

The Causal Theory of Properties

Ann Katherine Whittle
University College London
PhD 2003

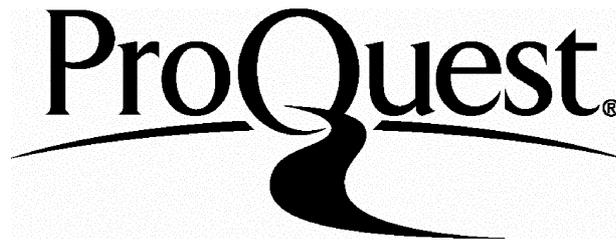
ProQuest Number: 10016078

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



ProQuest 10016078

Published by ProQuest LLC(2016). Copyright of the Dissertation is held by the Author.

All rights reserved.

This work is protected against unauthorized copying under Title 17, United States Code.
Microform Edition © ProQuest LLC.

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Contents

<i>Abstract</i>	4
<i>Acknowledgements</i>	5
Introduction	6
1. The Causal Theory of Properties	
1.1 The Theory Under Scrutiny	8
1.2 Varieties of CTP	11
1.3 The Scope of CTP	15
1.4 The Dispositions Debate	23
1.5 Individuation	32
2. Generalising Functionalism	
2.1 Functionalism Outlined	42
2.2 CTP as Functional Realiser Theory	53
2.3 CTP as Functional Role Theory	58
2.4 A Problem	61
2.5 A Metaphysical Worry	64
3. Grounding Causal Powers	
3.1 The Grounding Intuition	67
3.2 CTP in Trouble Again	79
3.3 A Solution	82
3.4 More Solutions	87
4. The Relata of Causation	
4.1 CTP's Causes and Effects	91
4.2 Davidson's Rival Ontology	95
4.3 The Slingshot	104
4.4 Omissions	109
5. Singularism	
5.1 The Intrinsicity Assumption	113
5.2 The Generality Assumption	123
5.3 Anscombe's Challenge	127

5.4	Armstrong's Reply	129
5.5	An AT-Singularist's Response	132
5.6	CTP's Solution	135
5.7	Causal Relata Again	139
6.	Properties and Their Instances	
6.1	CTP Expounded	144
6.2	Determinables and Determinates	150
6.3	MacDonald versus Yablo	158
6.4	The Nature of Tropes	166
7.	The Laws of Nature	
7.1	Modal Implications	172
7.2	A Case for Metaphysically Necessary Laws	176
7.3	CTP and the Contingency Thesis	183
7.4	The Conceivability Argument	187
	Conclusion	194
	<i>References</i>	196

Abstract

This thesis investigates the causal theory of properties (CTP). CTP states that properties must be understood via the complicated network of causal relations to which a property can contribute. If an object instantiates the property of being 90°C, for instance, it will burn human skin on contact, feel warm to us if near, etc. In order to best understand CTP, I argue that we need to distinguish between properties and particular instances of them. Properties should be analysed via the causal relations their *instances* stand in, it is *this oven's* being 90°C which causes my skin to burn, etc.

The resulting CTP offers an illuminating analysis of properties. First, it provides a criterion of identity for properties, their identity being analysed via the causal roles property instances realise. It also offers an account of how property instances are sorted into genuine kinds, in cases of determinables and determinates. I show how we can distinguish between genuine and non-genuine similarity via the property instances of objects.

The implications of CTP for an analysis of causation are then investigated. I argue that the proposed CTP offers a plausible causal ontology. The fine-grainedness of property instances enables us to capture the subtleties involved in questions concerning what causes what. But, even more importantly, CTP enables us to reconcile two highly attractive theses concerning the causal relation. The first of these is the generalist's thesis. This states that causal relations are part of more general patterns. The second of these is the singularist's thesis. This states that the causal connection between two entities, doesn't depend upon anything extraneous to that relation. I argue that by combining CTP with an ontology of tropes, we can thereby respect what is driving both singularism and generalism.

Acknowledgements

This thesis was paid for by the Arts and Humanities Research Board and by the Philosophy Department of University College London. I am extremely grateful for their financial assistance

An enormous philosophical debt is owed to my primary supervisor, M.G.F. Martin (aka. The Pope). All that is good about this thesis can be directly attributed to him. I would also like to thank my secondary supervisor, Tim Crane, for his astute comments and down to earth, philosophical good sense. Thanks also go out to countless staff and students of the University of London. Many have been bored by my musings on the subject of this thesis, but I have benefited enormously from their contributions.

