

The Generality of Particular Thought

Suppose that you have the capacity to think that Jack is asleep. Then it is intuitive that you must also have the capacities to think a range of thoughts about Jack – that Jack is hungry; Jack is happy; Jack is awake; and so on. It is also intuitive that you must have the capacity to think a range of thoughts about being asleep – that Jill is asleep; that *he* (pointing at a man) is asleep; and so on. So it seems (intuitively) that we have the following¹:

Necessarily, if a subject, *S*, can think that *a* is *F*, *S* has the capacity to think, for a reasonable range of properties, Φ , that $\lceil a \text{ is } \Phi \rceil$.

Necessarily, if a subject, *S*, can think that *a* is *F*, *S* has the capacity to think, for a reasonable range of objects, α , that $\lceil \alpha \text{ is } F \rceil$.

Together, these claims entail that, necessarily, the thought that *a* is *F* lies at the intersection of two series of thoughts. One is the series of thoughts that *a* is *F*, that *a* is *G*, that *a* is *H*, and so on, for some reasonable range of properties. The other is the series of thoughts that *a* is *F*, that *b* is *F*, that *c* is *F*, and so on, for some reasonable range of

¹ '*a*', '*b*',... are object language names; '*F*', '*G*',... are object language predicates; '*F*', '*G*' are names for the properties these predicates stand for; ' α ' and ' Φ ' are variables ranging over object language names and predicates respectively; ' Φ ' is a variable ranging over properties; 'The thought that $\lceil a \text{ is } \Phi \rceil$ ' and 'The thought that $\lceil \alpha \text{ is } F \rceil$ ' abbreviate "The thought expressed by the sentence got by inserting '*a*' into the argument place of the value of Φ " and "The thought expressed by the sentence got by inserting the value of α into the argument place of '*F*'" respectively.

objects. I shall call the first claim (according to which the thought that a is F is part of a series of thoughts extending along the ‘property’ axis) ‘Property Axis Generality’. I shall call the second claim (according to which the thought that a is F is part of a series of thoughts extending along the ‘object’ axis) ‘Object Axis Generality’. This paper is about why these intuitive claims about the generality of particular thought should be accepted.

To guard against misunderstanding, it is important to distinguish the generality claims I am concerned with here from two nearby, but much weaker, generality claims.

The first of these weaker claims is an empirical claim about the actual structure of our cognitive capacities: the claim that, as a matter of empirical fact, an ordinary adult human who can think that ‘ α is Φ ’ can also think a range of other α -thoughts and a range of other Φ -thoughts. This empirical claim is not my concern here. Property Axis Generality and Object Axis Generality are claims about the nature of all thought for all subjects. They cannot be defended by appeal to empirical evidence about how our own cognition works.

The second weaker claim is a claim about how to explain a central fact about human cognition. The central fact is that there are an infinite number of thoughts we are able to grasp, even though our minds have only finite processing power. Philosopher’s favourite explanation for this fact is that our finite processing power endows us with the capacity to grasp a finite number of thought components, but these components can be combined into an infinite number of thoughts². It is standard, as part of this explanation, to claim that thinking that a is F involves the combined exercise of two capacities, one

² For a canonical statement of the parallel to this line of thought for the case of language learnability and understanding see the first few pages of D. Davidson, ‘Truth and Meaning’, in his *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press, 1984).

redeployable in thinking that a is G , that a is H , and so on, the other redeployable in thinking that b is F , that c is F , and so on. But, again, this (widespread) weaker generality claim does nothing to establish the (comparatively seldom discussed) stronger ones.

To bring out exactly why not, suppose that we are dealing with a subject whose repertoire of thoughts is much smaller than our own. This subject can think about only five objects and ten properties, and can form only logically simple thoughts (thoughts to the effect that an object from among the five has a property from among the ten). This subject can think only a finite number of thoughts. So, for the case of this subject, the standard motivation for saying that thinking that a is F involves the exercise of separately deployable capacities does not arise. But according to the strong generality claims which are my concern in this paper, even this subject's thought that a is F must lie at the intersection of ranges of thoughts extending along the object and property axes, for this kind of generality is essential to all thought about particular things. This paper is about what the grounds for this essentiality claim, might be.

In outline, here is what I want to suggest.

Philosophers who have thought about the generality claims³ have assumed that if you can argue for them at all the argument will have its source in some basic truth about the nature of thought about particulars. So the assumption has been that if you want to argue for the generality claims you need to find some such basic truth – that thoughts about objects are essentially structured; that thoughts about objects have, essentially, a given range of inferential properties – and show that the generality claims follow.

³ See, for example, G. Evans, *The Varieties of Reference* (hereafter *VOR*) (Oxford: Oxford University Press, 1982); C. Peacocke, *A Study of Concepts* (Cambridge Mass.: MIT Press, 1992); J. Campbell, 'Conceptual Structure', in C. Travis (ed.), *Meaning and Interpretation* (Oxford and New York: Blackwell, 1986).

As a result of this initial assumption, philosophers endorsing the generality claims have made three further assumptions about them. The first assumption is that the claims belong together: Property Axis Generality and Object Axis Generality follow from the same basic underlying principle about the nature of thought about particulars, so each is true iff the other is. The second assumption is that the generality claims apply equally to thought about every category of object we can think about: because the generality claims arise from something shared by all thoughts about objects, our thoughts about ordinary objects, numbers, other abstract objects, subjects of consciousness, and the objects posited by scientific theories are all ‘general’ in the same way. The third assumption is that the ‘reasonable’ ranges of properties and objects in the statements of the generality claims are very wide ranges.⁴ If the generality claims arise from something foundational about the nature of thought, it is natural to think that the generality at issue reaches all the way to the edges of thought, so that if I can think that a is F , and $\lceil a \text{ is } \Phi \rceil$ and $\lceil \alpha \text{ is } F \rceil$ are thinkable, the only possible barriers to my being able to think these thoughts are ignorance (I might not know what is involved in being Φ or that there is such a thing as α) or pathology (some thought I am capable of forming might be too horrible for me to contemplate).

