

Thought About Properties: Why the Perceptual Case is Basic

Dominic Alford-Duguid

Forthcoming in *Philosophical Quarterly*

(Penultimate Draft: Please Cite Published Version)

Abstract: This paper defends a version of the old empiricist claim that to think about unobservable physical properties (such as spin or charge), a subject must be able to think perception-based thoughts about observable properties (such as red or circularity). The central argument builds upon foundations laid down by G. E. M. Anscombe and P. F. Strawson. It bridges the gap separating these foundations and the target claim by exploiting a neglected connection between thought about properties and our grasp of causation. This way of bridging the gap promises to introduce substantive constraints on right accounts of perception and perception-based thought.

How much of our thought about the world depends upon our being able to both perceive the world and think about what we perceive? One version of this old question concerns thought about objects. It asks how much of our thought about objects (e.g. thought about quasars or bosons) depends upon our capacity for perception-based thought about ordinary objects (such as sofas and cats). This paper focusses on a parallel question for thought about properties: How much of our thought about properties (e.g. charge and mass) depends upon our capacity for perception-based thought about observable properties (e.g. colours and shapes)? Common sense suggests that perception-based thought is more basic. Yet what does this basicness come to? This paper spells out—and argues for—a sense in which perception-based thought about observable properties is more basic than thought about unobservable properties. The argument exploits a deep but neglected connection between thought about properties and our grasp of causation.

Let's begin with some necessary terminology. Suppose a ball bounces into view. This perceptual encounter puts you in a position both to think about the ball and to ascribe bounciness to it, each in a particular way. Call these ways of thinking (of the ball on the one

hand, and bounciness on the other) ‘perceptual demonstration’ and ‘perceptual ascription’. In addition, let us say that a subject ‘identifies’ an object iff she stands in an identifying relation to it. An ‘identifying relation’ introduces an object as the object upon whose properties the truth or falsity of (at least one of) a subject’s thoughts depends. Perceptual demonstration is a form of identification. It contrasts with another form of identification, which I shall call ‘descriptive identification’. A person identifies an object ‘descriptively’ just if the identifying relation holds in virtue of (a) the subject grasping some descriptive condition and (b) the object being the unique satisfier of this descriptive condition (for example, someone might identify Frege descriptively as the author of the *Begriffsschrift*).

To consolidate this terminology, let’s use it to state an old and attractive picture of the structure of thought about objects. On this familiar view, our capacity for perceptual demonstration grounds our capacity to descriptively identify objects. Call the view ‘ATOMISM’.¹

Few have investigated whether an analogue of ATOMISM applies to perceptual ascription.² This paper defends one such analogue:

PREDICATIVE ATOMISM (Approximate Version) Our capacity to ascribe unobservable properties (such as spin or charge) ultimately depends upon our capacity to perceptually ascribe observable properties (such as round or red).

PREDICATIVE ATOMISM is more conservative than ATOMISM. While the latter posits a quite general relationship between two different ways of identifying objects, PREDICATIVE

¹ ATOMISM inherits its name from Russell’s (1918/1956) Logical Atomism. Aside from Russell, defenders of ATOMISM include Strawson (1959), Dummett (1973/1981), Evans (1982), Peacocke (1983), Brewer (1999), Campbell (2009), and Dickie (2010b). As formulated above, ATOMISM remains neutral with respect to the question of whether descriptive identification of an object counts as a way of thinking about an object. This question drives the contemporary debate over whether thoughts involving descriptive identification can count as singular thoughts (Cf. Jeshion (2002) and Recanati (2012)).

² Russell (1910; 1918/1956) is one of the few. He suggests that an analogue of ATOMISM applies to the perceptual ascription of properties. Yet Russell’s ability to back up this suggestion has been met with scepticism, most recently in Bostock (2012: Ch. 7).

ATOMISM (when made suitably precise) focusses on how our way of ascribing one type of property (i.e. unobservable properties) depends upon a way of ascribing another type of property (i.e. observable properties). Yet just as ATOMISM secures an ineliminable role for perceptual demonstration, PREDICATIVE ATOMISM entails that perceptual ascription is essential to our general capacity to ascribe properties (i.e. our general capacity to think thoughts we standardly express with sentences of the form $\lceil \alpha \text{ is } \phi \rceil$). My argument for PREDICATIVE ATOMISM not only secures this highly intuitive result, but also explains *why* perceptual ascription plays such an essential role.

The paper has three parts. §1 clarifies the content and scope of PREDICATIVE ATOMISM. §2 argues for the view. §3 defends a central premiss of the argument of §2.

§1 Understanding Predicative Atomism

I introduced PREDICATIVE ATOMISM as the claim that our capacity to ascribe unobservable properties (such as spin or charge) ultimately depends upon our capacity to perceptually ascribe observable properties (such as round or red). But this formulation isn't as transparent or precise as the version of PREDICATIVE ATOMISM I plan to defend. This section sharpens the initial formulation. I focus on three aspects of PREDICATIVE ATOMISM: how subjects ascribe unobservable properties; how they perceptually ascribe observable properties; and how the first kind of property ascription depends upon the second.

I shall start by introducing two ways of ascribing properties, what I shall call ascribing a property 'as dispositional' and 'as categorical'. This division roughly parallels a metaphysical distinction between two kinds of property. Metaphysicians divide properties on the basis of what it takes for an object (or some other particular) to instantiate them. A property is 'dispositional' when its instantiation conditions are exhausted by an object's

possession of a particular set of causal dispositions or causal powers (e.g. an electron's tendency to repulse other negatively charged particles). In contrast, a property is 'categorical' when its instantiation conditions are not exhausted by these causal dispositions.³

Unlike this metaphysical distinction, our division between two ways of ascribing properties does not depend upon the instantiation conditions of the ascribed properties. Instead it depends upon what a speaker must know in order to ascribe a property. A subject can ascribe a property only if she can identify objects, or other particulars, as bearers of the property (e.g. 'these chairs are red, those are not'). A kind of classificatory knowledge underwrites this capacity, knowledge manifested by a sensitivity to conditions under which a property is instantiated. Different ways of ascribing a property require distinct kinds of classificatory knowledge. A subject ascribes a property 'as dispositional' iff her classificatory knowledge consists in knowledge of the causal powers conferred by the property. If a subject does not ascribe a property as dispositional, she might instead ascribe it as categorical (the distinction isn't exhaustive). A subject ascribes a property 'as categorical' iff she ascribes it both (1) without relying only upon knowledge of causal powers it confers, and (2) by treating the property as a ground of its bearer's dispositions (e.g. someone who ascribes circularity as categorical can make sense of a circular object being inclined to roll on smooth inclines).

