

The Sortal Dependence of Demonstrative Reference

Imogen Dickie

Abstract - ‘Sortalism about demonstrative reference’ is the view that the capacity to refer to things demonstratively rests on the capacity to classify them according to their kinds. This paper argues for one form of sortalism. §1 distinguishes two sortalist views. §2 argues that one of them is false. §3 argues that the other is true. §4 uses the argument from §3 to develop a new response to the objection to sortalism from examples where we seem to succeed in referring even though we get sortal classification wrong, or do not attempt to classify at all.

‘Sortalism about demonstrative reference’ (hereafter just ‘sortalism’) is the view that the capacity to refer to things demonstratively rests on the capacity to classify them according to their kinds. This paper argues for one form of sortalism. The paper has four parts. §1 distinguishes two sortalist views. §2 argues that one of them is false. §3 argues that the other is true. §4 uses the argument for sortalism from §3 to develop a new response to an old and intuitive objection to sortalism – the objection from the claim that sortalism is disproved by examples where we succeed in referring even though we get sortal classification wrong, or do not (apparently) attempt to classify at all.

I shall restrict the discussion to cases of ‘perceptual demonstrative reference’ – cases where you are in a position to refer to an object using ‘this’ or ‘that’ in virtue of a current perceptual link with it. I take these to be the central cases of demonstrative reference. There are hard questions about how what I say in this paper might extend to less central cases, for example, use of ‘this’ or ‘that’ to refer to something you remember, or to complex demonstrative reference (reference using terms like ‘that tree’ and ‘that dog’). But it is not possible to consider these questions here.

§1 Appropriation Sortalism vs. Rich Referential Sortalism

As a first step towards the distinction between sortalist views that I want to introduce, consider the following accounts of what it takes for a speaker to be using an expression to ‘refer to’ an object.

First (minimal) account: S is using α to ‘refer to’ o (in a context) iff S’s intention, in using α in the context, is to mark o ’s relevance to the conversation from S’s point of view. (For example, S may use α intending to greet o or to draw someone else’s attention to o ; S may use α intending to indicate o as an object desired or sought; or S may combine α with other expressions intending to say something about o .)

Second (rich) account¹: S is using α to ‘refer to’ o (in a context) iff S’s intention, in using α in the context, is at least to introduce o for the rest of the sentence containing α to say something about.

(Note that this distinction between rich and minimal senses of ‘refer’ is a distinction at the level of ‘speaker’s referent’ as opposed to ‘semantic referent’. A token expression’s ‘speaker’s referent’ is the object the speaker intends to be talking about in using it. A token expression’s ‘semantic referent’ is the object that the linguistic conventions associated with the expression together with relevant features of context determine as its contribution to whether the sentence containing it is true.² It is a familiar observation that speaker’s referent and semantic referent may be distinct. For example, suppose that you nod towards a man who looks like Russell and say ‘Bertrand Russell is paying the bill for the whole table’. Then the speaker’s referent for your token use of ‘Bertrand Russell’ is the man you are indicating, even though the semantic referent for your use of this term is Russell. My concern in this paper is with what kind of relation between speaker and object is required if the object is to be the speaker’s referent of a token perceptual demonstrative used by a

subject. Nothing I have to say here requires taking sides with respect to how possible divergences between speakers' referential intentions and linguistic convention should be treated. Throughout the paper, ' α refers to o ' should be read as ' o is the speaker's referent of α '.)

Borrowing some terminology from the literature on complex demonstratives³, I shall say that there is an object 'appropriated' by S's use of an expression iff as a result of this use the object is introduced as relevant to the conversation from S's point of view. In these terms, the minimal account says that appropriation is both necessary and sufficient for reference, while the rich account says that appropriation is necessary but not sufficient for reference. According to the rich account, a use of α that appropriates o is a referring use iff it introduces o as the thing the proposition S intends to communicate⁴ is about.

Given the two accounts of reference, and the notion of appropriation, the two sortalist views that I want to consider can be introduced like this:

First sortalist view: 'Appropriation sortalism' – the view that the capacity to refer demonstratively depends on the capacity to classify because the capacity to appropriate objects depends on the capacity to classify them.

Second sortalist view: 'Rich referential sortalism' – the view that the capacity to appropriate objects does not depend on the capacity to classify them, but the capacity to refer demonstratively depends on the capacity to classify because the capacity to use a demonstrative to say something about an object depends on the capacity to classify it.

If appropriation sortalism is true, even the capacity for minimal reference depends on the capacity to classify. If rich referential sortalism is true, the capacity for rich reference depends on the capacity to classify, but the capacity for minimal reference does not.

The next two sections argue that appropriation sortalism is false and rich referential sortalism is true.

§2 Appropriation Sortalism is False⁵

Appropriation sortalism is the view that the capacity to refer demonstratively depends on the capacity to classify because the capacity to appropriate objects depends on the capacity to classify them. I shall suppose that you are in a position to appropriate *o* iff you are in a position to attach an expression to *o* as intended referent, and that your perceptual processing puts you in a position to attach an expression to *o* as intended referent iff it presents you with *o* as a possible object of cognition – a thing that you might go on to have a thought about. Given these suppositions, appropriation sortalism stands or falls with the claim that our perceptual systems present us with particulars as possible objects of cognition only with the aid of our capacities for conceptual classification. And there is strong empirical evidence that this claim is false.

The first part of this section sets out some of the relevant empirical evidence. The second considers in more detail how this evidence impacts on appropriation sortalism.

The relevant empirical evidence concerns the nature of visual selective attention. ‘Selective attention’ is the process by which a specific part of the visual field is highlighted as a part information from which is available for conceptual thought. For example, you can see a large number of words on this page. All the words you can see are in your visual

field. But at each moment you will be able to tell what only a few of the words are. The words you are currently in a position to recognise are the ones you are currently attending to. Because what you are attending to determines which aspects of what is presented in your visual field you are in a position to think about, the processes involved in preparing the field over which attention roams, and in locking attention to one part of this field, are pre-conceptual. But psychologists have amassed strong evidence that visual selective attention is often directed to objects (in some sense of ‘object’). Since attention is the gatekeeper for conceptual thought, this is evidence for the claim that our capacities to parse scenes into objects, and select particular objects as potential objects of thought, are pre-conceptual too: if attention delivers up objects for conceptual thought, our perceptual processing can parse scenes into objects and lock onto specific objects before conceptualisation, and, in particular, classification under sortal concepts, come into play.