Given these three assumptions, the intuitive generality claims I started with are combined and strengthened into what Gareth Evans calls the ‘Generality Constraint’:

⁴ Both Evans and other philosophers who have endorsed his Generality Constraint take ‘thinkability’ to be a matter of categorical appropriateness, so that there is no demand that, for example, I be able to think that Jack is a multiple of seven or that eight is hungry. E. Camp, ‘The Generality Constraint, Nonsense, and Categorical Restrictions’, *Philosophical Quarterly*, 54 (2004), pp. 209-231, begins by endorsing the Generality Constraint, but argues that it should not be restricted in this way. Since my concern here is with whether we should accept anything like the Generality Constraint in the first place, I shall not discuss her argument here.

If S can think that a is F , then, for every α and Φ for which $\ulcorner \alpha \text{ is } F \urcorner$ and $\ulcorner a \text{ is } \Phi \urcorner$ are thinkable, if S can think about α , S can think $\ulcorner \alpha \text{ is } F \urcorner$, and if S can think about Φ , S can think $\ulcorner a \text{ is } \Phi \urcorner$ ⁵.

So the assumption that the way to argue for the generality of particular thought is to start with foundational claims about the nature of thought about particulars leaves Property Axis Generality and Object Axis Generality looking like artificially separated and artificially weak shadows of a much stronger claim. Given this assumption, it seems that if you can argue for any generality claim at all it is going to be Evans's Generality Constraint. So, given the assumption, it seems that there is no point in considering Property Axis Generality and Object Axis Generality in their own rights.

But I think it is fair to say that the search for an argument from a foundational claim about the nature of thought to the Generality Constraint has not been a happy one. And, in the absence of a satisfactory argument of this kind, I think it is at least worth revisiting Property Axis Generality and Object Axis Generality and asking how arguments for these less ambitious conclusions might go.

The rest of the paper has four parts. §1 sets out the two main arguments for the Generality Constraint from foundational claims about the nature of thought about objects which are currently on the table, and shows that there are considerable difficulties which would have to be overcome before either could be made to work. §2 presents an

⁵ *VOR*, §4.3. Evans's actual statement of the constraint (104) mentions only Property Axis Generality, but he clearly means to include Object Axis Generality too – see note 21.

argument for Property Axis Generality for the case of thought about ordinary material things. §3 argues for a weakened and restricted form of Object Axis Generality. §4 makes some preliminary remarks about the importance of the shift in conception of the generality of particular thought that I want to propose.

§1 The arguments for the Generality Constraint

The first argument for the Generality Constraint that I shall consider is Evans's argument in *The Varieties of Reference*.

Evans argues like this. The capacity to think that a is F is essentially structured – thinking that a is F involves, essentially, the joint exercise of two distinct capacities: the capacity to think about a and the capacity to think about F ⁶. But if these capacities are genuinely distinct they must be separately deployable: each must be deployable in a range of thoughts which does not involve exercise of the other⁷. So when you think that a is F you are deploying a capacity (the capacity to think about a) which you could also deploy in thinking that a is G ; that a is H ; and so on. And you are deploying a capacity (the capacity to think about F) which you could deploy in thinking that b is F ; that c is F ; and so on.⁸

I shall raise four objections to this argument. I should emphasize that I do not take any one of these objections on its own, or all four together, to show that no argument for the Generality Constraint along Evans's lines could be made to work in the end. My aim

⁶ *VOR*, p. 100: 'It seems to me that there must be a sense in which thoughts are structured'; *VOR*, p. 102: 'But it simply is not a possibility for the thought that a is F to be unstructured – that is, not to be the exercise of two distinct abilities'.

⁷ See *VOR*, p. 102 on the 'commitments' we incur by claiming that the thought that a is F is structured.

⁸ *VOR*, p. 103: 'Someone who thinks that John is happy must, we might say, have the idea of a *happy man* – a situation instantiated in the case of John (he thinks), but in no way tied to John for its instantiation.'; see also p. 104 note 21.

here is just to show that Evans's argument as it stands falls far short of establishing the Generality Constraint, and, therefore, that it is at least worth considering what other support for the intuitive generality claims I started with might be available.

The first objection to Evans's argument is that we need a reason to accept the claim that thought about objects is essentially structured. (And remember that we are looking for a reason to think that all thought about objects is structured, not just that if I can think an infinite number of thoughts about objects my thoughts about objects must be structured.) Thoughts about objects do seem to have a kind of complexity that thoughts not about objects lack. For example, consider the contrast between thinking that Jack is asleep and thinking that it is misty. The truth of the first thought requires that an object (Jack) have a property (the property of being asleep). The truth of the second requires only that a condition (the condition of being misty) be met in some region of the world: it does not require that there be any object in that region of the world meeting the condition. So the truth conditions of the thought that *a* is *F* have a two-component structure not found in the truth conditions of mere 'feature placing' thoughts (like the thought that it is misty over there). But there is a long way to go from this claim to the claim that thinking that *a* is *F* involves the exercise of distinct capacities. For we still need a reason to think that grasp of structurally complex truth conditions is itself structurally complex.

The second objection to Evans's argument is that even if we accept that the thought that *a* is *F* is essentially structured (and so involves the exercise of distinct capacities) it is not clear why distinct capacities have to be capacities which can be deployed separately by a subject who has them. There is a hard question about how capacities should be individuated. But one plausible-seeming possibility is that they

should be individuated by what they are capacities to do. Given this way of individuating capacities, there is no reason why a subject could not have distinct capacities which share their activation conditions (so that S's capacity to do X is activated iff S's capacity to do Y is, even though the capacities are distinct). For example, I might be set up in such a way that my capacity to regard another person as a moral agent is activated iff my capacity to feel empathy for that person is. And this link between the activation conditions for my two capacities might be robust across physically possible worlds, so that in every physically possible world I regard as moral agents all and only those people with whom I have empathy. In this case, my capacities to regard people as moral agents and to feel empathy would not be separately deployable. But it would not follow that the capacities were the same. Similarly, I might be set up in such a way that I cannot think of Tahiti without thinking that it is sunny there, and cannot think of any example of a sunny place except Tahiti. And this connection might be modally robust. But if capacities are individuated by what they are capacities to do, this would not undermine the claim that in thinking that Tahiti is sunny I am exercising two distinct capacities.

Of course, a defender of Evans might reply that capacities should be individuated in such a way that distinct capacities are separately deployable. But if this claim is to be used to fill out Evans's argument we need a reason to accept it other than just that it is required to make the argument work. I do not suggest that no such reason could be found. I do suggest that if there is such a reason it will be hard to find and, therefore, that Evans's argument as it stands does not get past this second objection.