Perceptual experience is our primary, if not our only means of acquiring the classificatory knowledge that permits ascription of a property as categorical. Visual experience, for instance, seems to present some properties directly, unmediated by awareness of causal powers these properties confer upon the objects that instantiate them. Shape and

³ For a particularly clear statement of this distinction, see Armstrong (1997: Ch. 5).

size properties arguably belong in this category. Vehement disagreement persists over whether colours also belong to it.⁴

By contrast with the ascription of properties as categorical, both observable and unobservable properties may be ascribed as dispositional. However, unobservable physical properties must be ascribed as dispositional. For we only ever know unobservable physical properties such as negative charge or spin through their causal roles.⁵ Perceptual experience cannot provide the kind of direct epistemic access that would enable us to ascribe these properties without relying exclusively upon their causal roles. And once we step beyond perceptual experience, we seem left with no option but to identify unobservable physical properties through their effects on perceptually detectable phenomena. When classifying something as negatively charged, for example, we can only ever rely upon our knowledge that a negatively charged object will possess certain causal dispositions (e.g. an aversion to positively charged entities). Hence to ascribe an unobservable physical property as dispositional requires that a subject be capable of ascribing the causal relations that make up the causal role of the property (these capacities may be interdependent: neither need enjoy priority). This last result will serve as the first premiss of my argument for PREDICATIVE ATOMISM in §2.

⁴ Nothing in this paper hangs on whether we ascribe colours as categorical or as dispositional. Indeed, my definition of what it is to ascribe a property as categorical opens up a potential intermediate position (see §3.2.2 for discussion).

⁵ This epistemological claim is independently plausible, and has been argued for by others (most notably, Shoemaker (1980: 236)), though a proper defence of it lies outside the scope of this paper. However, one might worry that this claim fails to leave open the possibility that once a subject is able to ascribe some unobservable properties as dispositional, she could combine this capacity with knowledge of some non-causal facts (e.g. symmetries in the structure of properties) to ascribe other unobservable properties without relying upon knowledge of their causal roles (thanks to an anonymous referee for raising this worry). I believe we can set this worry aside. For even if it were possible to ascribe unobservable properties in the manner described, that fact would not rule out a weaker version of my original claim, namely that we can ascribe unobservable properties only if we can ascribe at least some of them as dispositional. The argument of the paper would still go through even with this weaker claim.

We can now take a step towards the precise version of PREDICATIVE ATOMISM that I shall defend:

PREDICATIVE ATOMISM (Quasi-Precise Version) Our capacity to ascribe unobservable properties depends upon our capacity to perceptually ascribe properties as categorical.

A question remains about the nature of the dependence relation invoked by PREDICATIVE ATOMISM. The relation is not merely nomological. If it were, PREDICATIVE ATOMISM would be an empirical claim about either psychological or developmental priority. While these sorts of empirical claims are of interest, they are not my target. PREDICATIVE ATOMISM is instead a thesis about the necessary structure of understanding. A subject must be able to exercise one capacity (to perceptually ascribe properties as categorical) in order to exercise another (to ascribe unobservable properties).

Yet the dependence relation in PREDICATIVE ATOMISM remains ambiguous. On one construal, any exercise of our capacity to ascribe unobservable properties must involve the exercise of our capacity to perceptually ascribe properties as categorical. But this version of PREDICATIVE ATOMISM seems implausibly strong. The version that I shall defend is weaker. It says that the dependence relation obtains not between token exercises of the respective capacities, but between each type of exercise. The presence of the one capacity requires the presence of the other, but we can exercise the capacity to ascribe unobservable properties without thereby exercising our capacity to perceptually ascribe properties as categorical. Call this ‘type-dependence’.

We now arrive at our canonical formulation of PREDICATIVE ATOMISM:

PREDICATIVE ATOMISM (Precise Version) Our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe properties as categorical.

Notice that this thesis carries no commitment to a *metaphysical* asymmetry between observable and unobservable properties. Unobservable properties may still be more fundamental than observable properties. I thus intend to remain neutral on the question of whether a right metaphysics must explain facts about the one sort of property in terms of facts about the other sort of property.

§2 The Argument for Predicative Atomism

This section sets out my argument for PREDICATIVE ATOMISM. I introduced the first premiss in §1. Recall that when we ascribe an unobservable property (such as spin) we must exploit our knowledge of its causal role. A property's causal role is its contribution to the causal behaviour of entities that instantiate it. But to grasp a property as the bearer of such-and-such causal role, we must have a grasp of causation. Putting this together we get:

CAUSALISM: Our capacity to ascribe unobservable properties exploits our capacity to ascribe the causal relations that hold in virtue of these properties (i.e. causal relations that constitute the causal roles of these properties).

CAUSALISM forges a tight connection between our capacities to ascribe unobservable properties and to grasp causation as a generic relation, since it is the latter capacity that underwrites our capacity to ascribe the causal relations that hold in virtue of unobservable properties. Yet to what do we owe this grasp of causation as a generic relation? I propose that we owe it to our grasp of what we shall call 'special causal relations' (call this new claim

‘CAUSAL ATOMISM’).⁶ These causal relations are those we both standardly express with ‘thick’ causal terms (e.g. boulder *flattens* a hut, a child *kicks* a ball, and so on) and may know of on the basis of experience. As with PREDICATIVE ATOMISM, the priority CAUSAL ATOMISM affords special causal relations lies at the level of understanding, not at the level of metaphysics (and we need not move directly from this kind of priority to a metaphysical priority thesis).

While I cannot provide a full defence of CAUSAL ATOMISM in this paper, I suggest that the following observation suffices to establish its plausibility. Notice that it is difficult, if not impossible, to imagine someone who grasps causation as a generic relation, but who nevertheless cannot ascribe special causal relations (i.e. someone who can think *A caused B*, yet is unable to think anything like *the boulder flattened the hut*). In contrast, nothing strikes us as incoherent about the possibility of a subject who grasps flattening, kicking, or other special causal relations, but cannot ascribe a generic causal relation. It is hard to see why this should be the case unless CAUSAL ATOMISM is true.⁷

My argument for PREDICATIVE ATOMISM now runs as follows:

1. CAUSALISM: Our capacity to ascribe unobservable properties exploits our capacity to ascribe the causal relations that hold in virtue of these properties (i.e. causal relations that constitute the causal roles of these properties).