Finally, a special thanks to Joel Smith. My live in philosopher and now husband has endured much in the making of this thesis. To him I owe so much.

Introduction

Ann sits down at her desk to write her thesis. The enormity of the task ahead fills her with anticipation and fear. This causes her mind to go completely blank, making her unable to write anything.

Here we have a causal story. It cites properties of Ann (namely, her anticipation, fear and blank mind) and properties of the task ahead of her (namely, its enormity). Furthermore, it claims that there is a causal relation between the enormity of the task and her mind being blank. Causal stories like this raise a number of questions. First, there are issues surrounding the nature of causation: What are the entities which cause and affect each other? What is it for one thing to cause another? Second, there are issues surrounding the nature of properties: What are properties, such as Ann's *being fearful* or the task *being enormous*, like? What relation do they bear to their objects? What role do they play in causation?

These questions are all ones that I intend to investigate during the course of this thesis. In particular, I shall be focusing on the role that properties occupy in causation. There certainly seems to be a close relationship between properties and causation. If I burn my hand by accidentally brushing it against a hot stove, we will pick out the temperature of the stove as being particularly causally relevant to the burn. For we know that we could have changed the stove's weight, make, colour, shape etc. without this altering the effect. What seems to be making the difference here, is the stove's temperature. How, then, should we account for this special role that properties appear to have in causation? Should we try to minimise it, or make much of it?

The issues raised by the causal story are difficult and far ranging. But there will be a common thread running throughout the discussion, as I intend to look at what perspective the causal theory of properties can offer on these matters.¹ In chapters one and two, I spend sometime considering how CTP should be formulated. A number of different proposals are investigated, but many of these, as we'll see, prove problematic. In chapter three, I consider an oft-cited objection to CTP. This leads us

¹ Throughout this thesis, I shall refer to the causal theory of properties as CTP for short.

into the territory of causal powers. I argue that, contrary to the claims of many, CTP is not committed to an implausible analysis of causal powers. By so doing, I suggest a way of viewing causal powers which coheres well with CTP's commitments.

Chapters four and five continue by looking at CTP's consequences for a theory of causation. I begin by arguing that CTP theorists should defend the claim that property instances are the (fundamental) causal relata, as this is strongly suggested by CTP. Considerations for and against this thesis are thus canvassed. Chapter five argues that CTP, combined with a causal ontology of *sui generis* property instances, has the resources to reconcile two plausible theses regarding the nature of the causal relation. Although a comprehensive analysis of the causal relation is still a long way off, the proposed CTP is shown to have plausible implications for such an analysis.

Chapter six returns to CTP's analysis of properties. I consider different ways of developing that theory, given the constraints that have been imposed by the preceding discussion. I also look at whether the resulting account can deal adequately with the relationship between determinables and their determinates. Finally, I come to an issue which has dominated discussions of CTP, namely, its commitment to the thesis that the laws of nature are metaphysically necessary. I look at arguments for and against this thesis, and consider whether there is any way a CTP theorist could avoid this commitment.

My hope is that by providing a careful formulation of CTP, and looking to see what its consequences are for an analysis of causation and properties, CTP will thereby be placed in a more positive light.

1. The Causal Theory of Properties

1.1 The Theory Under Scrutiny

What makes a property the property it is, what determines its identity, is its potential for contributing to the causal powers of the things that have it (Shoemaker 1980a, p.212).

Since Shoemaker wrote his paper “Causality and Properties” (1980a), increasing attention has been directed towards CTP. This interest has led to a number of different permutations of CTP. But before getting embroiled in difficult questions about how this thesis should be understood, I shall first try to offer a taster of what CTP is all about.

The causal theory of properties is so called because it claims that properties should be analysed via their causal features. Shoemaker, for instance, in the quote above, states that a property’s identity is determined by the causal powers of the objects which instantiate it. The thought is that it is the *properties* of objects which empower those objects to behave in certain ways. Therefore, we can individuate a property or determine its identity by its causal contributions to objects. Take, for instance, the property of being ten stone. If a boulder instantiates this property, then it will have the power to do certain things in virtue of being ten stone. It will, amongst other things, be able to fall to the ground when dropped, smash fragile items, act as a doorstep and so on. Of course, often one property will not, on its own, be sufficient for an object to have a certain causal power. The object will only have that power conditionally upon it instantiating some other properties besides the one in question. But this doesn’t matter, as properties can still be defined via the causal powers they are able to contribute to.