Here is some of the evidence for the claim that attention is often directed at objects (in some sense of ‘object’) that psychologists have found.⁶

§2.1 First piece of evidence – the automatic spread of attention



fig.1

fig.2

(both from Scholl 2002; ‘C’ stands for ‘cue’; ‘S’ for ‘same rectangle’; ‘D’ for ‘different rectangle’)

In a basic ‘automatic spread of attention’ experiment, subjects are presented with either two (*fig.1*) or three (*fig. 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 a decrement on the same rectangle as the initial cue than 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 every side. For the effect remains in *fig. 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 object boundaries – boundaries around one thing which is seen as lying behind another.⁷

§2.2 *Second piece of evidence - amodal completion*

'Amodal completion' is the process by which shapes are completed into apparent objects. For example, consider *fig. 3*:

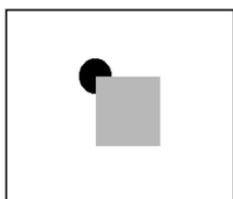


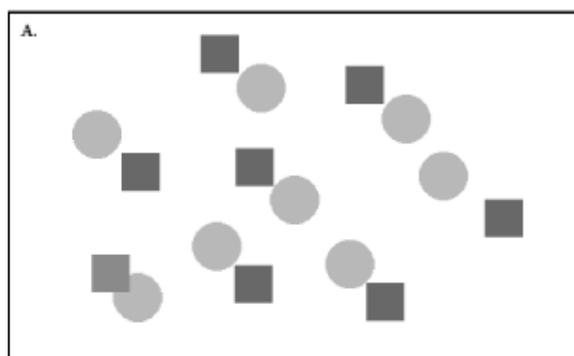
fig. 3 (from Driver et. al 2002)

When you look at *fig. 3*, your visual system completes the partially visible circle into a full circle that you see as lying behind an occluding square. Now consider *fig. 4* and *fig. 5*. And suppose that in each case you are asked to find the notched circle.

fig. 4 (from Driver et. al 2002)



fig. 5 (from Driver et. al 2002)



The notched circle is much harder to find in *fig. 5* than in *fig. 4*. And it has been found that this gap between *fig. 5* type cases and *fig. 4* type cases increases with the number of distracters in the diagram. The more extra shapes are in the frame, the harder it is to find the target shape in a *fig. 5* type case. In *fig. 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. Psychologists have concluded that the reason the notched circle is harder to find in *fig. 5* is that amodal completion is pre-attentive. When you look at *fig. 4*, your pre-attentive processing delivers 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. 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. 5* than in *fig. 4*, and why the gap increases with increasing numbers of distracters.⁸

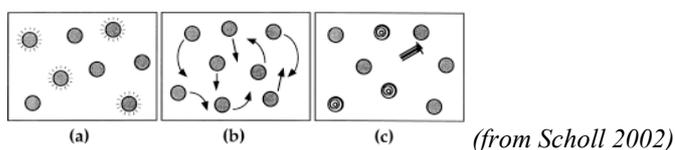
But amodal completion is completion into objects: when you look at *fig. 5*, your visual system completes the notched circle into a full circle seen as lying behind a square. So if attention roams over a field in which amodal completion has already taken place, our pre-conceptual processing is delivering a world already divided into objects of some kind.

§2.3 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. 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 attempt at once, the slower you will get⁹. 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.

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 dots are ‘flashed’ to distinguish them

from the others (a). The dots then move randomly around the display (b). When the motion stops the subject is asked to say whether a given dot flashed at the outset (c).



If parsing into apparent objects were post-attentive (so cognition-involving), you would expect us to be much better at this task when asked to track only one 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 performance 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.¹⁰

§2.4 *The argument for the falsity of appropriation sortalism*

With this empirical evidence in place, the argument for the falsity of appropriation sortalism can be stated like this:

1 Our pre-attentive perceptual processing parses the visual field into units (‘visual objects’) which attention selects as potential objects of thought. [Empirical claim]

2 Pre-attentive processing of perceptual input is pre-conceptual processing. [Consequence of the claim that what you are attending to determines what is available for conceptual thought.]

3 Pre-conceptual processing does not depend on conceptual thought.

4 Classifying objects according to their kinds (bringing them under sortal concepts) is a process at the level of conceptual thought.

So

5 The claim that perception presents us with particulars as potential objects of thought only with the aid of the conceptual capacity to classify them according to their kinds is false.

[From 1, 2, 3, 4]

But

6 If perception is presenting you with o as a potential object of thought, you are in a position to appropriate o . [Assumption from the start of the section.]

So

7 Appropriation sortalism (the view that sortalism is true because the capacity to appropriate objects depends on the capacity to classify them according to their kinds) is false. [From 5, 6]

To consolidate the case against appropriation sortalism, it is useful to consider how the empirical evidence described in this section impacts on one traditional argument for sortalism: the argument from the claim that pointing is ambiguous. Philosophers advancing

this argument say that the referent of a demonstrative is fixed by a mental or physical act of pointing which singles out one object as salient. But, they say, mental or physical pointing can single out a determinate object only if accompanied by a specification of kind. If I just point without specifying a kind, there will be no boundaries around the spatio-temporal part of the world I am pointing to. And in this case I cannot count as pointing at a unique object – there is nothing to make it the case that I am pointing at the chair across the room, rather than the arm of the chair, or the point on the chair’s arm that a line emerging from the end of my pointing finger would hit, or a momentary temporal stage of the chair, or the mereological sum of the chair and the table at which it is placed, or the bit of metal of which the chair is made (which may survive a change in which the chair itself is destroyed, or be destroyed by a change the chair survives). According to philosophers who advocate sortalism on the ground that unsupplemented pointing is ambiguous, the role of sortal concepts in enabling us to achieve demonstrative reference is to remove this ambiguity. These philosophers think of mental and physical pointing as having a two-part structure. First we indicate a direction. Then we provide a sortal concept to draw spatio-temporal boundaries around a part of the world in that direction. The object pointed at is the object with these spatio-temporal boundaries.¹¹

The empirical evidence I have set out in this section undermines this traditional argument’s central premiss – the claim that conceptually unaided mental or physical pointing is ambiguous. For this evidence suggests that our pre-conceptual processing parses the visual field into visual objects. So it suggests that the boundaries that proponents of the traditional argument think sortal concepts draw are already there in the pre-conceptual deliverances of our perceptual systems. And in that case, mental and physical

pointing do not have the two-part structure (indicate a general direction; provide a sortal concept to draw a boundary in that general direction) that proponents of the traditional sortalist argument supposed. ‘Mental pointing’ is attending to a visual object. Physical pointing serves to direct attention to a visual object. When you attend to a visual object there is a fact of the matter about the spatio-temporal spread of your attention. For example, if you are attending just to the arm of the chair, you will be quick to notice changes occurring on the chair’s arm, and slow to notice changes on the chair’s leg. (Compare the ‘automatic spread of attention’ experiments: if you are attending just to the chair’s arm, the boundaries around the chair’s arm are functioning like the boundaries of the initially cued rectangle in one of these experiments. Your information processing favours the attended region, so you are quicker at picking up changes within it.) Similarly, it is only if you were attending to the piece of metal of which the chair was made, rather than to the chair itself, that you would be able to keep attending through a change that the piece of metal but not the chair survives.

So the empirical evidence suggests that both appropriation sortalism, and the central premiss in the traditional ‘pointing is ambiguous’ argument for sortalism, are false. Our pre-conceptual processing has already worked out where the boundaries around the objects we are pointing (or attending) to are before conceptual thought comes to the table.¹²

§3 Rich Referential Sortalism is True

Recall the rich definition of ‘refer’ from §1: S is using α to ‘refer to’ o (in a context) iff S’s intention, in using α ’s in the context is at least to introduce o for the rest of

the sentence containing α to say something about. ‘Rich referential sortalism’ is the view that the capacity to refer in this rich sense depends on the capacity to classify. The argument for rich referential sortalism that I want to present goes like this. (For economy of expression, from now on I shall use ‘refer’ in the rich sense only unless otherwise stated.)

1 To refer to an object you have to be in a position to say things, right or wrong, about it.

2 To be in a position to say things, right or wrong, about an object you have to be in a position to ascribe properties to it.

3 Our capacities to ascribe properties to objects depend on our capacities to classify them.

So

4 Our capacities to refer to objects depend on our capacities to classify them.

1 follows from the rich definition of ‘refer’. I take 2 to be analytic: to ‘say something about’ an object just is to ascribe a property to it¹³. The argument is valid. So to use the argument to establish its conclusion it remains to establish 3 – the claim about the sortal dependence of property ascription.