The third objection to Evans's argument is that even if we accept that a subject who can think that *a* is *F* can deploy the capacity to think about *a* and the capacity to

think about F separately, the Generality Constraint does not follow. For the range of thoughts in which the subject can deploy the capacity to think about a might be the thoughts that $a = a$; $a \neq b$; $a \neq c$; and so on. And the range of thoughts in which the subject can deploy the capacity to think about F might be the thoughts that being F entails being G ; that being F entails being either H or I ; and so on. So the demand for separate deployability does not entail the demand for deployability in the ranges of thoughts that the Generality Constraint lays down.

The fourth objection is that, even if the other objections are set aside, Evans has given us no reason to think that the capacities involved in thinking that a is F are redeployable to the extent that the Generality Constraint specifies. The Generality Constraint says that if I can think that a is F , and $\lceil \alpha \text{ is } F \rceil$ and $\lceil a \text{ is } \Phi \rceil$ are thinkable, only ignorance or pathology (or failure to get around to it) can stand in the way of my thinking them. But you might allow that the capacities involved in thinking that a is F are redeployable along the object and property axes without thinking that their redeployability must extend this far.

Again, I do not contend that these objections are unanswerable. I do suggest that the objections show that, as it stands, Evans's argument for the Generality Constraint falls a long way short of establishing its conclusion.

The second argument for the Generality Constraint that I shall consider is what I shall call the 'Inferential Argument'⁹. The Inferential Argument takes off from the claim that to grasp a thought you must grasp some range of its inferential relations to other thoughts: for each thought, there is a range of inference patterns involving the thought

⁹ This kind of argument is proposed by John Campbell in 'Conceptual Structure', and *Past, Space, and Self* (Cambridge, MA: MIT Press, 1994).

such that if you do not recognise the correctness of inferences conforming to these patterns you do not grasp the thought. For the thought that a is F , this range of inference patterns will include inferences of these forms (this is not intended to be an exhaustive list):

$\frac{a \text{ is } F}{a = b}$ $b \text{ is } F$	$\frac{a \text{ is } F}{a \text{ is } G}$ $\text{Something is both } F \text{ and } G.$	$\frac{a \text{ is } F}{\text{Something is } F}$	$\frac{\text{Everything is } F}{a \text{ is } F}$
---------------------------------------------------	-----------------------------------------------------------------------------------------	--------------------------------------------------	---------------------------------------------------

The proponent of the inferential argument then claims that we are able to make inferences of these kinds only because our thinking meets the Generality Constraint. For example, you are able to make the inference in the first column only because (given the Generality Constraint) the fact that you understand the first premiss (a is F), and are able to think about b , guarantees that you are also capable of understanding the conclusion (b is F).¹⁰ And you are entitled to make the inference in the second column only if you are entitled to take it that the occurrences of a in the two premisses are tokens of the same name. This entitlement cannot rest on your knowledge of an extra hidden identity premiss ($a = a$). For to point to a hidden identity premiss as the source of the entitlement is just to invite the question ‘What is the source of your entitlement to think that either occurrence

¹⁰ See Campbell, ‘Conceptual Structure’, p. 166. Campbell takes this kind of argument to show only that grasp of the truth conditions of the thought that a is F must satisfy the ‘permutation’ constraint (the object-axis part of the Generality Constraint). His concern in the paper is actually to establish a different conclusion – that grasp of the truth conditions of the thought that a is F exploits grasp of the thought’s conceptual structure. The Generality Constraint leaves open the possibility that grasp of the structure of a thought is derivative from, rather than exploited by, grasp of the thought’s truth conditions (see p. 169). He has a separate argument (pp. 169-172) for the conclusion that grasp of thoughts about spatio-temporal particulars exploits grasp of these thoughts’ conceptual structure, and so satisfies the Generality Constraint. But as far as I can see this argument establishes only that grasp of, for example, the thought that a is F involves recognising that the state of affairs which makes this thought true (at a time) is repeatable and subsumable under laws. There is still a gap between this conclusion and the claim that grasp of the thought must satisfy the Generality Constraint.

of ‘ a ’ in ‘ $a = a$ ’ is an occurrence of the same term as occurs in either premiss?’. Rather, it seems that your entitlement to take the occurrences of ‘ a ’ to be tokens of the same term must rest on the fact that grasp of the ‘ a ’-part of each premiss involves exercise of the same capacity – a capacity that is exercised in thinking that a is F when you grasp the first premiss, and that a is G when you grasp the second.¹¹ So the suggestion is that to grasp a thought you must grasp its basic inferential connections with other thoughts, and your grasp of ‘ a is F ’ puts you in a position to grasp the required inferential connections iff your thinking meets the Generality Constraint.

Again, I do not want to suggest that no version of the Inferential Argument could be made to work. But as it stands the argument does not establish its conclusion.

To see why not, look first at the inference in the first column, and ask yourself what the process of grasping this inference’s validity is like. You start with the first premiss. You move on to the second. Then you look at the conclusion, and ask whether accepting the truth of the premisses would compel you to move to the conclusion. This process requires that you understand the conclusion by the time you get around to considering it. But it does not require that you already understand the conclusion back at the beginning of the process, when you are considering the first premiss. It may be that your understanding of the conclusion is generated by your understanding of the two premisses combined: in making the inference you are supposing that a is F and that $a = b$; these two suppositions generate an account of what it would be for b to be F . The process of grasping the inference’s validity does not even require that you already have the capacity to grasp the conclusion back when you are thinking about the first premiss. It may be that the capacity to grasp the conclusion is also generated by your understanding

¹¹ Compare Campbell, *Past, Space, and Self*, p. 203.

of both ' a is F ' and ' $a = b$ ': you are in a position to understand ' b is F ' because you already understand ' a is F ' and ' $a = b$ '. But the Generality Constraint says that if you understand ' a is F ', and can think about b , that alone is enough to give you the capacity to understand ' b is F ', without any extra ingredient that the supposition that $a = b$ might supply.