Some philosophers are wary of this talk of causal powers. Humeans in particular may suspect that it is illicitly introducing the most “obscure and uncertain” ideas as “power, force, energy or necessary connection”.¹ Mackie, for instance, argues that it is not “at all helpful to say that things have causal powers: the concept of power needs to be elucidated in terms of causation rather than causation in terms of powers”

¹ Hume, 1975 §VII, part I.

(1974, p.86). How is this transition supposed to be made? Instead of talking about causal powers, we can speak of what causal relations an object would be engaged in, given certain circumstances. This interpretation of the causal power talk is still problematic, as it does not rid us of modal notions such as, this object *would* have caused this had it been in these circumstances. But if this terminology is more to our taste, then rather than defining properties via the causal powers which they give rise to, we can characterise them via the causal relations which they can contribute to. So, for instance, if I drop a ten stone boulder onto a fragile vase, there will be a causal relation between my dropping the boulder and the breaking of the vase. Again, that this causal relation occurs (or that the causal power is manifested) is not just conditional upon the boulder being ten stone. The vase wouldn't have broken in a gravity free zone or if there had been an obstacle shielding the vase from the boulder. But this poses no threat. For a property can be characterised by all the complicated and particular causal relations it can contribute to, rather than by that which it alone is responsible for. I shall call this nexus of causal relations that a property contributes to the *causal profile* of a property.²

According to CTP then, the causal profile of a property provides an analysis of properties. But what sort of analysis does it provide? Answers to this question vary quite widely. One suggestion, originally put forward by Shoemaker, is to say that the nexus of causal interactions described in the causal profile just is what properties are. So the thought is that “properties just are clusters of conditional powers” (1980a, p.213). However, many, including Shoemaker in more recent papers, have rejected this claim, arguing instead that the causal profile serves to individuate properties. Shoemaker, for instance, writes, “properties are individuated by their causal features – by what contribution they make to the causal powers of the things that have them” (1998a, p.61). The idea here is that the essence of a property is captured by its causal profile. Therefore, a property wouldn't be the property it was if it didn't give rise to the causal nexus described by its causal profile.

The issue of how CTP should be interpreted will be considered shortly. For the purposes of this taster, however, it is enough to note that CTP theorists believe that the causal profile of a property is somehow crucial to understanding the nature of that

² Elder (2001) and Hawthorne (2001) use this terminology.

property. The key characterisation of a property is not given by something which is intrinsic to that property. Instead, CTP provides properties with an extrinsic analysis, since properties are characterised by their actual and possible relations with entities external to themselves. This pits CTP against those accounts of properties which claim that properties are essentially characterised by their intrinsic natures.³

Wherein lies the attraction of CTP? Two aspects of the account initially drew my interest. First, CTP's analysis of properties is different from those usually offered. Most accounts concentrate on the question, 'What are properties?'. Are they, for instance, universals, or sets of tropes, or sets of possible particulars? CTP, however, begins with a different question, 'What are properties like?'. What is it, for instance, to be the property of F-ness rather than, say, the property of G-ness or H-ness? CTP's answer to these questions promises to make properties accessible to us, as we can discover, through empirical investigation, what the identities of properties are by seeing what causal contributions they make. In contrast, a property's intrinsic nature seems hidden from the eyes of science. It is far from clear how we could ever go about unveiling these intrinsic natures, so the identities of properties are obscured to us.

The initial appeal of CTP does not end there, however. Another reason to be interested in this theory is because it links, in a way not done previously, the notions of a property, a causal power and a causal relation. Although it has often been observed that there is a close connection between properties and causation, little has been said about what exactly this relationship is. CTP attempts to rectify this, by offering an account of this relationship which explains why the notions of causation and property have to be understood in tandem. This, I hope to show, has positive implications for both an analysis of properties and of causation. For it results in both a well-worked out analysis of properties and a plausible causal ontology. Before embarking on the task of showing why this is, however, we must first try to get a clear sense of what sort of theory or analysis CTP is purporting to offer. So in the next section, I shall look at some of the different theses which are classified as causal theories of properties. After which, I shall discuss, in some detail, different ways of trying to make sense of what's on offer.

³ For an account of this kind, see Armstrong (1999a).