As it stands, 3 is ambiguous between a strong reading

3a The capacity to ascribe a property to o depends on knowledge which kind of thing o is.

and a weak reading

3b The capacity to ascribe a property to o depends on knowledge that o is a thing of some determinate kind, but not knowledge which kind this actually is.

This ambiguity in 3 generates a corresponding ambiguity between strong and weak sortalist conclusions:

4a (Strong sortalism) S can refer to o only if there is some kind κ such that S knows that o belongs to κ .

4b (Weak sortalism) S can refer to o only if S knows that, for some kind κ , o belongs to κ .

The argument for the sortal dependency of property ascription that I am going to present is an argument for 3b. So it contributes to establishing weak sortalism (4b), rather than strong sortalism (4a). I shall consider the importance of the distinction between strong and weak sortalism when I discuss the intuitive anti-sortalist objection in §4.

The argument for the sortal dependency of property ascription takes its rise from the need to solve an old puzzle about what constitutes a genuine property. The puzzle goes like this.

Suppose

Generosity about properties – for every predicate, Φ , there is a property F such that an object o satisfies Φ iff o has F .

This generous criterion entails that objects have many properties that we do not ordinarily consider. For example, it entails that at the time of writing I have not only the properties of being human and being 170 cm tall, but also the properties of being shorter than my younger brother, being such that George W. Bush is president of the United States, and being either 160cm tall or such that Toronto is the capital of Ontario. In addition, suppose

Generosity about change – o changes between times t_1 and t_2 iff, for some property F , o has F at t_1 but not at t_2 , or has F at t_2 but not at t_1 .

Generosity about change combines with generosity about properties to entail that objects undergo many more changes than we ordinarily think they do. For example, this combination entails that I changed when my younger brother grew taller than me (I acquired the property of being shorter than him) and I will change when George Bush ceases to be president (I will lose the property of being such that George Bush is president).

It is a terminological question whether to allow that the generous definitions are really defining legitimate notions of ‘property’ and ‘change’. But the examples show that the generous notions of ‘property’ and ‘change’ are not our ordinary notions. We ordinarily take it that being human and being 170 cm tall are ‘properties’ of objects in a

way that being such that George Bush is president is not. And we ordinarily take it that when George Bush ceases to be president or my brother grows taller than me they (not I) are the ones who change. So our ordinary notions of 'property' and 'change' are much more restricted than the generous notions. The old problem about what constitutes a genuine property (and a genuine change) is the problem of what to say about the status of this restriction: What, other than just the fact that intuition favours them, is so special about ordinary properties and ordinary change?

Philosophers' attempts to answer this question have focussed on the connection between the notions of property, change, and cause¹⁴. Causal relations are generally taken to be relations between events. But they are relations that hold in virtue of the objects which participate in the related events and these objects' properties. One event causes another iff the objects involved are such that a change in one object's properties brings about a change in another's. So, for example, X's pulling of the trigger causes Y's death in virtue of the fact that a change in one of X's properties (the amount of pressure X is exerting on the trigger) brings about change in Y's properties. Philosophers looking for an anchor for the favoured status of ordinary properties have pointed out the connection between the notions of cause, change, and property. And they have argued that only ordinary properties and ordinary change relate to causation in a way that respects this connection. So the suggestion is that, for example, my ceasing to be such that George Bush is president can never be causally relevant in the way that changes in ordinary properties are.

Of course, it would be open to define a notion of 'generous causation' which would count generous properties as causally relevant and generous changes as causes and effects.

The appeal to the connection between the notions of cause, change, and property to secure the status of ordinary properties owes its attractiveness to the fact that it seems much more obvious that ‘generous causation’ would not be real causation than it is that a generous property is not a real property, or a generous change not a real change. For we appeal to causal relations to explain why events in the world unfold the way they do. And a ‘generous’ notion of causation would be a notion of causation without explanatory power.

There are many questions about how this connection between the ordinary notions of property and change and the notion of causation is to be made out. But I want to say just enough about the causal account of the status of ordinary properties to show how it entails the sortal dependency of property ascription. And all that is required for this purpose is the general claim that the special status we assign genuine properties rests on our taking it that an object’s genuine properties, and not its spurious properties, contribute to determining its causal path through the world.

As a first step towards showing how this general claim generates the sortal dependency of property ascription, let us make it more precise. The causal implications, for an object, of possessing a property always depend on which other properties the object has. For example, a boulder’s being round entails that it will roll down a plane inclined at a given angle and with a given surface iff it is also heavy, rigid, and smooth enough to roll given the angle and surface of the plane (and has no other property which would stop it rolling under the circumstances). An artefact’s being made of iron entails that it will, usually, cut wood if it has a sharp edge (and is constructed strongly enough), that it will hold water if it has an appropriate shape, that it will dissolve in acid at a given rate, depending on its surface area, and so on. Possession of a property endows an object with

‘conditional causal powers’ – powers to behave in various ways in various circumstances, conditional on what its other properties are¹⁵. So the claim that a thing’s genuine properties are all and only its causally relevant properties can be precisified into

The Causal Relevance Principle (hereafter ‘CRP’): *F* is a genuine property of object *o* iff possession of *F* endows *o* with some range of conditional causal powers.

Note that the CRP is neutral between the (extreme) metaphysical claim that the nature of a genuine property is exhausted by the causal implications of possessing it¹⁶ and the (more moderate) metaphysical claim that a genuine property is what generates (or provides the ‘categorical grounds’) for a cluster of causal implications¹⁷. The CRP is also neutral with respect to whether we must appeal to conditional causal powers to individuate genuine properties. The claim is only that the genuine properties endow objects with conditional causal powers. So, for example, it is consistent with the CRP to say that the property of being spherical is individuated purely geometrically (it is the property of having a volume determined by a specific equation), but that sphericity owes its status as a genuine property to the causal difference that it makes.

Now let us ask exactly how the notion of a conditional causal power must be understood if the CRP is to provide a solution to the problem about the status of ordinary properties.

If the CRP is to provide a solution to this problem, the notion of ‘conditional causal power’ must be understood in such a way that spurious properties (like being fifty miles south of a burning barn and being such that George Bush is president) do not endow things

with any. So, for example, the following will not do as an elucidation of the notion of a conditional causal power:

(CCP1) Possession of property F endows o with the power to X in situations of kind Σ given that o also has other properties Γ iff, if o has F and Γ , it will X in a situation of kind Σ .

For consider the following statement:

If o is such that George Bush is president, and o has the properties {being a boulder, being spherical, weighing 200kg} then if o is placed mid-way down a grassy slope o will roll to the bottom.

This statement is true. And it combines with CCP1 to entail

Possession of the property of being such that George Bush is president endows o with the power to roll to the bottom of a grassy slope if o also has properties {being a boulder, being spherical, weighing 200kg}.

So, given CCP1 and the CRP, the property of being such that George Bush is president emerges as a genuine property.

The problem with CCP1 is obvious enough. The left hand side of CCP1, with its use of ‘endows’, treats whether o has F as making a difference to which powers o will

have, provided that it also has the properties in Γ . But this claim about making a difference has been stripped from CCP1's right hand side. So CCP1 leaves out an essential part of the relation between possession of a genuine property and possession of conditional causal powers: the fact that possession of a genuine property makes a difference to which conditional causal powers an object has.