Now consider the inference in the second column. The proponent of the inferential argument is surely right to say that my entitlement to take occurrences of names occurring in different premisses of an inference to be tokens of the same name cannot always rest on grasp of a hidden identity premiss: if it did, there would be no end to the hidden identity premisses we would have to suppose that inferences like that in the second column contained. And the only plausible account of where else this entitlement could come from does seem to be the account the proponent of this argument wants: if I am to be in a position to move from ' a is F ' and ' a is G ' to 'Something is both F and G ', my grasp of the ' a '-part of each premiss must involve exercise of the same capacity in each case. So it does seem that if I am to be licensed in making this kind of inference, the capacity to think about a that I am exercising when I grasp the thought that a is F must be redeployable in thinking that a is G . But this line of thought will establish the Generality Constraint only given the suppositions that (i) if I can think that a is F , I can think that a is G , and (ii) if I can think these thoughts I am entitled to make the move to 'Something is both F and G '. And to suppose (i) is just to suppose a restricted version of the Property Axis of the Generality Constraint, which is what the argument is trying to establish. It will not do to object just that the capacity to make this kind of inference is essential to the capacity to think that a is F . There are other inference patterns on the table as potential

anchors for the claim to be grasping the thought that a is F – for example, the inferences in the other three columns. (Note that applying the line of thought about entitlement to treat distinct occurrences of ‘ a ’ as occurrences of the same name to the inference in the left hand column gets us only that in thinking that a is F I am exercising a capacity to think about a which can be redeployed in thinking that $a = b$.) So we would need a reason to think that this particular pattern of inference is essential to all thought about particulars. And it is hard to see how you could establish this essentiality claim without first establishing that if I can think that a is F I can think that a is G – that is, without first establishing the Property Axis of the Generality Constraint.

Now consider the inference in the third column. The process of grasping the validity of ‘ a is F therefore Something is F ’ requires that if you grasp ‘ a is F ’ you grasp ‘Something is F ’. But, again, this requirement falls short of the Generality Constraint. For you might understand both ‘ a is F ’ and ‘Something is F ’ without understanding ‘ α is F ’ for any name α distinct from ‘ a ’. Here is an example to bring out this possibility.

Suppose that a censorship committee is set up to consider which books are to be banned for their unsuitable contents. The committee works through books of various genres. And part of what has to be decided is what counts as ‘unsuitable’ within a genre. So, for example, at the time at which it has been decided that *Middlemarch* is unsuitable it is still up for grabs what would count as unsuitability for the *Grundlagen*¹². Then there might be a time at which there is only one book for which standards of unsuitability have been set. Let us say that this book is in fact *Middlemarch*. In that case, speakers will be in a

¹² J. Tappenden uses examples like this to argue that natural language negation is wide-scope negation at ‘Negation and Language Change in Philosophical Logic’, in D. Gabbay and H. Wansing (eds.), *What is Negation?* (Kluwer Academic Press, 1999), §4.

position to grasp both ‘*Middlemarch* is unsuitable’ and ‘Something is unsuitable’ without grasping ‘ α is unsuitable’ for any α other than ‘*Middlemarch*’.

Finally, consider the inference in the fourth column. You will be able to grasp the validity of ‘Everything is F therefore a is F ’ only if the fact that you understand ‘Everything is F ’ endows you with the capacity to understand ‘ a is F ’. But this provides no reason to think that if you understand ‘ a is F ’ and understand ‘ b ’ you must have the capacity to understand ‘ b is F ’, which is what the Generality Constraint requires.

At this point a defender of the Inferential Argument might object that we will find inferential patterns which vindicate the argument by expanding the range of inferential behaviour regarded as essential to grasp of ‘ a is F ’. In particular, you might think that we should be looking at both the introduction and the elimination rules for the quantifiers (rather than just at Existential Introduction (the pattern second from the right) and Universal Elimination (the rightmost pattern)). Recall that the Existential Elimination rule says

$\exists x\Psi x$	$X, [\Psi b]$	Given ‘ $\exists x\Psi x$ ’, and a derivation of P from the assumption ‘ Ψb ’ (for arbitrary ‘ b ’) and auxiliary premisses X , discharge the assumption and move to P from X and ‘ $\exists x\Psi x$ ’.
	⋮	
	⋮	
	P	
P		

Now suppose that a subject who grasps ‘ a is F ’ must recognise the validity of basic inferences involving this sentence and both the introduction and the elimination rules for the existential quantifier. This looks like a reasonable supposition: if I am in the business of talking about particulars at all, I must be in the business of generalising about them,

6 $\exists x(x \text{ is asleep} \ \& \ x \text{ is a dog})$. [1, 3, 5 by \exists Elimination; assumptions at 1 and 2 carried forward]

This argument purports to derive the conclusion that some dog is asleep from the premisses that something is asleep and Jack is a dog. It is invalid. The problem is that the move from 5 to 6 (which takes us from the claim that ‘ $\exists x(x \text{ is asleep and } x \text{ is a dog})$ ’ follows from 2 and 3 to the claim that this conclusion follows from 1 and 2) is legitimate iff the assumption at 3 brings no more information into the proof than the assumption at 1 does. And given that the name ‘Jack’ occurs in line 2, this condition does not hold. Given line 2, the proof already contains the information that Jack is a dog. So at 3 we are not supposing that some object, about which we know nothing else, is asleep. We are supposing that an object which we already know to be a dog is asleep: given 2, the assumption at 3 has consequences that the assumption at 1 does not, so the move from 5 to 6 is illegitimate. The stipulation that the assumption in the sub-proof of \exists -elimination must involve an ‘arbitrary’ object (P must be derived from X and assumption ‘*b is F*’ for arbitrary ‘*b*’) is there to isolate the assumption from any information about specific objects that the proof might contain¹³.