1.2 Varieties of CTP

The increasing proliferation of causal theories of properties⁴ (and interpretations of those theories⁵) has made it impractical to outline each theory in its turn. Instead, I shall try to offer a general overview, which picks out the main theses and trends that have emerged.

One divide which runs through causal theories of properties, and is evident in Shoemaker's corpus, is between what I shall call strong and weak versions of CTP. The contrast between these two positions was mentioned earlier. The key theses expounded by strong and weak CTP can be stated as follows:

Strong CTP: the causal profile of a property is all that there is to a property.

Weak CTP: the causal profile of a property provides us with the most crucial characterisation of that property.

The most notable example of strong CTP is found in Shoemaker's 1980 papers. There, he seems to be arguing that there is nothing more to a property than what it does. CTP is thus presented as a metaphysical thesis about the nature of properties, which states that "all properties are causal powers" (Shoemaker 1980a p.210). Many commentators on CTP have picked up on this strand and thus present CTP in its strong form. Armstrong, for instance, identifies the core of CTP with the claim that "properties are exhausted by their causal role" (1999a, p.26). While Hawthorne characterises CTP as denying that "there is something to a property – call it its quiddity – over and above its causal profile" (2001, p.362).⁶ Both these outlines of CTP strongly suggest that the causal profile exhausts the nature of a property. So it looks as if they are expressing a version of strong CTP.

Some variant of weak CTP, however, has proved to be the most popular position to adopt. Instead of plumping for the fascinating yet rather puzzling claim that

⁴ See, for example, Achinstein (1974), Shoemaker (1980a and b, 1981, 1998a, 1999), Swyer (1982), Fales (1990), Ellis and Lierse (1994), Ellis (2001). Mellor and Oliver (1997) and Campbell (manuscript copy) offer related views.

⁵ See, for example, Armstrong (1997, 1999a) Elder (2001), Hawthorne (2001) and Durham (2002).

⁶ It is not clear whether or not Hawthorne endorses CTP, therefore, I shall refer to him as a commentator, even though his paper does seem to be sympathetic towards CTP.

properties are just (conditional⁷) causal powers, tamer suggestions have been put forward. These usually involve the claim that the causal profile is essential to the property. It is this which individuates it or provides it with a criterion of identity. Swoyer, for instance, argues that “the essential features of a property lie in its relationships to other properties” (1982, p.214). But he does not rule out the possibility that there could be something over and above these causal features. He writes, “if we subtract the active and passive dispositional powers that a property bestows upon its instances, whatever is left would not enable it to affect our sensory apparatus” (p.214). So Swoyer seems to be endorsing a weak version of CTP. The same point is brought out clearly in Shoemaker’s later papers and in Campbell’s discussion.⁸ Instead of talking about properties *being* causal powers, Shoemaker moves to the claim that they are individuated by their causal powers. He writes, “properties are individuated by the contribution they make to the causal powers of their subjects” (1999, p.297). Any suggestion that this causal contribution is all that there is to a property has been removed. Campbell more explicitly denies the strong CTP’s thesis, arguing that “shape properties are individuated by their causal roles – so that the dispositions which a shape property grounds are individuated of that property – but...nonetheless the shape properties are the categorical grounds of those dispositions” (manuscript, p.2). Here, again, we find a willingness to go only so far as weak CTP.

But what exactly is the difference between strong and weak CTP? How deep is the divide that separates them? As yet, we are unable to give satisfactory answers to these questions, because we’re far from understanding the content of strong and weak CTP. Philosophers, however, have come up with a number of different suggestions about what sort of theory CTP is. So in the rest of this section, I shall outline various ways we might try to make sense of the account on offer. A detailed discussion of these proposals will be postponed to later sections.

One way of understanding CTP is as a theory about dispositions. This approach was inspired by Shoemaker’s own statement of his thesis, as at one point he

⁷ This terminology is Shoemaker’s. It just expresses the idea that, usually (perhaps always), an object has to instantiate a number of different properties in order to possess a causal power. So we can say that an object instantiating property P will possess causal power C conditional on that object having a number of other properties also.