As a first step towards rectifying this problem, we might move from CCP1 to

(CCP2) Possession of property F endows o with the power to X in situations of kind Σ given other properties Γ iff, if o has F and Γ , it will X in a situation of kind Σ , and o will not X in this kind of situation given Γ alone.

But CCP2 will not do either. For consider this claim:

If Jack is such that George Bush is president, and Jack has the property of being able to discriminate person-sized objects against different coloured backgrounds at distances of up to 1000 feet, then Jack will see the president if looking in the right direction in a situation where George Bush is 1000 feet away in appropriate lighting conditions; and if Jack lacks the property of being such that George Bush is president he will not see the president in this situation.

This claim is true. And it combines with CCP2 to entail that the property of being such that George Bush is president does endow Jack with a conditional causal power. So, like CCP1,

CCP2 combines with the CRP to entail that being such that George Bush is president is a genuine property.

Again, it seems obvious enough what has gone wrong with CCP2. CCP2 fails to recognise that to count as endowing an object with a conditional causal power the possession of a property must make a difference to how the object would behave against a sufficiently varied range of background Γ 's and in a sufficiently varied range of Σ 's. So we need something like

(CCP3) Possession of property F endows o with the power to X in situations of kind Σ given Γ iff,

- (i) if o has F and Γ , o will X in a situation of kind Σ , and o will not X in this kind of situation given Γ alone, and
- (ii) for some $\Gamma_1, \dots, \Gamma_n$ some X_1, \dots, X_n and some $\Sigma_1, \dots, \Sigma_n$, where the Γ_i , X_i , and Σ_i are sufficiently varied, if o has F and Γ_i it will X_i iff introduced into a situation of kind Σ_i and will not X_i in this kind of situation given Γ_i alone.

CCP3 captures the fact that to count as endowing an object with a conditional causal power possession of a property must make a difference to the object's causal embedment in the world – to how the object will behave given a range of other properties and across a range of situations.

And now recall the various glosses philosophers have given on the notion of a 'sortal concept'. These include the following. κ is a sortal concept iff the question 'How many things falling under κ are there in such and such a region of the world?' always has

a determinate answer¹⁸. κ is a sortal concept iff for all x , if x falls under κ no part of x also falls under κ ¹⁹. κ is a sortal concept iff for all x , if x falls under κ the question ‘What is it?’, when asked with respect to x , may be answered by saying that x falls under κ ²⁰. κ is a sortal concept iff the fact that x is a κ determines which changes x can undergo without ceasing to exist²¹. κ is a sortal concept iff the things falling under κ share a characteristic kind of way of behaving²².

Lack of clarity about the relations between these various notions of sortal concept, and about which notion is at issue in any particular instance, is a long-standing source of confusion in the debate about sortalism. It is not possible to trace the various confusions, or clarify the relations between philosophers’ various notions of sortal concept here²³. But the sortalist argument I am constructing does not require a detailed survey of this kind. It requires only an observation that we are now in a position to make: the observation that, given CCP3, the CRP entails that the conditional causal powers with which possession of a genuine property endows an object depend on the object’s sort in a sense of ‘sort’ coinciding with at least some of the glosses on ‘sortal concept’ that philosophers have given. CCP3 supposes that o instantiates some background pattern of causal embedment, and explains the causal implications, for o , of possessing F in terms of the difference between variations on that pattern – the variation that o will instantiate if it has F , and the variation it will instantiate if it does not. So CCP3 entails that the causal powers with which possession of F would endow o do not depend just on o ’s other properties. They also depend on which background pattern of causal embedment o instantiates. But to give an object’s background pattern of causal embedment – an account of the ways it will behave given a variety of properties in a variety of situations – just is to give (one) answer

to the question ‘What is this object?’, and to give (one) classification of the object in terms of a characteristic way of behaving.²⁴ So examination of the notion of a conditional causal power has led us to the conclusion that whether possession of a property endows an object with conditional causal powers, and which conditional causal powers these are, depends on the object’s kind.

This conclusion entails a special form of the sortal dependence of property ascription (claim 3 in the argument from the start of the section). Here is how. Let us say that a ‘genuine’ property ascription is an ascription to an object of a property that would endow the object with conditional causal powers if the object actually had it. (So ‘I am 170cm tall’ is a genuine property ascription and ‘I am such that George Bush is president’ is not.) In addition, let us say that S ‘grasps a property ascription as genuine’ iff S knows that it is a genuine property ascription. And now let us ask what epistemic situation S must be in if S is to grasp an ascription of F to o as genuine.

One possibility is that for some kind κ , S knows that o is a κ , and knows a range of the causal implications of possession of F for κ ’s. In this case, S is in a position to grasp an ascription of F to o as genuine in virtue of knowing a specific range of conditional causal powers with which possession of F would endow o .

A second possibility is that, though not knowing o ’s kind, S knows that there is some kind κ such that members of κ would be endowed with conditional causal powers by possession of F , and o is a κ . In this case, S is in a position to grasp an ascription of F to o as genuine in virtue of knowing that there is some range of conditional causal powers with which possession of F would endow o , though S does not know which range this is.

But now suppose S neither knows, for some κ , that o is a κ , nor knows that there is some κ for which possession of F is causally significant such that o is a κ . Then S does not know any specific causal implications of F -ness for o , and S does not know that F -ness even has causal implications for o . And in that case S does not know that an ascription of F to o would be a genuine property ascription. For S does not know that possession of F would endow o with any conditional causal powers. (Suppose I know the following (and no more than the following) about sphericity. Sphericity is a property that endows o with the powers to roll down a smooth surface if o is a suitably heavy stone; bounce straight up when dropped on the floor if o is a ball made of suitably dense rubber; squash into *this* shape given *that much* pressure if o is a jelly. And suppose I do not know that o is a stone or a bouncy ball or a jelly. Then I do not know that possession of sphericity would endow o with any conditional causal powers.)

So to grasp an ascription of F to o as genuine, S must either know o 's kind, or know that o belongs to some range of kinds for which F -ness is causally significant. And that is the weak sortal dependence of property ascription for the special case of property ascriptions grasped as genuine.

To complete the argument for rich referential sortalism, it remains to show that claim 2 in the argument stated at the start of the section (the claim that to be in a position to say things, right or wrong, about an object you have to be in a position to ascribe properties to it) can be upgraded, for the case of things picked out perceptually, to

2* To be in a position to say things, right or wrong, about an object picked out perceptually you must be in a position to grasp some ascriptions of properties to the object as genuine.

But if the distinction between genuine and spurious properties is allowed its intuitive bite in the first place, it carries this upgrade with it. For the intuitive notion of a ‘genuine’ property – the notion we appealed to causal relevance to explain – is just the notion of a property that you would have to mention to give a complete account of what there is in the spatio-temporal world. The ‘genuine’ properties and changes are the properties and changes which would have to be included in a complete account of the history of the spatio-temporal world; the ‘spurious’ properties and changes could be omitted without loss of information because they are just superfluous constructions out of the properties spatio-temporal things really have and the changes they really undergo. But we can pick out perceptually only objects that are apparently located in the space around us. So to suppose that 2* is false is to suppose that you could be in a position to understand claims about a spatio-temporal thing even though you had no grasp of any difference that the truth of any of these claims might make to a right account of the history of the spatio-temporal world.

Putting together the pieces now on the table, the argument for rich referential sortalism that I want to propose goes like this:

1 To refer to an object you have to be in a position to say things, right or wrong, about it.

[The rich sense of ‘refer to’ from §1.]