⁶ Anscombe (1971), later followed by Strawson (1992: Ch. 9), argues for this claim about our grasp of causation.

⁷ It lies outside the scope of this paper to explore the nature of the connection between our grasp of special causal relations and our grasp of generic causal relations. However, Strawson (1992: 120) provides the following rough gloss on the connection: ‘In general, then, the search for causal theories is a search for modes of action and reaction which are not observable at the ordinary level (or not observable at all, but postulated or hypothesized) and which we find intelligible because we model them on, or think of them on analogy with, those various modes of action and reaction which experiences presents to gross observation...’

2. CAUSAL ATOMISM: Our capacity to ascribe causal relations that hold in virtue of unobservable properties type-depends upon our capacity to ascribe special causal relations.
3. Our capacity to ascribe unobservable properties type-depends upon our capacity to ascribe special causal relations. [From 1 and 2]

But what might our capacity to ascribe special causal relations require? I'll argue for the following answer, one that bridges the gap between 1–3 and PREDICATIVE ATOMISM:

4. BRIDGE PREMIS: Our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe observable properties as categorical.

3 and 4 entail PREDICATIVE ATOMISM:

5. PREDICATIVE ATOMISM: our capacity to ascribe unobservable properties type-depends upon our capacity to perceptually ascribe observable properties as categorical. [From 3 and 4]

In upholding the entailment from 3 and 4 to 5, I'm assuming that type-dependence is transitive. This transitivity should not be controversial. For I introduced type-dependence to capture what should be a transitive explanatory connection between cognitive capacities.

Before I go on to defend this argument in detail—the task of the next section—I wish to answer a pair of worries about the argument. Addressing these worries will bring the argument into sharper focus.

The first worry concerns the role of CAUSAL ATOMISM. It jumps off from the claim that CAUSAL ATOMISM is plausible partly because of our prior commitment to the old empiricist idea that thought about the unobservable depends upon thought about the observable. Yet surely—the worry continues—a theorist who accepts this old empiricist idea

does so in part because she already accepts a version of PREDICATIVE ATOMISM. Circularity seems to threaten.

However, this worry admits of an answer even if we grant (for the sake of argument) the alleged connection between CAUSAL ATOMISM and the old empiricist idea. The worry mistakes the nature of my project. I do not seek to establish the bare empiricist claim that thought about unobservable properties depends upon thought about observable properties. I instead aim to defend a particular *elucidation* of the bare empiricist claim (i.e. PREDICATIVE ATOMISM). I thus require only that CAUSAL ATOMISM help explain PREDICATIVE ATOMISM. Yet it can perform this role even if its plausibility somehow rests upon the bare empiricist claim (and thus upon the old empiricist idea that thought about the unobservable depends upon thought about the observable).

The second worry I wish to address concerns the relationship between my argument and an old debate (going back at least to Hume 1739) over whether we perceive causation. Hume argued that perception does not present causal relations. To think otherwise, Hume claimed, is to mistake perception of mere regularity for perception of genuine causality. But a growing cadre has sought to resist Hume's insistence that we do not perceive causation.⁸ Some may worry that my argument for PREDICATIVE ATOMISM takes a side in this old but fierce debate.

Yet this worry is misplaced. My argument remains neutral about whether we perceive causation. PREDICATIVE ATOMISM is a claim about the structure of thought, not a claim about what we can perceive. My argument for it turns on claims about the requirements on thinking certain kinds of thoughts. But whether we perceive causation or not fails to tell us if our capacity to ascribe special causal relations depends upon our capacity to perceptually

⁸ These rebels include Anscombe (1971), Strawson (1992: Ch. 9), and Siegel (2010).

ascribe properties as categorical. To see or feel causation may be one thing, a capacity to ascribe special causal relations—for instance to think of one object as flattening (or denting, or pushing...) a second object—may be quite another. Finally, special causal relations are not distinguished by being perceived, but by the fact that subjects may know about them on the basis of perception.

§3 A Foundation for Predicative Atomism

I have already briefly motivated two of the three premisses that support PREDICATIVE ATOMISM: CAUSALISM in §1, and CAUSAL ATOMISM at the beginning of §2. In this section I defend the BRIDGE PREMISS. More could be said in defence of the other two premisses. I focus on the BRIDGE PREMISS because it has yet to receive a decisive defence. My defence of the premiss is my principal positive contribution to the argument for PREDICATIVE ATOMISM.⁹ It is also this aspect of my argument for PREDICATIVE ATOMISM that promises to generate implications for perception and thought.

§3.1 Whence the Basicness of Perception-Based Thought?

Before I argue for the BRIDGE PREMISS, I shall briefly step back and highlight its position within the larger debate over the role of perception-based thought.

My defence of PREDICATIVE ATOMISM aligns me with those who regard perception-based thought as importantly basic. Attempts to work out this basicness have largely focussed upon perceptual demonstration, and in particular upon ATOMISM. Extant defences of ATOMISM standardly explain the primacy of perceptual demonstration over descriptive

⁹ Others have defended theses in the neighbourhood of the BRIDGE PREMISS. Strawson (1992: 122) appears to defend the nearby claim that our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe observable properties as dispositional (though Roessler 2011 argues that Strawson should have appealed instead to our capacity to ascribe these properties as categorical). Also, Peacocke (2011: 162–3) endorses something like the BRIDGE PREMISS (though he does not argue for it).

identification by appeal to a priority of reference over quantification.¹⁰ By contrast, my defence of PREDICATIVE ATOMISM explains the basicness of perceptual ascription by appeal to the role of such ascription in our grasp of causation—a role the BRIDGE PREMISS captures.

I am not alone, however, in defending the basicness of a category of perception-based thought by appeal to the role of such thought in our grasp of causation. John Campbell (2002, p. 247) defends the basicness of perceptual demonstration by arguing that our capacity to ascribe causal relations type-depends upon our capacity to perceptually demonstrate objects. I shall briefly sketch Campbell's line of thought, as it will provide a helpful contrast with my eventual argument for the BRIDGE PREMISS.