And given that this is the role that the stipulation about arbitrariness is playing, it would just be a mistake to think that deployment of the assumption that an arbitrary object is *F* in an instance of \exists -elimination requires grasp, for each object in the domain,

¹³ Natural Deduction systems differ with respect to whether they use \exists -elimination (‘Fitch systems’) or \exists -instantiation (‘From $\exists x(\dots x \dots)$ infer $(\dots \alpha \dots)$, with restrictions on α ’) (‘Quine systems’). They also differ with respect to whether they achieve the ‘arbitrariness’ of the assumed or inferred instance by using free variables (with restrictions on which variables are eligible), ordinary names (with restrictions on when a name may be used), or a special class of ‘arbitrary’ or ‘ambiguous’ names. See F. J. Pelletier, ‘A History of Natural Deduction and Elementary Logic Textbooks’, in J. Woods and B. Brown (eds.), *Logical Consequence: Rival Approaches vol. 1* (Oxford: Hermes Science Publishers, 2000), for a summary of how the main texts on Natural Deduction divide up with respect to these and other choice points. The point I have made about the role of the claim that some arbitrary object is *F* would be stated slightly differently for different kinds of system. The differences are not important for the purposes of this paper.

of what it would be for that object to be F . The force of the assumption is not (as someone making this mistake is thinking) “Let us suppose that some arbitrary one of ‘ a is F ’; ‘ b is F ’; ‘ c is F ’.... is true”. Rather, it is ‘Let us suppose that some object is F and make no other specific assumptions about it’. And to make this assumption you do not need to know what it would be for each of ‘ a is F ’; ‘ b is F ’; ‘ c is F ’.... to be true.

To consolidate this point, consider again the case in which a committee has been set up to decide criteria of unsuitability, and this committee has got only as far as laying down conditions which determine that *Middlemarch* is unsuitable, but which it is completely unclear how to apply to anything else. So in these circumstances we do not yet know what would be required for ‘is unsuitable’ to be satisfied by anything other than *Middlemarch*. In this situation the following line of inference would still be perfectly intelligible: ‘*Middlemarch* is unsuitable. So something is unsuitable. But whatever is unsuitable is banned. So something is banned.’ Grasp of how ‘*Middlemarch* is unsuitable’ interacts with basic general claims about unsuitability does not require grasp of ‘ α is unsuitable’ for any α other than ‘*Middlemarch*’.

So the situation is this. There are two main arguments for the Generality Constraint currently on the table. Neither succeeds as it stands. It would be premature to suggest that neither can, in the end, be turned into an argument which survives the objections I have raised. But in each case coming up with a satisfying defence would be a long, hard job.

In the rest of the paper I want to explore the possibility of a different kind of argument for the generality of particular thought. The argument I want to propose does not begin with a supposedly foundational claim about the nature of all thought about

particulars. And it does not attempt to reach all the way to the Generality Constraint. Instead, it begins with a problem about how our thought about ordinary material objects works, and aims to show only that the solution to this problem generates Property Axis Generality and a form of Object Axis Generality for the case of thought about this specific subject matter. So the suggestion is that, for the case of thought about ordinary objects, the generality intuitions I started with can be traced to a reasonably basic claim about how thought about this subject matter works, and, in default of an argument for a more deeply rooted and more comprehensive generality claim, it is at least worth entertaining the possibility that this weaker kind of generality is all the generality that is essential to thought about ordinary material things.

The next two sections set out the arguments for Property Axis Generality and Object Axis Generality that I want to propose. The final part of the paper is about the significance of the move I am suggesting, away from Evans's Generality Constraint and towards these more humble generality claims, for the attempt to provide an account of how our thought about the world works.

§2 Property Axis Generality

Property Axis Generality says that if you can think that a is F you can think a range of thoughts about a extending along the property axis: the thoughts that a is G , that a is H , and so on for some reasonable range of ' G ', ' H ', In this section I present an argument for Property Axis Generality for the case of thought about ordinary material things. (From now on I shall leave the 'for the case of ordinary material things'

qualification implicit, and use ‘object’ or ‘particular’ to mean only ‘ordinary material thing’.)

The argument for Property Axis Generality that I want to present takes its rise from the old problem of what distinguishes ‘genuine’ properties, like being square, and weighing twenty pounds, from so-called ‘spurious’ or ‘mere Cambridge’ properties like having been slept in by George Washington, being green and looked at before midnight or blue and looked at afterwards, or being five miles south of a burning barn. In outline, what I want to suggest is this. The solution to the problem entails that our grasp of what it is for an object to have a genuine property comes in clumps: if being F is a genuine property and you understand what it is for an object to be F , you must also understand what it is for the object to be G or H or, ..., for some reasonable range of genuine properties. But if you have the capacity to think about an object at all you have the capacity to think about the object as having at least one genuine property. So if you have the capacity to think about a at all you have the capacity to think, for a range of genuine properties Φ , that ‘ a is Φ ’.

The old problem about how to distinguish ‘genuine’ from ‘spurious’ properties arises like this.

Suppose

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

This generous criterion entails that objects have many properties that we do not ordinarily consider. For example, it entails that 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, and being such that George W. Bush is president of the United States. 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 in which 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 questions 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 which hold in virtue of the objects which participate in the related events and these objects' properties. Events consist in changes in objects. 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 argued that only ordinary properties and ordinary change interact with the notion of causation in the appropriate way. So the suggestion is that, for example, the fact that X ceased to have the property of being such that George Bush is president is not relevant to whether or not X's trigger-pulling caused Y's death, and 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 and ordinary change 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

¹⁴ For canonical statements of the problem and of this kind of solution see S. Shoemaker, 'Causality and Properties' and 'Identity, Properties, and Causality', in his *Identity, Cause, and Mind (Expanded Edition)* (Oxford: Oxford University Press, 2003).

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 are to be made out. But I want to say just enough about the causal solution to the problem about the status of ordinary properties and ordinary change to show how the solution entails Property Axis Generality. And to bring this out we need only the following general claim:

The Causal Relevance Principle (hereafter ‘CRP’): Our grasp of properties as genuine is grasp of properties as determining the causal implications of possessing them.

Note that the CRP is an epistemic principle. Grasping is a kind of knowing. And knowledge is factive. So the CRP does have a metaphysical consequence – the consequence that the properties we take to be genuine determine the causal implications of possessing them. But it is not a metaphysical claim about what genuine properties are. In particular, it is neutral between the (extreme) metaphysical claim¹⁵ that the nature of a genuine property is exhausted by the causal implications of possessing it and the (more moderate) metaphysical claim that a genuine property is what generates (or provides the ‘categorical grounds’) for a cluster of causal implications¹⁶.

¹⁵ This is Shoemaker’s view. See, for example, ‘Causality and Properties’, p. 213.

¹⁶ For defence of this view see Campbell, *Reference and Consciousness* (Oxford: Oxford University Press, 2002), ch. 12.

