two thoughts are about the same thing is transparent to me: there is no extra information that I need to warrant the move to 'He is both hungry and tired'; it would be rational for me to make the move, and irrational for me to resist it.

Given that there does seem to us to be a distinction between these two kinds of case, to argue that there is in fact no such distinction is to deny that differences in content which make a difference with respect to how it is rational to deploy our thoughts need be visible to us. And, again, I take it that this result is a cost (though one that many philosophers have been willing to absorb).

But it is hard to see how to characterize the distinction between the two kinds of case without loss of the claim that perception-based thoughts about ordinary objects are acquaintance-based thoughts. For when you have an acquaintance-based thought about a thing there is no descriptive or conceptual mode of presentation to mediate your thought: you are thinking about the thing 'as itself'. So if you think two acquaintance-based thoughts about a thing you will be thinking of it in the same way (as itself) on each occasion. And in that case, there seem to be no resources with which to draw the distinction between cases where a pair of acquaintance-based thoughts are transparently about the same object and cases where they are not.

So, once again, we seem to be faced with a choice between abandoning the claim that our perception-based thoughts about ordinary objects are acquaintance-based thoughts, and abandoning the claim that normatively significant differences between thoughts are transparent to us.

The proposal I have made offers a solution to this problem because it provides a way to distinguish the cases without loss of the claim that perception-based thoughts are directly about objects. According to this proposal, perceiving an object puts you in a position to think about it because it puts you in a position to open a file of beliefs whose content is **(p.244)** derived from the information you receive through the perceptual link, and whose structure is inherited from the object the perceptual link connects you to. So this is a view on which perception puts you in a position to have a 'direct' thought about a thing in that when you perceive an object you have an open channel to it. The cases in which it is transparent to you that  $\tau_1$  and  $\tau_2$  are about the same object are just cases in which  $\tau_1$  and  $\tau_2$  are thoughts you can think in virtue of the openness of the same channel: if the information 'is hungry' and 'is tired' comes to you down the same channel, it ends up in the same file. That is why you do not need any extra information to conclude that something is both hungry and tired. But there is nothing to stop you from having two, or more, channels to the same object. This is what happens in cases where the fact that  $\tau_1$  and  $\tau_2$  are about the same thing is not transparent to you. In this case, 'is F' and 'is G' are delivered, through different channels, to different files: you would need to discover the truth of the identity statement to bring the two pieces of information together as information about a single thing.

## References

Bibliography references:

Campbell, John (2002). *Reference and Consciousness*. Oxford: Oxford University Press.

Cytowic, Richard (1995). 'Synesthesia: Phenomenology and Neuropsychology: A Review of Current Knowledge', *Psyche* 2. Online.

Driver, John, Greg Davis, Charlotte Russell, Massimo Turatto, and Elliott Freeman (2001). 'Segmentation, Attention, and Phenomenal Visual Objects', in B. Scholl (ed.), *Objects and Attention*, 61–95. Cambridge, MA: MIT Press.

Evans, Gareth (1982). *The Varieties of Reference*. Oxford: Oxford University Press.

— (1985). 'The Causal Theory of Names', in G. Evans, *Collected Papers*. Oxford: Oxford University Press.

Fine, Kit (2007). *Semantic Relationism*. Oxford: Blackwell.

Frege, Gottlob (1979). *Posthumous Writings*, ed. Hans Hermes, Freidrich Kambartel, and Friedrich Kaulbach; trans. Peter Long and Roger White. Oxford: Blackwell.

Grice, H. P. (1969). 'Vacuous Names', in Donald Davidson and Jaakko Hintikka (eds), *Words and Objections*. Dordrecht: Reidel.

Harrer, Heinrich (1976). *The White Spider*. London: Hart-Davis.

Hoffman, Donald D. (1998). *Visual Intelligence*. New York: Norton.

Jeshion, Robin (2001). 'Acquaintanceless *De Re* Belief', in J. Campbell, M. O'Rourke, and D. Shier (eds), *Topics in Contemporary Philosophy: Truth and Meaning*. New York: Seven Bridges Press.

**(p.245)** Locke, John (1975). *An Essay Concerning Human Understanding*, ed. P. Nidditch. Oxford: Clarendon Press.

Palmer, Stephen E. (1999). *Vision Science*. Cambridge, MA: MIT Press.

Pylyshyn, Zenon (2003). *Seeing and Visualizing*. Cambridge, MA: MIT Press.

— (2007). *Things and Places: How the Mind Connects with the World*. Cambridge, MA: MIT Press.

Quine, W. V. (1960). *Word and Object*. Cambridge, MA: MIT Press.

Recanati, François (1993). *Direct Reference*. Oxford: Blackwell.

Russell, Bertrand (1956). 'The Philosophy of Logical Atomism', in Robert Marsh (ed.), *Logic and Knowledge*. London: Routledge.

Scholl, Brian (ed.) (2001a). *Objects and Attention*. Cambridge, MA: MIT Press.

— (2001b). 'Objects and Attention: the State of the Art', in B. Scholl (ed.), *Objects and Attention*. Cambridge, MA: MIT Press.

— (2007). 'What Have We Learned about Attention from Multiple Object Tracking (and Vice Versa)?', in D. Dedrick and L. Trick (eds), *Computation, Cognition, and Pylyshyn*. Cambridge, MA: MIT Press.