Campbell defends his claim about perceptual demonstration on the grounds that it solves an old puzzle, introduced by Shoemaker (1988, pp. 211-15), which challenges any view in which our capacity to think about objects rests on our capacity to ascribe causal relations. Shoemaker's puzzle begins with a question: what resources are required to capture our ordinary conception of physical objects (e.g. tables, houses, etc.)? We ordinarily conceive of the class of physical objects as a non-arbitrary ontological category. Thus a chair counts as a physical object, as does a boulder, but the mereological sum of the chair and the boulder does not. But if we assume a picture in which we think of objects as defined geometrically in terms of totalities of actual and possible space-time trajectories—areas of space whose contours change over time—we cannot generate a non-arbitrary reason to think that the objects picked out by our ordinary conception are special.¹¹ Consider the set of space-time points that specifies an object's trajectory in the actual world. There exists a set of points just like the original, but which begins five minutes earlier. Call this collection of points a

¹⁰ Cf. Russell (1918/1956).

¹¹ This picture has its roots in Quine (1960: §36).

‘shadow’ of the original object. We can define counterparts of a shadow object in every possible world in which the original object exists. Furthermore, one could repeat this shadow-making process for every object with which we interact, and for every object that these objects interact with, and so on. On our (temporarily) assumed picture, we have no non-arbitrary grounds for excluding these shadows from the class of physical objects.

The obvious next step, after abandoning our first assumed picture, would be to claim that physical objects have a particular kind of causal unity. It is thinking of objects in terms of their causal role that supposedly supports a conception of physical object that allows us to non-arbitrarily distinguish physical objects from their ‘shadows’. Yet this suggestion also quickly runs into difficulty. The shadow objects will exhibit the same regularities in the actual world, and their counterparts the same regularities across possible worlds, as the ordinary objects and their counterparts. So if a regularity or counterfactual account of causation were true, to think about objects as causal unities would fail to privilege physical objects over their shadows.¹²

Campbell argues that Shoemaker’s puzzle arises only if it is assumed that the question of which objects exist must be settled by appeal to causal relations we can ascribe without thinking about objects (except as the occupants of certain causal roles). He proposes

¹² To see why this holds, consider Lewis’s (1973) counterfactual account of causation. His account has two parts. First, an event e ‘causally depends’ on an event c iff (1) had c not occurred, e wouldn’t have; and (2) had c occurred, e would have. Second, c causes e iff there is a finite chain of events $e_1 \dots e_n$ that is such that e_1 causally depends on c , e_2 causally depends on e_1 , and so on until e causally depends on e_n . Now let o be a physical object, and o_s its shadow. If E is the set of events involving o (e.g. o ’s being pushed into another object o_1), then there exists a corresponding class of shadow events E_s involving o_s (i.e. if an event involving o also includes another object o_1 , the event’s shadow will involve o_1 ’s shadow; and if the first event starts at time t , the second will start at a ‘shadow’ time t_s ; and so on). Any counterfactual true of $e \in E$ will also be true of a corresponding event $e_s \in E_s$, since the counterfactuals true of events in E depend for their truth on the behaviour of counterparts of o (and other physical objects), and this behaviour will be mirrored by counterparts of o_s (and the other shadow objects). It follows, on Lewis’s account, that if $e \in E$, and e causally depends on c , then there will be an event $e_s \in E_s$ that causally depends on c_s (the shadow of c). And so if $e \in E$, and e causes c (or c causes e), then there will exist an event $e_s \in E_s$ such that e_s causes c_s (or c_s causes e_s). Thus any causal unity displayed by an ordinary object will be matched, on Lewis’s account of causation, by a parallel causal unity displayed by the object’s shadow.

we drop this assumption, and hold instead that our capacity to ascribe causal relations type-depends upon our capacity to perceptually demonstrate objects. Perceptual demonstration exploits our perceptual encounters with objects, rather than our grasp of a descriptive condition. It therefore enables us to identify objects without merely identifying them as the occupants of certain causal roles.

Campbell's argument provides an important structural contrast with my eventual defence of the BRIDGE PREMISS. Whereas Campbell seeks to capture our ordinary conception of physical objects, my defence of the BRIDGE PREMISS will seek to capture core features of our grasp of special causal relations. By relying directly upon the nature of our capacity to ascribe special causal relations, I provide a quite different sort of foundation for the basicness of our capacity to perceptually ascribe observable properties as categorical.

§3.2 Defending the Bridge Premiss

The BRIDGE PREMISS is the claim that our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe observable properties as categorical. I shall defend it by arguing for three nested component claims: (1) our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties; (2) we must ascribe at least some of the properties as categorical; and (3) this ascription must be perceptual. I shall defend each claim in turn.

§3.2.1 Our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties

One way to argue for (1)—the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties—would be to first assume a particular conception of causation. We could then argue that the ascription of special causal relations, construed according to our assumed conception of causation, requires that a thinker have

the capacity to ascribe properties. To ascribe a relation is to think that an entity stands in the relation to some other entity (e.g. that the table is next to the chair). And plausibly a subject can think this sort of thought only if she can think about (or ascribe) the entities the ascribed relation relates. Consequently, (1) will be nearly trivial if we assume a conception of causation on which causal relations relate properties (rather than events, states of affairs, or objects).

As trivial as (1) might seem on this account of causation, it appears much less plausible on others. For instance, what if causal relations relate events (e.g. the boulder falling) rather than constituents of these events (i.e. objects and properties)? Our capacity to think about an event may not depend upon our capacity to ascribe properties (e.g. it may be possible to perceptually demonstrate events). So it appears that one could accept that ascription of a causal relation requires a capacity to think about its relata, yet at the same time deny (1).

However, (1) would be considerably less interesting if it were found to rest on a particular account of what causal relations relate. This dependence would also undermine the interest of PREDICATIVE ATOMISM. My goal in this subsection is to argue that acceptance of (1) does not depend upon what we take causal relations to relate. I show that it has a deeper anchor.

An argument for (1)—the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties—that does not depend upon a particular account of what causal relations relate must spell out a connection between our capacity to ascribe special causal relations, on the one hand, and our capacity to ascribe properties, on the other. This connection needs to be quite general: it must hold of both our sophisticated ascriptions of causal relations and our most primitive.

My argument for (1) rests on a minimal constraint on our general capacity to ascribe relations in thought. I contend that the link between property ascription and the ascription of special causal relations has its roots in the generality implicit in any ascription of a relation. Consider the Generality Constraint defended by Evans (1982): necessarily, to count as thinking that a is F , a thinker must also have the capacities to think that a is H , a is G , and so on (for some suitable range of H , G , etc.), and to think that b is F , c is F , and so on (for some suitable range b , c , etc.). This constraint extends to the ascription of relations: necessarily, to count as thinking that a stands in R to b , a thinker must also have the capacities to think that a stands in R_2 to b , a stands in R_3 to b , and so on (for some suitable range of R_2 , R_3 , etc.), and to think that a_1 stands in R to b_1 , a_2 stands in R to b_2 , and so on (for some suitable range of pairs a_1 and b_1 , a_2 and b_2 , etc.).