The CRP is also neutral with respect to exactly how our grasp of genuine properties as determining the causal implications of possessing them works. One possibility is that we grasp genuine properties as determining the causal implications of possessing them because if I grasp *F* as a genuine property I treat *F* as distinguished from other properties by its causal implications (so that, for example, knowing which property is the property of being square involves knowing enough of the causal implications of being square to distinguish squareness from the other genuine properties). Another possibility is that knowledge which property *F* is does not require knowledge of the causal implications of being *F*, but knowing what it is for an object to be *F* does involve knowledge of these implications. On this second kind of view, my knowledge which property is the property of being square does not involve grasp of the causal implications of being square – maybe it consists just in my capacity to recognise squares when presented with them. But I have to grasp at least some of the causal implications of being square to know what it is for an object to be square.

Because it is neutral in both of these respects, the CRP is actually quite a weak claim. But it is strong enough to establish the following step along the way to Property Axis Generality. Understanding of ascriptions of genuine properties comes in clusters: if you grasp *F* as a genuine property, and know what it would be for *a* to possess *F*, there is a range of other properties, *G*, *H*,... such that you grasp these properties as genuine and know what it would be for *a* to have these properties as well.

The reason the CRP has this result emerges as soon as we ask how knowledge of a property's causal implications is going to work. For the causal implications of being *F* for an object depend on the object's other properties. An object'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 object'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. How an object will behave if it has *F* is always conditional on the object's other properties.¹⁷ So the CRP can be expanded to

Grasp of *F* as a genuine property involves knowledge that an object which has *F* will behave in way w_1 iff it also has set of properties Π_1 ; in way w_2 iff it also has set of properties Π_2 ; and so on, for a range of sets of properties with which being *F* might be combined.

So knowing what it is for an object to have *F* involves knowing what it is for the object to have *G*, *H*, *I*, ..., and so on, where these are the properties in Π_1 , Π_2 , But the properties in Π_1 , Π_2 , and so on are also genuine properties – otherwise they would not be properties whose possession or not is relevant to how an object will behave if it has *F*. So the CRP entails that grasp of ascriptions of genuine properties comes in clusters: if I can think that *a* is *F*, where '*F*' stands for a property *F* which I grasp as genuine, then I can

¹⁷ Compare Shoemaker, 'Causality and Properties', pp. 212-213; 'Identity, Properties, and Causation', pp. 248-250, but note that Shoemaker is proposing a metaphysical account of what a genuine property is (he thinks that a genuine property is a cluster of conditional causal powers) while I am concerned with a parallel epistemological account – an account of how our knowledge what it is for an object to have a genuine property is structured.

also think that a is G , that a is H , that a is I , ..., where I grasp G, H, I, \dots as genuine properties too.

To get from this point to Property Axis Generality we need one final claim: the claim that if you can think of a as having any property you can think of it as having at least one genuine property. But if the distinction between genuine and spurious properties is allowed its intuitive bite in the first place, it carries this claim 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 you would have to mention as part of a complete account of what there is in the world: the genuine properties are the properties you would have to mention in giving such an account; the spurious properties (like being such that George Bush is president) are the ones which could be left out without losing any information. So to suppose that you could think of a without understanding what it would be for a to have any genuine property is to suppose that you could still count as thinking of a even though you had no grasp of any difference that what a was like could make to a right account of the way the world is. But to be committed to the existence of an ordinary object is to take it that what the object is like will make a difference to the way the world is¹⁸. And if you are thinking about a you are committed to a 's existence. So if you are in a position to think about ordinary object a at all, there must be at least one genuine property such that you can grasp what it would be for a to have this property.

Putting together the pieces we have now assembled, the argument for Property Axis Generality that I want to propose can be put like this:

¹⁸ You might deny the parallel claim for the case of mythical, fictional, or other abstract objects.

1 Grasp of properties as genuine involves grasp of properties as causally relevant. [The Causal Relevance Principle]

2 For all Φ , if subject S can think that $\lceil a \text{ is } \Phi \rceil$, where Φ stands for property Φ which S grasps as genuine, then for some $\Psi_1, \Psi_2, \Psi_3, \dots$ which stand for properties distinct from Φ and grasped by S as genuine, S can think that $\lceil a \text{ is } \Psi_1 \rceil, \lceil a \text{ is } \Psi_2 \rceil, \lceil a \text{ is } \Psi_3 \rceil, \dots$ [From 1].

3 If S can think about a , then for some Φ such that Φ stands for a property that S grasps as genuine, S can think that $\lceil a \text{ is } \Phi \rceil$. [From the explanatory role of the notion of a genuine property.]

4 If there is some Φ such that S can think that $\lceil a \text{ is } \Phi \rceil$, there is a range of $\Psi_1, \Psi_2, \Psi_3, \dots$ which stand for properties distinct from Φ such that S can think that $\lceil a \text{ is } \Psi_1 \rceil, \lceil a \text{ is } \Psi_2 \rceil, \lceil a \text{ is } \Psi_3 \rceil, \dots$ [From 2 and 3]

§3 Object Axis Generality

Object Axis Generality is the claim that if you have the capacity to think that a is F you also have the capacity to think that $\lceil \alpha \text{ is } F \rceil$, for some range of α such that $\lceil \alpha = a \rceil$

is false. The tendency to think of Object Axis Generality and Property Axis Generality as united into the single Generality Constraint has masked the fact that Object Axis Generality is actually much less intuitive than Property Axis Generality. It is hard to see how you could think of an object as having (and able to have) only one property: if you think of it as round it seems that you must be recognising that it is located somewhere and its boundaries are filled in in some way; if you think of it as purple it seems that you must be allowing that it is a thing of some shape. This is one source of the intuitive plausibility of Property Axis Generality. But the possibility of thinking of a property as possessed (and possessable) by only one object (even though you can think of other objects as possessing other properties) seems less problematic. The thoughts that God is maximally perfect and that 2 is an even prime number seem to provide examples of thoughts of this kind.

In this section I extend the line of argument about our grasp of possession of properties from §2 to establish a restricted form of Object Axis Generality. The restriction vindicates the suspicion that Object Axis Generality is somehow on shakier ground than Property Axis Generality to the following extent. The anchor for Property Axis Generality that we have found will support only a very restricted version of Object Axis Generality. So anybody wanting to uphold Object Axis Generality in a stronger form still needs to find a reason to do so.