2* To be in a position to say things, right or wrong, about an object picked out perceptually, you have to be in a position to grasp some ascriptions of properties to the object as genuine.

3b* Grasp of some ascriptions of properties to o as genuine requires at least knowledge that, for some κ where possession of the property endows κ 's with conditional causal powers, o is a κ .

therefore

4 S can refer to an object, o , picked out perceptually only if S knows that, for some kind κ , o belongs to κ ²⁵.

It is important to emphasize that the sense of 'refer to' in this argument is the rich sense from §1, and that the argument is (therefore) an argument for rich referential sortalism rather than appropriation sortalism. So it is consistent with 4 to allow that a subject may be in a position to introduce an object picked out perceptually as relevant to the conversation from his or her point of view without being in a position to ascribe properties grasped as genuine to the object, and without knowing that for some κ the object belongs to κ . In a case of this kind, the subject is in a position to 'appropriate' o (that is, 'refer to' o in the minimal sense distinguished in §1), but not to 'refer to' o in the rich sense. And I take the argument from §2 to have established that in fact the capacity to appropriate perceived objects does not rest on the capacity to classify them. So I am allowing both of the following:

(i) A perceptual link with an object may put an ordinary adult subject in a position to introduce o as relevant to the conversation from his or her point of view even though he or she does not know that, for some kind κ , o is a κ .

(ii) A subject who has yet to acquire, or has lost the capacity to ascribe predicates grasped as genuine may have the capacity to introduce a perceived object as relevant to the conversation from his or her point of view.

§4 The objection to sortalism from examples where reference succeeds while classification fails

So far in this paper I have distinguished two varieties of sortalism and argued that one (appropriation sortalism) is false, while the other (rich referential sortalism) is true. Now I want to consider how the version of sortalism I have proposed fares in the face of the obvious intuitive objection to any sortalist view: the objection from examples which seem to present clear cases of reference succeeding where classification fails.

Here is a range of examples of this kind²⁶.

CASE 1 You stumble upon a thing of a kind you have not come across before. You have no idea where it might fit in your system of sortal classification – whether it is animal, vegetable, or mineral; whether it is an artefact or something that came to exist without human intervention. But it seems that you can refer to it, in the richer sense of ‘refer’: there it is, sitting on your desk, and you say ‘That is yellow; it is cubical; I wonder what it is’.

CASE 2 You are receiving a lot of perceptual information from a thing, and much of it is veridical. But, carelessly, you misapply your capacity to classify, and reach a wrong conclusion about the kind of thing it is. (You conclude that it is a person when actually it is a very life-like waxwork.) But it seems that you can refer (richly) to the thing: when you say ‘He’s standing inside the rope-line’, you seem to be ascribing a property to the thing you are mistaken about.

CASE 3 You can see a thing moving around in the distance. It is too far away for you to tell what kind of thing it is. But it seems that you can refer (richly) to it (you and your friend might be wondering whether it is a bird, a plane, or superman, and trying to work out how fast it is actually moving.)

CASE 4 Your attention has been caught by what seems to you to be an animal in the shadows. Actually there is no ordinary object there. The ‘thing’ (or ‘visual object’) that you are attending to is just a thickening in the pattern of light and dark. It seems that you are in a position to refer (richly) to the thing: otherwise you would not be in a position to speculate about its size and likely fierceness.

The standard sortalist response to these examples has been to make a blanket appeal to the distinction between strong and weak versions of sortalism drawn at the start of §3. Philosophers making this response acknowledge that the examples do undermine strong sortalism:

If S can refer to o then, for some kind κ , S knows that o belongs to κ .

But they contend that the examples do not undermine weak sortalism:

If S can refer to o then S knows that, for some kind κ , o belongs to κ .

Proponents of the standard response then claim that they are endorsing only weak sortalism, and that in each of cases 1 – 4, though the strong sortalist demand (for knowledge which kind o falls under) is not met, the weak sortalist demand (for knowledge that there is some kind that o falls under) is.²⁷

At first sight, the argument for sortalism I have presented seems to provide an ideal set-up for this standard response. For the argument is an argument for weak sortalism. So the standard response – suggesting that the examples are not counterexamples to weak sortalism – looks like a natural move.

However, I want to suggest that the standard response is a mistake. As a first indication why, consider the fact that, on the face of it, Cases 1-4 seem to be of quite different kinds. Cases 1 and 3 are cases of classification apparently withheld, rather than cases of misclassification. And they are cases in which you withhold classification for different reasons. In Case 1 you know a lot about the thing – all its observable properties are transparent to you – it is just that you cannot classify it because you cannot tell where it fits in your system of sortal classification. In Case 3 your inability to classify seems to trace to a quite different source – you cannot classify the thing because it is so far away that you cannot discern its basic observable properties (like colour, shape, texture, and

degree of solidity). Cases 2 and 4 are cases of misclassification, rather than classification withheld. But, again, they are cases in which classification seems to fail for very different reasons. In Case 4 you are hazarding a classification on the basis of sketchy information (the thing seems to have a shape, colour, and degree of solidity which it in fact lacks). In Case 2, your information about the thing is mostly veridical – you are right about many of the thing’s basic observable properties. And most of the time if you classify as a person a thing with these basic observable properties you will be right. It is just that in this situation your usually reliable method of classification goes awry. So cases 1 – 4 look, on the face of it, very different. But the standard sortalist response gives them a uniform treatment – it treats them all as cases in which you know that the thing falls under some kind without knowing which.

I suggest that the blanket appeal to weak sortalism in response to cases 1 – 4 obscures the way the relation between demonstrative reference and sortal classification actually works. The rest of this section uses the argument for rich referential sortalism from §3 to develop an alternative treatment of the problem examples that sets this mistake right.

The alternative treatment is built around the role that argument of §3 assigns to the capacity to classify a thing in enabling a subject to refer to it demonstratively. According to this argument, the capacity to refer demonstratively depends on the capacity to classify because to refer to a thing demonstratively you need to be able to grasp ascriptions to the object of some properties as genuine; grasp of an ascription to *o* of *F* as genuine requires either knowing some of the specific conditional causal powers with which possession of *F* would endow *o*, or knowing that possession of *F* would endow *o* with some conditional

causal powers though you do not know which. The more specific type of knowledge requires knowing the object's kind, and the less specific requires knowing that the object is a member of a kind for which F is causally significant.

So the argument I have proposed generates criteria for drawing two distinctions between examples of apparent successful demonstrative reference without successful classification.

The first is a criterion for the distinction between cases of successful reference and cases of reference failure. According to the argument I have proposed, your perceptual link with o puts you in a position to refer to o demonstratively only if it puts you in a position to know, for some range of properties, that possession of these properties would endow o with conditional causal powers.

The second is a criterion for distinguishing cases in which you have knowledge of classification (cases where there is some κ such that you know that o belongs to κ) from cases in which you have mere knowledge of classifiability (cases where you know only that, for some κ , o belongs to κ). According to the argument I have proposed, you have knowledge of classification iff you know a range of specific causal implications with which possession of various properties would endow the object. You have mere knowledge of classifiability iff you know no specific causal implications, but know that the object belongs to a kind for which a specific range of properties is causally significant.

Armed with these two criteria, let us consider each kind of problem example in turn.