Controversy persists over whether and how to restrict the two ranges of thoughts.¹³ But my argument relies on only a minimal constraint entailed by any relational version of Evans's Generality Constraint. This minimal constraint is the claim that a subject counts as ascribing a relation only if she could apply this relation to other relata. Put formally: necessarily, to think that a stands in R to b , a subject must have the capacity to think that a_1 stands in R to b_1 (for at least one pair a_1 and b_1 distinct from the pair of a and b). Applied to special causal relations, my minimal constraint entails that to ascribe a particular special causal relation a subject must know what it would be for a different causal interaction to involve that same causal relation. But to apply a relation to other causal interactions in this fashion a subject must grasp the relation as *repeatable*.

¹³ Recent advocates of (versions of) the Generality Constraint include Campbell (1986), Peacocke (1992), Camp (2004), Dickie (2010a), Beck (2012), and Recanati (2012). Evans restricts the ranges of thoughts to those that obey categorial restrictions (e.g. we need not be able to think that Julius Caesar is prime), a restriction Camp (2004) resists. Dickie (2010a) pushes in the other direction: she shows that extant arguments fail to establish even Evans's restricted constraint.

What does it take to grasp a special causal relation as repeatable? Objects, for example, enter into a wide range of special causal relations: an object can flatten, disperse, discolour, and so on. So to ascribe a particular special causal relation as repeatable a subject must be able to know what makes it the case that *this* relation, rather than some other one, obtains from one moment to the next (e.g. what makes it the case that, during a given interval of time, an object is flattening a hut rather than toppling it).

Whether one special causal relation or another obtains over an interval depends on the sorts of changes involved. Some relations hold in virtue of one set of properties, others in virtue of different properties. Change some of these properties, and one sort of causal relation will obtain; change another set of properties, and a different sort of causal relation may result. For example, a causal interaction counts as an instance of flattening in virtue of properties of both the flattening object (the *agent*) and the flattened object (the *patient*). A subset of the agent's properties brings about a particular change in the shape and size of the patient (i.e. the patient is flattened). If a different subset of properties of the patient were changed, the causal interaction wouldn't be a flattening (it might instead be a discolouring, if the patient's colours change).

I'm now in a position to offer my argument for the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties. We begin with the application of my minimal constraint to the ascription of special causal relations:

1. For a subject to ascribe a particular special causal relation she must know what it would be for a different causal interaction to involve that same causal relation—she must grasp the relation as *repeatable*.
2. But to grasp a special causal relation as repeatable requires that a subject be able to think, of some other causal interactions, that they involve instances of the ascribed

causal relation (for instance, to grasp what it is for one thing to *push* another, I must be able to think, of some other scenario, that it involves one thing pushing another).

However,

3. Being able to think this way requires that a subject distinguish (in thought) that relation from other special causal relations.

After all, I cannot entertain the thought *Judy pushes Jeremy* unless I can distinguish (in thought) Judy pushing Jeremy from Judy tripping Jeremy. But

4. To isolate a *particular* special causal relation a subject must distinguish a *particular* dimension of change.

1–4 generate an initial conclusion about the ascription of special causal relations:

5. In order to ascribe a particular special causal relation a subject must be able to distinguish a *particular* dimension of change. [From 1–4]

Yet

6. For a subject to distinguish a particular dimension of change it is not sufficient to think of a bearer of properties as *somehow* altered; she must think of it as changed in a particular way.

Notice that 6 (and thus my argument as a whole) does not presuppose a particular conception of what special causal relations relate. While my illustrative examples used objects that serve as agent and patient, 1–6 remain silent as to the relationship between these objects and the relata of special causal relations. The relata could be events that have these objects as parts, or facts with the objects as constituents; they could even be the objects themselves. So my argument preserves an important degree of neutrality about the nature of causation. However, we still require a gloss on what it is for a bearer of properties to change in a particular way:

7. These ways for bearers of properties to change are distinguished by the properties these changes involve. If we want to think of a bearer of properties as changed in a particular way, we must be able to think about what distinguishes this way of being changed.

So 6 and 7 entail 8:

8. To distinguish a particular dimension of change a subject must be able to ascribe properties. [From 6 and 7]

And 5 and 8 entail 9:

9. A subject who ascribes special causal relations must be able to ascribe properties. [From 5 and 8]

If ϕ -ing requires being able to X , our capacity to ϕ type-depends on our capacity to X . So 9 entails (1): the claim that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties.

This argument represents a real step forward. It establishes a deep and direct connection between the ascription of properties and the ascription of special causal relations—a connection that remains of independent interest even when we ignore its link to PREDICATIVE ATOMISM. The argument also avoids presupposing a specific account of what special causal relations relate. It instead derives the connection from a general claim about the conditions under which a subject can ascribe a relation.

§3.2.2 Our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties as categorical

This subsection lays the second stone on the road to the BRIDGE PREMISS. We must show that the ascription of special causal relations type-depends on the ascription of properties as

categorical. As we've already seen, our capacity to ascribe special causal relations depends upon our capacity to ascribe properties, in part because the latter capacity enables us to isolate *particular* special causal relations. I shall argue that this central role for property ascription cannot be fulfilled by a capacity to ascribe properties as dispositional. I offer two arguments for this claim: the first relies upon a claim about our cognitive limitations; the second reaches for a deeper anchor. With the claim in hand, I conclude that the ascription of special causal relations type-depends on the ascription of properties as categorical.

This conclusion follows because the central role for property ascription carved out in §3.2.1 must be filled either by a capacity to ascribe properties as dispositional or by a capacity to ascribe them as categorical. For if a subject were to ascribe a property neither as dispositional nor as categorical, she would ascribe it without relying only upon knowledge of causal powers it confers, and without treating the property as a ground of its bearer's dispositions. She would therefore grasp the property neither as a direct contributor to causal interactions (as she would were she ascribing it as dispositional) nor as a ground for these interactions (as she would were she ascribing the property as categorical). Yet a subject who grasps a property in this way cannot use it to isolate a particular special causal relation within a diverse range of causal interactions, since she cannot treat changes in that property as shaping how those causal interactions unfold.