Scholl, Brian, Zenon Pylyshyn, and Jacob Feldman (2001). 'What is a Visual Object? Evidence from Target Merging in Multiple Object Tracking', in B. Scholl (ed.), *Objects and Attention*. Cambridge, MA: MIT Press.

Spelke, Elizabeth (1988). 'Where Perceiving Ends and Thinking Begins: The Apprehension of Objects in Infancy', in A. Yonas (ed.), *Perceptual Development in Infancy*. Hillsdale, NJ: Erlbaum.

Strawson, P. F. (1974). *Subject and Predicate in Logic and Grammar*. London: Methuen.

Wittgenstein, Ludwig (1961). *Tractatus Logico-Philosophicus*, trans. D. F. Pears and B. F. McGuinness. London: Routledge and Kegan Paul.

— (1953). *Philosophical Investigations*, trans. G. E. M. Anscombe. London: Macmillan.

Wolfe, Jeremy (2000). 'Visual Attention', in K. De Valois (ed.), *Seeing*, 2nd edn. San Diego, CA: Academic Press.

## Notes:

(\*) Thanks to audiences at Sussex University, Warwick University, and the University of Waterloo. Thanks also to the University of Toronto Working Papers in Metaphysics and Epistemology Group. Particular thanks to John Campbell, Gary Ebbs, Robin Jeshion, Kirk Ludwig, and David Velleman for comments on and discussion of earlier drafts.



(1) For the OEV in the empiricist tradition see, for example, Locke (1975: Book II ch. 1): '*Light, and colours, are busie at hand every where, when the Eye is but open; sounds, and some tangible Qualities fail not to solícite their proper Senses, and force an entrance to the Mind . . .*'; Russell (1956 [1918]: 273): 'all the ordinary objects of daily life are extruded from the world of what there is, and in their place as what there is you find a number of passing particulars of the kind that one is immediately conscious of in sense'. The OEV has its behaviourist counterpart in Quine's claim that an observation sentence's content is fixed by the 'ocular irradiation patterns' which prompt assent to and dissent from it (1960: 32). For a psychologist stating the OEV see, for example, Spelke 1988: 229: 'Perceptual systems do not package the world into units. The organization of the perceived world into units may be a central task of human systems of thought. . . . The parsing of the world into things may point to the essence of thought and its distinction from perception. Perceptual systems bring knowledge of an unbroken surface layout . . . .'

(2) For surveys of evidence for this claim see Wolfe 2000: §2.5; Scholl 2001b; Pylyshyn 2003: esp. ch. 4, and 2007.

(3) See Pylyshyn 2003: ch. 2, for this sense of 'encapsulated'.

(4) For references to automatic spread of attention experiments see Scholl 2001b: §2.4. For a textbook account, see Palmer 1999: 547–548.

(5) See Driver et al. 2001: esp. §9. For additional discussions of amodal completion and the conditions under which it occurs see Hoffman 1998: ch. 3; Palmer 1999: 293–294 and 554–555.

(6) See Palmer 1999: 554 for a textbook statement of this claim.

(7) For accounts of multiple object tracking experiments, see Pylyshyn 2003: §5.3.1; Scholl 2001b: §2.5 and 2007.

(8) See Scholl 2001b: §6 and Scholl, Pylyshyn, and Feldman 2001.

(9) Grice 'Vacuous Names'; Strawson 1974; Evans 1982 and 1985; Recanati 1993; Jeshion 2001; Campbell 2002.

(10) The best-known explicit occurrence of this kind of constraint in the historical debate about representation is in the *Tractatus*, where Wittgenstein claims that *v* is a name for *o* iff for every possible sentence containing *v* there is a possible state of affairs containing *o* whose structure

matches the structure of the sentence, and for every possible state of affairs containing *o* there is a possible sentence containing *v* whose structure matches the structure of the state of affairs (1961: 2.1–2.225). The version of the constraint I am going to propose is weaker than Wittgenstein's version in important ways (see notes 12 and 13).

(12) In contrast, the *Tractatus* principle is an exact match principle, and so entails that I can think about *o* iff I know all of *o*'s possible occurrences in states of affairs: 'If I know an object I also know all its possible occurrences in states of affairs' (Wittgenstein 1961: 2.0123).

(13) In contrast, the *Tractatus* principle does make a sufficiency claim, and so entails that I cannot represent two objects which share all their possible occurrences in states of affairs as distinct.

(11) There is a hard question about what secures the unity of a file across time. One possibility is that files are individuated by their neurophysiological addresses. Another is that there is a specific relation of continuity which unites the slices of the same file. It would add too much complexity to consider this question here. But note that nothing in the view I am proposing requires that the notion of a file be treated as explanatorily prior to the notion of grasp of co-reference (cf. Fine 2007: 67–68).

(14) This view was proposed in Evans 1985.

(15) So the MCP provides an answer to the question raised by Evans (1982: §6.2)—the question of which information links with objects sustain demonstrative thoughts about them.

(16) For an account of this condition, see Cytowic 1995.

(17) Compare Wittgenstein's discussion of ostensive definition at the start of *Philosophical Investigations*, especially §§28–32.

(18) This is John Campbell's view (2002: 226).

(19) Though it is not possible to defend this claim here I think this proposal offers a way to make sense, for the case of sentences containing perceptual demonstratives, of Frege's claim that a sentence containing an empty name expresses a 'mock thought': see, for example, Frege 1979: 130. In the terms I have used here, a 'mock thought' will be a thought which is norm-governed but has no representational content.