First consider examples like Case 1. In this kind of example, a thing's observable properties are transparent to you, but you have no idea where to place it in your system of

sortal classification. The standard sortalist move has been to acknowledge that in this situation you do not know the thing's kind, but argue that you do know that it is a thing of a kind. But if the argument for sortalism I have proposed is right, this application of the standard move is wrong. For in Case 1 there is a range, though a small one, of conditional causal powers with which you can tell that specific genuine properties endow things of *o*'s kind. Suppose you can see that the thing is cubical and yellow. You pick it up. It holds its shape under the pressure of your grip. You put it on the desk, where it sits, yellow and cubical. You can push it around with a finger. In this case, you know that the thing has the power to stay at rest on the ground, or on a flat wooden surface, given its observed colour, shape, and mass; the power to hold a constant shape and colour under the kind of pressure exerted by your grip and through a change in position; the power to slide across a flat wooden surface given a small push, and, again, given the shape, colour, mass, and degree of smoothness that you have observed it to have. In this kind of case, where the thing's observable properties are transparent to you, you are watching a section of its causal history unfold in front of you. You can see its interactions with other objects. You can see how its properties are developing given external influences. So, for the properties you can see (or otherwise perceive) that the thing has, you know a basic range of conditional causal powers with which possession of these properties endows the thing. In this case you do have specific knowledge which kind of thing you are dealing with: it is a thing of *this* kind, a thing which behaves in *these* ways given *these* properties and *this* stable situation.

I shall call the kind knowledge you have when you are in this position 'ostensive' knowledge of kind.

Ostensive knowledge of kind can fill the role identified for knowledge of kind in the argument for sortalism that I have proposed. There are other roles, identified by other arguments for sortalism, that ostensive knowledge of kind cannot fill. For example, consider the traditional argument for sortalism from the alleged ambiguity of pointing sketched at the end of §2. According to this argument, the job of sortal concepts in enabling demonstrative reference is to draw a boundary around a spatio-temporal region in a direction indicated by a mental or physical demonstration. So philosophers whose sortalism is motivated by this argument will think that knowledge of kind must be knowledge of something comprehensive enough to delineate a determinate spatio-temporal region. Similarly, some philosophers have suggested that knowledge of kind involves knowledge where the thing fits into a system of sortal classification, so that to know a thing's kind you need to know which kinds of things count as more or less similar to it. But the argument for sortalism that I have proposed assigns no role to either of these more demanding forms of classificatory knowledge. So, from the point of view of the sortalist view I have proposed, the fact that ostensive knowledge of kind falls short of what some philosophers discussing sortalism have assumed knowledge of kind would involve is no objection: ostensive knowledge of kind is a type of knowledge of kind because it plays knowledge of kind's role.

Now consider Case 2. Here, as in Case 1, many of the thing's observable properties are transparent to you. But in Case 2 you are making a wrong classification. You are thinking 'That is a person', when actually the thing you are attending to is a waxwork.

I suggest that this case is actually a parallel to a familiar kind of situation involving 'knowledge which' or 'knowledge who'.

Suppose you and I are talking face to face, and I am under the misapprehension that you are the Secretary General of the United Nations. Then I know who you are ostensibly (you are *this* person). I also have a false belief about who you are. This false belief will combine with my background beliefs to generate many other false beliefs, for example, the belief that you will be addressing the United Nations tomorrow; the belief that you orchestrated the deployment of a peacekeeping force to such-and-such country last year; and so on. But there is no reason to think that my false belief about who you are (that you are the Secretary General of the United Nations) undermines my true, ostensive, one (that you are *this* person).

I suggest that Case 2 has just this structure. In Case 2, many of the observable properties of the thing you are attending to are transparent to you. So, just as in Case 1, you have ostensive knowledge what kind of thing it is. You also have a false belief about its kind: the belief that it is a person. So you believe both ‘That is of *this* kind’ and ‘That is a person’. And, because your beliefs about the causal implications of *F-ness* for a thing depend on your beliefs about its kind, your false belief about the thing’s kind carries with it false beliefs about some of the causal implications of *F-ness* for the thing. But the mistake about kind, and your false beliefs about some of the causal implications of *F-ness* for the thing, do nothing to undermine the fact that a portion of the thing’s causal history is unfolding in front of you. Though you have many false beliefs about the causal implications of *F-ness* for the thing, you also have many true ones. So you still have ostensive knowledge of the thing’s kind. And this specific knowledge of kind is putting you in a position to refer.

Next consider Case 3 (the case of the thing visible only as a speck moving around in the distance). Suppose that the thing seems to be moving in the way characteristic of material objects (for example, tracing out a continuous spatiotemporal path). And suppose that whatever qualitative changes the thing seems to be undergoing are explicable in terms of a combination of changes in viewing conditions (maybe it is getting closer or farther away, or passing through a semi-opaque cloud or a shadow) and changes in a causally unified thing. In this case, you would be justified in believing that the thing is an ordinary material object. So if it actually is an ordinary material object (as opposed to just appearing to behave like one) you are in a position to know that it is. But the class of ordinary material objects just is the class of things for which we know observable properties to have causal implications. (It is doubtful that observable properties have causal implications for non-ordinary ‘things’ such as shadows, reflections, and mereological sums. But even if they do, knowledge that they do does not play a role in ordinary demonstrative reference, for we as ordinary speakers have no such knowledge.) So in this case you are in a position to know that, for some κ for which observable properties are causally significant, o is a κ . And here the move that the standard response to the problem examples insists on across the board is the right one: your capacity to refer to the thing is secured by the fact that you know that it is classifiable, though you are not in fact in a position to classify it.

Because this is where my treatment of the supposed counterexamples to sortalism lies closest to the standard response, I shall take a little time to set out the differences between the standard response and my own appeal to weak sortalism. There are three main points of comparison.

The first is just the point about uniform treatment of apparently diverse examples that I have already made. The standard response to the problem examples appeals to weak sortalism across the board. I have appealed to weak sortalism in only one kind of case.

The second point of comparison concerns the fit between the appeal to weak sortalism and the argument used to establish sortalism in the first place. The argument for sortalism that I have presented both identifies a role for knowledge of specific classification when we have it, and entails that we can refer with mere knowledge of classifiability. From the point of view of this argument, the ready availability of examples where reference succeeds given only knowledge of classifiability, not knowledge of classification, is an indication that we are on the right track. In contrast, the standard extant arguments for sortalism seem, at first sight, to support strong sortalist conclusions. For example, consider again the argument from the alleged ambiguity of pointing. According to a proponent of this argument, a subject referring to a thing demonstratively needs to supply a sortal concept to determine where the spatio-temporal boundaries of the demonstrated object lie. But this job of delineating a demonstrated object seems to require a specific sortal classification – if pointing is ambiguous, pointing and saying ‘the thing of some kind in that direction’ will be ambiguous too. So when this traditional argument for sortalism is combined with the standard response to Cases 1- 4 the result just looks *ad hoc* – it looks like the sortalist is stepping back to a position weaker than the position demanded by the argument used to support sortalism in the first place.

The third point of comparison between the appeal to weak sortalism that I am suggesting and the blanket appeal made by the proponent of the standard response

concerns the need to provide an account of what knowledge of classifiability without knowledge of kind consists in.

Proponents of the standard response have tended to gesture towards the notions of ‘approximation to’ or ‘homing in on’ a kind to fill this role²⁸. But this gesture just does not seem to fit the problem examples. There is no sortal concept that you are homing in on when you switch back and forth among the ‘bird, plane, Superman’ hypotheses. When you say ‘I have no idea what this is’ you may really have no idea. And it seems that you can get a thing’s sort wildly wrong (for example, mistaking a waxwork for a person) while still managing to refer to it.²⁹

The argument for sortalism I have proposed generates an account of knowledge of classifiability without knowledge of classification that has nothing to do with ‘approximation’ or ‘homing in’. Given this argument, knowledge of classifiability without specific knowledge of kind is knowledge that the appropriated object falls under some kind for which you know, for some specific properties, that these properties are causally significant.