Argument From Cognitive Limitations A special causal relation is the realisation of a causal power possessed by an object. If a boulder flattens a hut, for example, this flattening is the realisation of a causal power belonging to the boulder: roughly, to flatten objects of such-and-such a size, when dropped from a certain height, when the flattened object is made of suitably fragile material, etc. As this example makes vivid, any proper specification of the

causal power must be enormously complicated, since whether the causal relation obtains will be conditional upon myriad complex background conditions.¹⁴

Now for our ascription of properties as dispositional to underwrite our capacity to distinguish particular special causal relations, it would have to be within our cognitive reach to isolate whichever highly conditional causal powers distinguish a particular special causal relation from other such relations. To isolate these causal powers would require knowledge of how their realisation is sensitive to changes in background conditions. And this does not seem like something a creature with our limited cognitive abilities could do. We can ascribe properties as dispositional without possessing the kind of detailed conditional knowledge required to isolate one of the causal powers conferred by an ascribed property. But without this fine-grained knowledge these property ascriptions will not underwrite any ascription of a special causal relation.

This issue of cognitive load does not afflict the ascription of properties as categorical. The problem with ascribing properties as dispositional was that we could not isolate the specific contribution of a property to change (and thus causation) without detailed knowledge of the conditions that must be met for the associated causal powers to be realised. By contrast, isolating the contribution of a property ascribed as categorical does not require that we inquire into the detailed realisation conditions of any associated causal powers. To consider a property as categorical is to consider it as something which acts as an autonomous basis of change and causation. For example, when a wheel rolls down a hill, an adequate account of what caused this behaviour may simply appeal to the roundness of the

¹⁴ On the account of dispositional properties developed by Shoemaker (1980), an object has a power conditional on a property or set of properties Q if it has some other property Φ such that the combination of Q and Φ is causally sufficient for having the power, while Q is not alone causally sufficient for having the power. Shoemaker's account treats dispositional properties as clusters of these conditional causal powers.

wheel. So only the ascription of properties as categorical can underwrite our ascription of special causal relations—at least for subjects with limited cognitive resources.

A Deeper Anchor While my first argument gets us the conclusion we want, we can do better. In what follows I develop a second argument that eschews reliance upon claims about our limited cognitive resources. The argument establishes an even more secure foundation for the claim that ascription of a special causal relation requires ascription of a property as categorical.

My first argument exploited the fact that any adequate specification of a causal power must be enormously complicated (since whether the causal relation obtains will be conditional upon myriad complex background conditions). Following Shoemaker (1980), call these background-sensitive causal powers ‘conditional causal powers’. The inevitable complexity of a proper specification of a conditional causal power prevents agents with our limited cognitive resources from acquiring sufficiently detailed knowledge of the cluster of causal powers conferred by a property.

I shall argue that this complexity introduces yet another obstacle for subjects who ascribe properties as dispositional—an obstacle that would face even those without our limited cognitive resources. Suppose I come to think that a rolling ball pushed a cup. To think that one object pushes another is to ascribe a special causal relation—pushing. We’ve already seen that to ascribe this relation I must know what it would be for some other pair of objects to be such that one pushes the other. Yet knowledge requires non-luckiness: when thinking of some other situation involving a pair of objects, my ability to isolate their interaction as an instance of pushing (rather than some other relation) cannot rest on mere luck. For instance, if I think of a pair of objects interacting in a certain way—where this interaction actually constitutes an instance of *pulling* rather than *pushing*—I must count as

unlucky if I think of it as an instance of pushing. (To be clear: the non-luckiness involved concerns a subject's ability to isolate a relation *in thought*, not her ability to recognise instances of this relation out in the world on the basis of perception, testimony, or the like.)

Yet what does it take for a subject to non-luckily isolate a special causal relation? Since this capacity rests on our ability to isolate (again in thought) the changes characteristic of the ascribed special causal relation, our isolation of these changes must also count as non-lucky. But we isolate these changes in terms of the properties they involve. So isolating them in thought requires an ability to non-luckily isolate the relevant properties—and (given our story about the ascription of special causal relations) to do so within a range of potential causal interactions. For example, suppose that the change characteristic of pushing involves a certain motion property. Isolating the former in thought will require isolating the latter. And this isolation will require non-luckily distinguishing that motion property from similar (but distinct) motion properties whose presence in a potential causal interaction would not help constitute the sort of change characteristic of pushing.

We now have a new constraint on the kind of property ascription that must underwrite our capacity to ascribe special causal relations: the classificatory knowledge characteristic of this kind of property ascription must enable a subject to non-luckily isolate a property in thought within a diverse range of potential causal interactions. In particular, the classificatory knowledge cannot leave open rationally-relevant potential causal interactions in which the property a subject isolates is not the ascribed property (where a circumstance is 'rationally relevant' iff a rational subject must take it into account).¹⁵ These potential causal interactions will often involve quite diverse background conditions. As a result, the conditional causal powers manifested in these cases will vary widely.

¹⁵ I owe this notion of rational relevance to Dickie (2015: 42).

I shall argue that only our capacity to ascribe properties as categorical can satisfy the new constraint. Hence only this capacity can underwrite our capacity to ascribe special causal relations.

1. Suppose all we have is a capacity to ascribe properties as dispositional.

It follows that

2. We isolate properties in thought by treating them as clusters of conditional causal powers. [From 1]

But we've seen that

3. Being able to ascribe a special causal relation requires being able to non-luckily isolate properties across a diverse range of potential causal interactions.

Yet

4. When attempting to isolate a property F within a potential causal interaction using a subset of its conditional causal powers, there will always be rationally-relevant potential causal interactions in which a non-F property manifests that subset of F's conditional causal powers.

Why accept 4? Suppose I want to use a subset of a property's conditional causal powers to isolate it within a diverse range of potential causal interactions. The members of this subset will also belong to other clusters of conditional causal powers—each of which individuates a property. For instance, nearby determinates of a determinable property will share many of the original property's conditional causal powers. If these determinates are reasonably different (e.g. circularity and marginal-oval-ness), subjects will often be able to discriminate them by the conditional causal powers they manifest. As a result, these nearby determinates will count as rationally relevant. Unlike rarely instantiated properties that share causal powers with an ascribed property (and which are rationally irrelevant due to their rarity), these

properties are pervasive. Hence they belong to the class of properties whose presence we must be able to distinguish from the presence of the original determinate property when we seek to non-luckily isolate the latter in thought. 4 follows once we admit these nearby-but-discriminable properties—those that share a subset of their conditional causal powers with the ascribed property. For when we try to isolate a property F within a diverse range of potential causal interactions, there will be rationally-relevant potential causal interactions involving the nearby-but-discriminable-from-F properties.