The argument for Property Axis Generality developed in the previous section turns on the claim that the CRP entails that grasp of ascriptions of genuine properties comes in clusters. Now I am going to argue that the CRP also entails that thinking that a is F involves recognition that a is a thing of some or other kind, where to recognise that a

is a thing of a 'kind' is to recognise that *a* instantiates a combination of properties at a time and over time which is characteristic of some kind of thing. The suggestion will then be that a restricted version of Object Axis Generality is true just because grasp of a kind is grasp of something 'general' (in a sense to be explained).

Consider some statements of conditional causal power:

(1) A knife-shaped object has the power of cutting wood at room temperature if knife-sized and made of steel.

(2) A knife-shaped object has the power of cutting butter at room temperature if knife-sized and made of plastic.

These statements of conditional causal power make appeal to kind concepts – the kinds of stuffs that the knife-shaped object is being credited with the (conditional) powers to cut, and the kinds of stuffs of which the object itself is being said to be made. And these kind concepts are performing an ineliminable role. For compare 1 and 2 with

(3) A knife-shaped object has the power of cutting soft substances if knife-sized and made of a sufficiently harder substance.

3 is not a satisfactory statement of conditional causal power because it makes it too hard to be knife-shaped. Consider substance X (let us not worry about what it is) which is soft at room temperature but hardens instantly on contact with plastic or wood or steel. Then

no knife-sized thing made of plastic, wood, or steel has the power to cut X. So, given 3, no knife-sized plastic, wood, or steel thing is knife-shaped. To avoid this kind of problem we need to specify that the soft substances in question do not harden on contact with plastic, steel, or wood; do not harden or melt away when you start to cut them; are not disposed to switch backwards and forwards between being hard and being soft; and so on. But to provide these qualifications just is to provide a more or less specific account of the kind of substance you are dealing with. A thing's conditional causal powers are powers to act on other things or substances of specific kinds.

And the line of thought which brings out the ineliminability, in a statement of conditional causal powers, of reference to the kinds of things on which the powers are powers to act also brings out the eliminability of reference to the kinds of things which have powers to do the acting. For consider an object of kind Y (again, let us not worry about what this kind is) which is made in such a way as to be usable for cutting butter only once – on the second contact with butter it will self-destruct (in the interests of hygiene). This example shows that even 1 and 2 need amendment. We need to say that a knife-sized thing has the power to cut butter if it is knife-shaped and of a kind whose shape, size, composition, and so on remain stable on contact with butter. So we need something like

A knife-shaped object has the power to cut butter if it is knife-sized, made of plastic, and an ordinary plastic artefact.

A statement of conditional causal power must contain information about the kind of the thing that has it.

Stepping back from the examples, the point I am urging about the connection between conditional causal powers and kinds can be put like this. There is an old distinction between the notion of an ‘attribute’ (a thing’s attributes include its shape, colour, texture, and so on) and the notion of a ‘substance’ (tables, dogs, and boulders are all substances of different kinds). Both our grasp of attributes and our grasp of substances involves grasp of associated conditional causal powers. But a thing’s causal behaviour in a situation always depends on both its attributes and what kind of substance it is. So taking a thing to have a conditional causal power involves recognising it as a thing of a kind.

I want to consolidate this account of the connection between attaching a conditional causal power to a thing and recognising that it falls under a kind by comparing what I have said with a key transition in one of the canonical discussions of the notion of a conditional causal power: Shoemaker’s discussion in his ‘Identity, Properties, and Causality’. Shoemaker notes that the notion of a conditional causal power has to be defined relative to ways that things generally behave. If we cannot take for granted the ways A and B would behave if they did not interact, we cannot make sense of the claim that their properties give them powers to bring about various changes in one another: grasp of the notion of the power to bring about change presupposes grasp of how things will go if the power is not exercised. But Shoemaker holds back from the appeal to grasp of kinds that I have suggested must be made at this point. Instead, he says that the work of securing the stability of an object’s properties (so that conditional causal powers

can be defined as powers to interfere with, or undergo alternation in, what would have unfolded otherwise) is to be done by the notion of ‘immanent causation’: the tendency of an object’s properties to perpetuate themselves in the absence of change in external influence. So the suggestion is that an object’s properties will be stable in the absence of change in external influence, and that conditional causal powers are then to be thought of in terms of powers to interfere with bundles of properties which are, of their nature, otherwise stable.¹⁹

But Shoemaker’s notion of immanent causation cannot do the job he wants it to. For an object’s properties just do not have the tendencies to perpetuate themselves in the absence of change of external influence that he is supposing. Shoemaker acknowledges that the claim that an object’s properties have a tendency to perpetuate themselves given stable external influences needs qualification to allow for phenomena like biological growth and radioactive decay, which involve a thing’s properties changing without change in external influence.²⁰ But exceptions to the claim that an object’s properties will perpetuate themselves without change in external influence are even more widespread than Shoemaker is supposing. Here are some examples:

An object subjected to a constant force moves faster and faster – its speed changes even though the relevant external influence (the force being applied to it) remains the same.

A hot object left outside on a cold day cools while the relevant external influence (the temperature of air with which the thing is in contact) remains unchanged.

¹⁹ ‘Identity, Properties, and Causality’, pp. 253-254.

²⁰ ‘Identity, Properties, and Causality’, p. 254.

The wind and weather blow on the stone tower in the same unspectacular but insistent way for 500 years. For 499 years its large-scale properties remain unchanged. In the 500th it begins to crumble.

The ball of snow rolls at a constant speed down the uniformly angled and uniformly snow-covered slope getting larger and larger as it goes.

It is just false that a thing's properties have the tendency to perpetuate themselves given stability in the external environment. So the stability or predictability in the properties of objects that is needed to get the notion of a conditional causal power off the ground cannot be provided by the claim that bundles of properties will perpetuate themselves unless external influences change. What I have suggested instead is that things' properties have tendencies to develop according to set patterns given steady external influences, and that these patterns provide the stability that the notion of a conditional causal power requires.