Finally, consider examples like Case 4 (where you mistake a thickening in the shadows for a dog). In this kind of case, you have a perceptual link with a thing (you are attending to a visual object). But your perceptual link does not put you in a position to know, for any F , a range of conditional causal powers with which being F would endow the thing. And it does not put you in a position to know, for any F , that the thing belongs to a kind for which F -ness is causally significant. So the view I have proposed treats this as a case of reference failure.

It may seem that to diagnose reference failure in Case 4 is to concede that here we have a counterexample to the sortalist view I have proposed. For someone might argue like this. ‘Ordinary speaker intuition wants to treat Case 4 as a case of successful reference. So an account of the conditions for reference which entails that reference fails in Case 4 does not capture how our ordinary practices of referring and ascribing reference work. Case 4 shows that we can refer when we lack both specific knowledge of classification and knowledge of classifiability. So it shows that the sortalist view proposed in this paper cannot be right.’

But how clear is it, really, that ordinary speaker intuition wants to treat Case 4 as a case of successful reference? Here are two ways that a conversation in a Case 4 situation might go:

Me – ‘That is a dog’

You (having got in a position where you can see better) – ‘No, it’s a shadow’.

Me – ‘That is a dog’

You (having got in a position where you can see better) – ‘No, there’s nothing there’.

In the first exchange you seem to be treating me as referring to a thing (a shadow) that I then misclassify. In the second you seem to be treating me as failing to refer.

Given that both exchanges are coherent, it seems that ordinary speaker intuition does not produce a consistent result with respect to whether cases like Case 4 involve successful reference. So these examples raise a challenge that an account of how

demonstrative reference works should be able to meet: to explain, in terms of the conditions for successful demonstrative reference, why ordinary speaker intuition should deliver conflicting results about this kind of case.

And once this challenge has been recognised, Case 4 emerges as raising just as much of a problem for anti-sortalists as it does for sortalists. For each kind of view generates a univocal verdict with respect to Case 4 examples. The view I have proposed entails that they involve reference failure. And an anti-sortalist seems to be stuck treating them as cases where reference succeeds. So, given that ordinary speaker intuition delivers conflicting results, anyone proposing either view has work to do to explain why the conflict arises.

This is the hard problem about the relationship between reference and classification that remains when the foundations for the new approach to sortalism I am advocating are in place. And it arises for sortalist and anti-sortalist alike.

It is not possible to provide a full discussion of this hard problem here. So I shall not try to. Instead I shall close with a sketch of the kind of solution to the problem that the view I have proposed generates.

The solution is most easily brought out by way of a parallel with an example involving 'I'. Suppose that the colleague in the next room has set up his or her computer to say 'I'm finished' in a human-sounding voice whenever it completes a task. You do not know this. One day you are talking with this colleague in your office and a token of 'I'm finished' rings loud and clear through the open doors. You say to your colleague 'He's finished'. Your colleague says 'Oh, there's nobody there – it's just my computer'.

Now consider the reference-fixing rule for 'I':

A token of 'I' refers to the agent who produces it.

Given this rule, the computer-generated token of 'I' does not refer. But it does appropriate. It appropriates the object that produces it. It is because this object does not meet an extra condition required by the reference-fixing rule (that the referent be an agent) that the token does not refer. And note that, even though the token of 'I' in the example does not refer, it serves as the antecedent of an anaphoric pronoun that does ('it' in 'it's just my computer'). So an indexical that merely appropriates an object can serve as the antecedent to an anaphoric pronoun that refers to it.

I want to suggest that Case 4 (where my attention has been caught by a thickening in the shadows) should be understood in this same way. When I say 'That is a dog', my use of 'that' appropriates the visual object I am attending to. I believe that *F-ness* would have one set of causal implications for the object (the causal implications that *F-ness* has for a dog). Actually, the causal implications of *F-ness* for the object that my use of 'that' appropriates would be very different. So I am not in a position to grasp ascriptions of genuine properties to the thing appropriated by my use of 'That'. So, given the argument of §3, this is a case of appropriation without reference. But an anaphoric pronoun can refer even though its antecedent merely appropriates (just as in the case of computer-generated 'I'). When you say 'No, it's just a shadow', your use of 'it' refers to the thing that you are in a position to refer to though I was not. The alternative response, 'No, it's nothing' is the parallel to 'It's nobody' in the 'I' case: a response that acknowledges that the initial demonstrative does not refer to the object it appropriates.*

Imogen Dickie,
 Philosophy Department, University of Toronto,
 170 St George St., Toronto, Ontario, Canada M5R 2M8.
 imogen.dickie@utoronto.ca

References

- Ayers, M. (1974), 'Individuals without Sortals', *The Canadian Journal of Philosophy*, 4 (1): 113-48.
- (1997), 'Is "Physical Object" a Sortal Concept? A Reply to Xu', *Mind and Language*, 12 (3-4): 393-405.
- Campbell, J. (2002), *Reference and Consciousness*. Oxford: Oxford University Press.
- (2006), 'Does Visual Reference Depend on Sortal Classification? Reply to Clark', *Philosophical Studies*, 127 (2): 221-37.
- Dickie, I. (2010), 'The Generality of Particular Thought', *Philosophical Quarterly*, 60 (240): 508-31.
- Driver, J., G. Davis, C. Russell, M. Turatto, and E. Freeman. (2002), 'Segmentation, Attention, and Phenomenal Visual Objects', in Scholl (ed) *Objects and Attention*. Cambridge: MIT Press: 61-95.
- Dummett, M. (1973), *Frege: Philosophy of Language*. London: Duckworth.
- (1981), *The Interpretation of Frege's Philosophy*. London: Duckworth.
- Evans, G. (1982), *The Varieties of Reference*. Oxford: Oxford University Press.
- Frege, G. (1950), *The Foundations of Arithmetic*. (J. L. Austin, Trans.). Oxford: Blackwell. (Original publication date 1884)
- Grandy, R. E. (2007), 'Sortals', *The Stanford Encyclopedia of Philosophy*, E. N. Zalta (ed.), <http://plato.stanford.edu/archives/sum2007/entries/sortals>.
- Kripke, S. (1977), 'Speaker's Reference and Semantic Reference', *Midwest Studies in Philosophy*, 2 (1): 255-76.
- O'Craven, K. M., P. E. Downing, and N. Kanwisher. (1999), 'fMRI Evidence for Objects as the Units of Attentional Selection', *Nature*, 401: 584-87.
- Palmer, S. E. (1999), *Vision Science*. Cambridge: MIT Press.