A property manifests only a subset of its associated conditional causal powers in any given situation (we cannot after all have every condition operative at once). But then given that our classificatory knowledge exploits only a subset of F's conditional causal powers, 2 and 4 entail that attempts to use this classificatory knowledge to isolate F in thought will fail to eliminate rationally-relevant potential causal interactions which involve non-F properties that manifest the relevant subset of F's conditional causal powers. So:

5. If a subject attempts to isolate F within a diverse range of potential causal interactions, she will fail to non-luckily isolate F.

If ascribing a property as dispositional prevents a subject from being able to non-luckily isolate other instances of the property across a range of potential causal interactions, 3 entails that this kind of property ascription cannot underwrite our capacity to ascribe special causal relations.

Note that 5 does not prevent subjects from ascribing properties as dispositional. For the range of potential circumstances within which we must isolate a property when simply ascribing it as dispositional (as opposed to ascribing it for the purpose of ascribing special causal relations) need not involve the kind of variation in background conditions that at 5 undermines non-luckiness. All I've shown is that our capacity to ascribe properties as

dispositional cannot underwrite our ascription of special causal relations. Nothing prevents this capacity from performing other explanatory roles.

In contrast to those who must ascribe properties as dispositional, a subject who can also ascribe properties as categorical faces no parallel obstacle to their capacity to ascribe special causal relations. If a subject ascribes a property as categorical, isolating this property in thought will rarely depend on luck. The classificatory knowledge that underwrites the ascription of a property as categorical enables a subject to non-luckily isolate the property within a diverse range of potential causal interactions. For this knowledge isn't restricted to causal powers a property might share (in whole or in part) with others. As a result, a subject who ascribes a property as categorical can non-luckily isolate the property within a range of potential causal interactions—an ability that will enable her to ascribe special causal relations.

This second argument has been more involved than the first. Yet I believe it succeeds in strengthening my case for the claim that ascription of special causal relations requires a capacity to ascribe properties as categorical.

Objection and Reply The way I established 4 opens up a potential objection to my second argument: while ascription of a special causal relation requires isolating a change characteristic of that relation, an opponent could insist that this does not require a subject to isolate causally-fine-grained properties (in particular, such an opponent could insist that a subject need not distinguish a property from others that would make the same contribution to the rationally-relevant causal interactions). But if ascribing a special causal relation were to require merely the ascription of causally-coarse-grained properties, we arguably would lose the threat to our non-lucky isolation of these properties in thought. For suddenly any property we might isolate using a set of conditional causal powers would count as an instance of the property we would be seeking to isolate in thought.

I believe this objection fails. To begin with, I'm sceptical that any special causal relation is such that ascribing it requires merely the ascription of causally-coarse-grained properties. I also believe that a more principled reply is available. In order to get this reply into view, we must begin with some background from the metaphysics of causation. Causal relations are immune to certain sorts of fragility.¹⁶ Suppose I throw a ball that breaks a window. Had I thrown the ball in a radically different way, or had I thrown a much lighter ball, or had I thrown it at a time when the window was not there, my throw would not have broken the window. Thus certain alterations will turn an instance of a causal relation into a scenario where that causal relation does not obtain. However, not every alteration has such a destructive effect on a causal interaction. It is in this sense that causation isn't fragile. For instance, if I threw the ball a moment earlier, or if the ball were an ounce heavier, or if the window were slightly closer, my throw would still break the window.

The next step in answering the objection involves a move from the metaphysics of causation to a right account of our grasp of special causal relations. We would look askance at a subject who claims to be able to ascribe a special causal relation, but who refuses to ascribe the relation except when conditions exactly match those under which the relation was first encountered. Hence, a subject can ascribe a special causal relation only if she is in a position to recognise the non-fragility of the relation.

Let us now return to the original ball-throwing example. When I throw the ball a different way, my breaking of the window involves changes that are distinct from those involved in the first throwing (though these changes must be of the same type, since we have typed special causal relations by the kind of changes they involve). Yet if we were only able to ascribe causally-coarse-grained properties, we would lose our cognitive grip on this

¹⁶ Cf. Lewis (2000).

distinction between the changes that involved in my original throwing, and the changes involved after causation-sustaining alterations. As a result, we can recognise the non-fragility of my throwing only if our grasp of these changes exploits the ascription of causally-fine-grained properties.

It seems that our capacity to grasp the changes involved in a special causal relation must do double duty: it must *both* permit us to isolate the relation *and* permit us to recognise its non-fragility. And we've just seen that for the capacity to do the latter, it must exploit a further capacity, namely a capacity to ascribe causally-fine-grained properties. Our capacity to isolate special causal relations therefore also depends upon our capacity to ascribe causally-fine-grained properties. For we can isolate a special causal relation only if we can isolate the changes that relation involves, and we can do that only if we can ascribe causally-fine-grained properties. In effect, my more principled reply to the original objection has used our grasp of causation's non-fragility to introduce a new constraint on a requirement for isolating special causal relations (namely the requirement that we isolate some particular dimension of change that is characteristic of the special causal relation).

§3.2.3 Our capacity to ascribe special causal relations type-depends upon our capacity to perceptually ascribe properties as categorical

I've shown that our capacity to ascribe special causal relations type-depends upon our capacity to ascribe properties as categorical. This final subsection takes us from that result to the claim that our capacity to ascribe special causal relations type-depends upon our capacity to *perceptually* ascribe properties as categorical. In light of what has been established in the previous subsections, to falsify the BRIDGE PREMISS at this stage a subject would have to ascribe a physical property as categorical without relying upon either perception of the

property or a capacity to perceptually ascribe other properties as categorical. Yet someone can ascribe a property as categorical only if she possesses a suitable kind of classificatory knowledge: it cannot be exhausted by her knowledge of the property's causal role. This much follows from our original definition of what it is to ascribe a property as categorical. But from whence could such knowledge derive, if not perceptual experience of the properties themselves?