And with the connection between conditional causal powers and kind concepts in place, the diagnosis I want to make with respect to Object Axis Generality can be put like this. Suppose that S can think that a is F . Then S can think that ' a is Φ ' for some causally relevant property Φ (otherwise there is no anchor for the claim that S is really thinking of a at all). But to think that a has a causally relevant property you have to recognise that a is a member of a kind (though you do not need to know what kind this is). So, for example, thinking that a is square is thinking that, for some kind κ , a is a κ

and a is square (so that the causal implications of a 's squareness are defined against the background of its κ -hood). And the presence of a kind concept in your grasp of what it is for an object to have a genuine property opens up the potential for generalisation along the object axis. For the causal implications of being square will be the same for anything of a 's kind as they are for a . So in grasping this set of causal implications you are grasping what it would be for any κ to be square: in grasping what it is for a to be square, you set yourself up for thinking that anything of a 's kind is square. So my grasp of what it is for a to be square is, essentially, generalisable along the object axis. But I might recognise that a is a κ without conceiving of the possibility that any other thing is a κ . (Compare 'God is a prime mover unmoved.') So we have yet to see any reason to think that my thought that a is F actually extends along the object axis: we have yet to see any reason to think that if I can think that a is F there is a range of objects, α , such that I can think that $\ulcorner \alpha \text{ is } F \urcorner$. If Object Axis Generality is to be defended in the form initially stated, the defence lies elsewhere.

§4 Consequences

So far in this paper I have argued for a shift away from the orthodox view of how an argument for the generality of particular thought might go. According to the orthodox view, if there is such an argument to be had at all it will begin with some foundational claim which holds for thought about any domain of particulars, and takes us all the way to the Generality Constraint. There are three main respects in which the argument I have suggested departs from this standard template. Firstly, the argument does not begin with a claim about the nature of all thought about particulars. It begins with a claim about

thought about ordinary material objects – the claim that our ascriptions of properties to ordinary objects are causally loaded. Secondly, the argument treats Property Axis Generality and Object Axis Generality separately (while an argument for the Generality Constraint will seek to establish these two generality claims as a package). Finally, even when the versions of Property Axis Generality and Object Axis Generality that are extractable from the argument are put together, they do not add up to the Generality Constraint for the case of thought about ordinary material things. The Generality Constraint says that if I can think that a is F then for every object α and property Φ that I can think about, if $\ulcorner \alpha \text{ is } F \urcorner$ and $\ulcorner a \text{ is } \Phi \urcorner$ are thinkable, I can think them. The argument I have set out establishes only that my capacity to think that a is F must generalise along the object and property axes far enough to let me grasp the causal loadedness of the property ascription.

In this final section I shall give a very brief indication of some lines of inquiry that this proposed shift in conception of the source and extent of the generality of particular thought opens up.

One direction of inquiry concerns the question of whether there is a genuine asymmetry between thought about objects and thought about properties. Intuition cries out that there is such a difference. But it has proven difficult to say where exactly this difference lies. The line of argument of this paper suggests one way to flesh out the asymmetry claim for the case of thoughts about ordinary objects and their properties: grasp of properties is essentially general in a way that grasp of objects is not. If you can think that a is F you must be able to think that a is G , a is H , and so on, for a reasonable range of properties (a range big enough to anchor your grasp of the causal loadedness of

the initial property ascription). But you need not be able to think that *b* is *F*. You need only recognise that in thinking about *a* you are thinking about an instance of a kind.

A second direction of inquiry concerns sortalism about thought about ordinary things. ‘Sortalism’ is the view that the capacity to think about ordinary objects depends on the capacity to classify them according to their kinds. The current debate about sortalism takes place against the background assumption that whatever motivation there might be for sortalism will come from the claims that to think about *a* you must know which thing *a* is, and that knowledge—which always depends on the capacity to classify. For example, the suggestion might be that sortalism is true because if you can think about *a* you know which thing *a* is, and knowing which thing *a* is involves knowing which changes *a* could undergo and which would destroy it²¹. Or it might be that sortalism is true because to think about *a* you must know which thing *a* is; your knowledge—which cannot consist just in knowing that *a* is the thing at such-and-such spatio-temporal location because there can be two things (a statue and a lump of clay) in a place at a single time; and that the extra element required for knowledge—which is knowledge of kind.²² But this paper has opened up a very different possible argument for sortalism – an argument from the claims that to think about a thing you must know what it would be for the thing to have a genuine property, and that this knowledge is kind-concept-involving. If this is how sortalism is motivated, the main intuitive objection to sortalism – the objection from the claim that we seem to be able to lock onto objects in perception as potential objects of thought while lacking any grasp of where they might fit in a system

²¹ Insofar as D. Wiggins has an argument for sortalism it seems to be an argument of this kind. See for example, *Sameness and Substance Renewed* (Cambridge: Cambridge University Press, 2001), ch. 2 §2.

²² Compare Evans, *VOR*, p. 107.

of sortal classification²³ – falls away. For, given the new motivation, you could allow that grasp of sortal concepts plays no role in enabling us to lock onto objects, but maintain that it has an ineliminable role in enabling us to think about objects because of its role in our grasp of genuine property ascription. And the new account of the role our grasp of sortal concepts plays will open up new questions about what sortal concepts, and grasp of sortal concepts, must be like if this role is to be filled.

The final direction of inquiry that I shall mention here concerns the similarities and differences between thought about ordinary objects and thought about other kinds of particular, for example, numbers, abstract objects, events, the objects posited by scientific theories. There is an old and hard question about the extent to which all of these things are ‘objects’ in the same sense. According to one traditional approach to this question, we should start with an account of how our thought about ordinary objects works, and then ask to what extent thought about other kinds of ‘object’ works in parallel ways. The line of thought I have set out in this paper opens up a hitherto overlooked point of comparison to consider in developing this approach: to what extent do these other domains admit of parallels to the causal loadedness argument for the generality of particular thought that I have set out in the home, ordinary object, case?

²³ See for example Campbell, ‘Does Visual Reference Depend on Sortal Classification’, *Philosophical Studies*, 127 (2006), pp. 221-237, §2. The standard sortalist response to this objection is to say that to refer to *o* S must know that there is some sort that *o* falls under, but need not know which sort this is. (See, for example, Evans, *VOR*, p. 178; Wiggins, *Sameness and Substance Renewed*, p. 7). Anti-sortalists reply that the capacity to pick a thing out does not require even this weaker involvement of sortal concepts. It is not possible to provide a full account of the objection here. My point is only that the debate about this intuitive objection in particular, and sortalism in general, will look very different given the motivation for sortalism that I have suggested.