- Pylyshyn, Z. (2003), *Seeing and Visualising – It's Not What You Think*. Cambridge: MIT Press.
- (2007), *Things and Places – How the Mind Connects with the World*. Cambridge: MIT Press.
- Quine, W. V. (1953), 'Identity, Ostension, and Hypostasis', in W. V. Quine, *From a Logical Point of View*. Harvard: Harvard University Press.
- Quine, W. V. (1960), *Word and Object*. Cambridge: MIT Press.
- Scholl, B. (2002), 'Objects and Attention: the State of the Art', in Scholl (ed.) *Objects and Attention*. Cambridge: 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: MIT Press.
- Shoemaker, S. (2003a), 'Causality and Properties', in Shoemaker, *Identity, Cause, and Mind*. Oxford: Oxford University Press.
- (2003b), 'Identity, Properties, and Causality', in Shoemaker, *Identity, Cause, and Mind*. Oxford: Oxford University Press.
- Siegel, S. and M. Glanzberg. (2007), 'Presupposition and Policing in Complex Demonstratives', *Nous*, 40 (1): 1-42.
- Strawson, P. F. (1959), *Individuals*. London: Methuen.
- Wiggins, D. (1980), *Sameness and Substance*. Oxford: Blackwell.
- (1997), 'Sortal Concepts: a Reply to Fei Xu', *Mind and Language*, 12 (3-4): 413-21.
- (2001), *Sameness and Substance Renewed*. Cambridge: Cambridge University Press.
- Xu, F. (2007), 'Sortal Concepts, Object Individuation, and Language', *Trends in Cognitive Science*, 11 (9): 400-406.

¹ The rich definition's pedigree includes Quine (see for example, (1960): 141-42, 177); Frege (I take it that the context principle (1950): 116 requires something like the rich definition); Evans (the rich definition is entailed by Evans's 'Generality Constraint' (1982): 100-105 and the importance he attaches to it); and many others.

² Compare Kripke (1977): 262-264.

³ See Siegel and Glanzberg (2007): 1.

⁴ Note that a rich referring use introduces *o* as relevant to the proposition *S* intends to communicate, rather than the proposition 'strictly and literally expressed' by *S*'s utterance. This is in keeping with the fact that distinction I have drawn between 'rich' and 'minimal' reference is a distinction between senses of 'speaker's reference'. See note 2 and accompanying text.

⁵ The argument of this section parallels the argument of Campbell (2006): 226-31, though Campbell does not distinguish the conditions for reference from those for appropriation, and thinks the argument establishes the falsity of sortalism about reference, not just the falsity of appropriation sortalism.

⁶ The evidence I cite is all 'behavioural' evidence for the object-directedness of attention. A separate group of studies have found neurophysiological evidence for this conclusion. Attention to one attribute of an object automatically entails 'enhanced processing' (which can be detected by measuring electrical activity in the brain) of other attributes of the same object, with task-irrelevant attributes of an attended object selected along with task relevant attributes (so if you ask someone to keep track of a thing's shape the processing devoted to its colour, size, etc. is also enhanced). Sameness of location on its own does not generate this kind of enhancement. This suggests that attention is being held by the object, rather than the location. See O'Craven, Downing, and Kanwisher (1999).

⁷ For references to automatic spread of attention experiments see Scholl (2002): 8-9. For a textbook account see Palmer (1999): 547-48.

⁸ See Driver et al. (2002): 76-79. For a textbook account of amodal completion results see Palmer (1999): 293-294, 554-555.

⁹ See for example Palmer (1999): 554.

¹⁰ For accounts of multiple object tracking experiments see Pylyshyn (2003): 223-27 and (2007): 34-58; Scholl (2002): 9-10 and (2007): 49-78.

¹¹ For this kind of argument for sortalism see, for example, Quine (1953): 67-68. Wiggins's (1997) claim that a subject can count as keeping track of a single object through motion and change only in virtue of keeping track of the object as a thing of some kind is a variant of this same 'pointing is ambiguous' line of thought adapted to the case of temporally extended pointing. Dummett (1973): 179, 233-234, 570-577, 582-583 and (1981): 218 and Strawson (1959): 180-213 both endorse two-stage (pointing indicates a general location; a sortal concept delineates the boundaries of an object in the location indicated) accounts of demonstrative reference-fixing.

¹² Note that to deny the traditional two-stage account of mental or physical pointing (pointing indicates a direction; sortal concepts draw boundaries around an object in that direction) is not to deny that sortal concepts may sometimes have a role to play in achieving determinate pointing. For it may be that sortal concepts sometimes play a role in orienting attention towards one rather than another visual object. Both the arm and the whole chair are there to attend to. One rather than another may be rendered salient by some cue (as in the automatic spread of attention and multiple object tracking experiments). If there is no such cue, a sortal concept might be needed to draw attention to a specific visual object. This suggestion is developed in Campbell (2002): 61-83.

¹³ This would be denied by philosophers who think that you can ‘say something about’ an object just by saying that it exists, and that existence is not a property. I take the argument for 2* at the end of this section to show that possible exceptions to 2 generated by this view do not undermine the argument for rich referential sortalism.

¹⁴ For canonical statements of the problem, and of this kind of solution, see Shoemaker (2003a) and (2003b).

¹⁵ This is Shoemaker’s term. See Shoemaker (2003a): 213.

¹⁶ This is Shoemaker’s view. See, for example, Shoemaker (2003a): 213.

¹⁷ For defence of this view see Campbell (2002): 235-54.

¹⁸ Frege (1950): 65-7. Wiggins raises a doubt about this claim (2001): 74.

¹⁹ Frege (1950): 65-7.

²⁰ Wiggins (2001): 21, 56.

²¹ Wiggins (2001): 70.

²² Wiggins (2001): 72, 77-106.

²³ See Grandy (2007) for a recent discussion of the various definitions of ‘sortal concept’ and the relations between them.

²⁴ The first part of this line of thought parallels Shoemaker’s claim about the need for a ‘baseline’ of stable behaviour against which the notion of endowment with a power may be explained. (Compare Shoemaker (2003a): 254: ‘In order to be able to specify the effects on a thing of various kinds of external influences, we need to be able to specify, as our baseline, the sort of history the thing would have if not subjected to these influences’.) But rather than looking to the notion of a characteristic pattern of behaviour to provide this baseline, Shoemaker claims that it is to be provided by a kind of essential inertia that he

supposes bundles of properties to have: ‘...it is typically built into the nature of a property that in the absence of external influences (or of a change in external influences) a thing having that property will continue to have it’ (Shoemaker (2003b): 254). But, though it is not possible to argue for this point properly here, Shoemaker’s notion of inertia among bundles of properties cannot provide the kind of ‘baseline’ that the notion of a conditional causal power requires. For an object’s properties just do not have the kind of stability that Shoemaker is supposing. I discuss this point in a little more detail in Dickie (2010): 527-528.

²⁵ A full discussion would bring out the connections between the sortalist proposal I have advanced in this section, the role of sortal classification in enabling us to reidentify objects after breaks in observation, and the developmental psychology literature on acquisition of the capacity to reidentify (see Xu (2007) for a summary). But it is not practicable to develop these connections here.

²⁶ For examples like this in the anti-sortalist literature, see Ayers (1974), (1997); Campbell (2006).

²⁷ See, for example, Wiggins (2001): 218; Evans (1982): 178.

²⁸ For example, here is Wiggins: ‘...for a thinker to single out or individuate a substance, [x] there needs to be something about his *rapport* with *x* or his relational state towards *x* and his practical sensibility in relation to *x*, which ... sufficiently approximates to this: the thinker’s singling out *x* as *x and* as a thing of a kind *f* such that membership in *f* entails some correct answer to the question “what is *x*?”’ (Wiggins (2001): 7). And compare Evans (1982): 178-179.

²⁹ Compare Campbell (2002): 72-73, (2006): 221-26.

* Thanks to John Campbell, Susanna Siegel, and an anonymous referee for this journal for comments on earlier drafts. Thanks to Hakwan Lau for discussion of the empirical material in §2. Thanks also to the Philosophy Department at New York University where this paper was completed during my time as a Bersoff Faculty Research Fellow.