A temptation will strike some opponents at this juncture. They will appeal to examples usually assumed to undermine broadly empiricist accounts of concept possession, such as Hume's missing shade of blue. Hume (1739, 1.1.1.10) imagines a subject who has seen all but one of the visible shades of blue. He claims that if these shades were placed before the subject in order of blueness, the subject would be able to think about the missing shade of blue (despite never having seen it). But while such examples undermine positions on which 'more complex' property ascriptions must be analysed in terms of perceptual ascription, they do not undermine a view on which the ascription of properties as categorical type-depends upon perceptual ascription of properties as categorical. A subject can ascribe Hume's missing shade of blue, for example, because of a prior capacity to perceptually ascribe a range of colours. Similar observations hold of potential parallels to the missing shade of blue for shape, size, and other observable properties. So we can accept that subjects may ascribe some observable properties as categorical in the absence of perceptual encounters with these properties, yet deny that these ascriptions are more basic than perceptual ascription of properties as categorical.

Concluding Remarks

The task of this paper was to defend PREDICATIVE ATOMISM:

PREDICATIVE ATOMISM: our capacity to ascribe unobservable properties type-
depends upon our capacity to perceptually ascribe observable properties as
categorical.

My major contribution to the argument for PREDICATIVE ATOMISM was my defence of the
BRIDGE PREMISS in §3.2:

BRIDGE PREMISS: Our capacity to ascribe special causal relations type-
depends upon
our capacity to perceptually ascribe observable properties as categorical.

I established the BRIDGE PREMISS by defending three nested component claims: (1) our
capacity to ascribe special causal relations type-
depends upon our capacity to ascribe
properties; (2) we must ascribe at least some of the properties as categorical; (3) and this
ascription must be perceptual.

My defence of PREDICATIVE ATOMISM promises to introduce substantive constraints
on right accounts of perception and perception-based thought. To see why, cast your mind
back to our original parallel between PREDICATIVE ATOMISM and ATOMISM—the view that
our capacity for perceptual demonstration grounds our capacity to descriptively identify
objects. ATOMISM carves out an explanatory role for perceptual demonstration. This role in
turn constrains a right account of such thought (e.g. it cannot explain perceptual
demonstration in terms of descriptive identification). The explanatory role also promises to
constrain our account of perception.¹⁷ A right account of perception must accommodate the
fact that perception permits perceptual demonstration. Hence an explanatory role that
constrains our account of perceptual demonstration will also indirectly shape our account of
the perceptual contact that permits perceptual demonstration.

¹⁷ Campbell (2002), for example, argues for his controversial Relational View of visual experience partly by
exploiting an ATOMISM-inspired explanatory role for perceptual demonstration.

I suggest that my defence of PREDICATIVE ATOMISM introduces a parallel set of constraints. In conjunction with the BRIDGE PREMISS, PREDICATIVE ATOMISM carves out a rich explanatory role for the perceptual ascription of observable properties as categorical. And this explanatory role will in turn constrain both our account of that capacity, and our account of the perceptual contact that permits perceptual ascription of observable properties as categorical. It lies outside the scope of this paper to investigate the precise character of these constraints, or to engage in substantial speculation about their downstream influence on right accounts of perception and thought. Yet my defence of PREDICATIVE ATOMISM and the BRIDGE PREMISS will serve as cornerstone for any such investigation.

References

- Anscombe, G. E. (1971) *Causality and Determination*. Cambridge: CUP.
- Armstrong, D. (1997) *A World of States of Affairs*. Cambridge: CUP.
- Beck, J. (2012) 'The Generality Constraint and the Structure of Thought', *Mind*, 121/483: 563–600.
- Bostock, D. (2012) *Russell's Logical Atomism*. Oxford: OUP.
- Brewer, B. (1999) *Perception and Reason*. Oxford: OUP.
- Camp, E. (2004) 'The Generality Constraint and Categorical Restrictions', *Philosophical Quarterly*, 54/215: 209–230.
- Campbell, J. (1986) 'Conceptual Structure', in C. Travis (ed.) *Meaning and Interpretation*. Oxford: Blackwell.
- Campbell, J. (2002) *Reference and Consciousness*. Oxford: OUP.
- Dickie, I. (2010) 'The Generality of Particular Thought', *Philosophical Quarterly*, 60/240: 508–531.
- Dickie, I. (2010) 'We Are Acquainted With Ordinary Things', in R. Jeshion (ed.) *New Essays on Singular Thought*. Oxford: OUP.
- Dickie, I. (2015) *Fixing Reference*. Oxford: OUP.
- Dummett, M. (1973/1981) *Frege: Philosophy of Language*, 2nd edn. Cambridge, MA: Harvard University Press.
- Evans, G. (1982) *The Varieties of Reference*. (J. McDowell, ed.). Oxford: OUP.
- Hume, D. (1739) *A Treatise of Human Nature*, 1896 edn. (L. A. Selby-Bigge, ed.). Oxford: Clarendon Press.
- Jeshion, R. (2002) 'Acquaintanceless De Re Belief', in J. Campbell, M. O'Rourke, & D. Shier (eds.) *Meaning and Truth: Investigations in Philosophical Semantics*. Seven Bridges Press.
- Lewis, D. (1973) 'Causation', *Journal of Philosophy*, 70: 556–567.
- Lewis, D. (2000) 'Causation as Influence', *Journal of Philosophy*, 97: 182–197.

- Peacocke, C. (1983) *Sense and Content*. Oxford: Clarendon Press.
- Peacocke, C. (1992) *A Study of Concepts*. Cambridge: MIT Press.
- Peacocke, C. (2011) 'Representing Causality', in T. McCormack, C. Hoerl, & S. Butterfill (eds.) *Tool Use and Causal Cognition*. Oxford: OUP.
- Quine, W. V. (1960) *Word and Object*. Cambridge, MA: MIT Press.
- Recanati, F. (2012) *Mental Files*. Oxford: OUP.
- Roessler, J. (2011) 'Perceptual Causality, Counterfactuals, and Special Causal Concepts', in C. Hoerl, T. McCormack, & S. R. Beck (eds.), *Understanding Counterfactuals and Causation: Issues in Philosophy and Psychology*. Oxford: OUP.
- Russell, B. (1918/1956) 'The Philosophy of Logical Atomism', in R. Marsh (ed.) *Logic and Knowledge*. London: Allen and Unwin.
- Shoemaker, S. (1980) 'Causality and Properties', in P. v. Inwagen (ed.) *Time and Cause*. Dordrecht: D. Reidel.
- Shoemaker, S. (1988) 'On What There Are', *Philosophical Topics*, 16: 201–223.
- Siegel, S. (2010) *The Contents of Visual Experience*. Oxford: OUP.
- Strawson, P. F. (1959) *Individuals*. London: Methuen.
- Strawson, P. F. (1992). *Analysis and Metaphysics*. Oxford: OUP.