⁸ Campbell, however, only talks about shape properties in his paper (manuscript), it is not clear that he would want to say the same thing about other properties.

writes, “all properties are dispositional properties” (1980a, p.210). Although careful reading of his paper indicates we should be wary of this formulation, a number of commentators have latched onto this claim. Armstrong, for instance, calls Shoemaker’s thesis “the Dispositional Theory of Properties”. According to his interpretation, CTP renders all properties dispositions, as they all “have a nature that is exhausted by their possible (*empirically* possible) manifestations” (1997, p.76).⁹ In similar vein, Ellis (2001) states his version of CTP as a thesis about dispositions. He wishes to defend what he calls “dispositional essentialism” (p.127). This states that dispositions are amongst the most fundamental properties of the world and that the causal profiles that identify these dispositions are essential to them. Other philosophers also employ talk of dispositions in order to convey a CTP style thesis. Campbell, for instance, argues that the dispositions of a shape property are individuating of it, while Swyer claims that the “key features of properties are dispositional” (1982, p.214). The suggestion that CTP is a thesis about the nature of dispositions, then, should be taken seriously.

More recently, however, the claim that CTP is a criterion of identity for properties or a thesis about how to individuate properties has proved to be more popular. This is largely due to the change of emphasis in Shoemaker’s own papers on CTP. He writes,

I would want to reject the formulation which says that a property is a cluster of conditional causal powers. That formulation has a reductionist flavour to it. And the reduction it seems to promise is a cheat (1998a, p.64).

He thus retreats to talk of properties being *individuated* by their causal profile, as these profiles are thought to be essential to the nature of properties. Shoemaker sees himself as providing “identity conditions” (1998a, p.66) for properties, which has led commentators to interpret CTP as a criterion of identity for properties. Durham, for instance, talks about “Shoemaker’s proposed criterion of identity” (2002, p.22), while Elder interprets CTP as the view that a “property’s identity consists in its nomic profile” (2001, p.249).

⁹ In a more recent paper (1999a), Armstrong presents a slightly different, and in my view more accurate, exposition of Shoemaker’s views.

The claim that CTP is (in the case of weak CTP) or results in (in the case of strong CTP) identity or individuation conditions for properties is central to all formulations of CTP. Very basically, the individuation thesis on offer seems to be something like this: two properties are identical at all times and places, actual and possible, iff they have the same causal profile.¹⁰ This rules out the possibility of a causal profile of a property changing across places, times or possible worlds (although which aspects of that causal profile are manifested at a world can of course vary from world to world), since property P at t_1 in possible world_w is the same property as property P* at t_2 in possible world_x iff P and P* have the same causal profile. Having the same causal profile is thus claimed to be a necessary and sufficient condition for property identity.

Many versions of weak CTP seem to rest content with this claim. Furthermore, although strong CTP seems to be saying something more than this, it looks like this individuation thesis is at least implied by it. For if the causal profile of a property is all that there is to a property, a property's causal profile will provide us with necessary and sufficient conditions for the existence of that property. It seems clear, then, that CTP is interested in making claims about the identity or individuation of properties, indeed, this is a common concern which runs throughout the different CTPs. However, just to say that CTP provides us with an account of individuation or a criterion of identity leaves many matters unsettled. Is, for instance, an account of individuation the same as a criterion of identity for properties? What sort of information about the nature of properties do these kinds of analyses offer? There are no straightforward answers to these questions because there are different sorts of individuation theses and criteria of identity. So in order to understand and evaluate the thesis on offer, more needs to be said about the content of an account of individuation/criteria of identity.

The final way of trying to make sense of CTP is, again, suggested by Shoemaker. In his 1981 paper, Shoemaker argues that functionalism has an important

¹⁰ As causal profiles state what causal relations a property *could* enter into, not just what causal relations it in actual fact enters into, the individuation thesis must be thought to hold across a number of possible worlds. For it may be the case that property P, for instance, *could* bring about X in such-and-such circumstances, but this power of P's is never manifested in the actual world, so there has to be a possible world in which P does bring about X. However, as most CTP theorists claim that the causal profiles of properties are essential to them, I shall assume (until chapter seven) that CTP is committed to the claim that this individuation thesis holds across all possible worlds.

bearing on CTP, because the latter can be understood as a generalised form of the former. The idea is that rather than just defining mental properties by their functional role, all properties can be so defined. CTP theorists are thus like functionalists in that they think that mental properties, like the belief it is raining, should be defined by the inputs that can cause this belief, its actual and possible connections with other mental states and the behaviour it can result in. But they go further than this, because they claim that (most) other non-mental properties can be similarly defined. The property of being knife-shaped, for instance, can be understood via the things which cause it, its connections with other properties and the behaviour it can result in.

This idea that CTP is a generalised form of functionalism is, I think, an interesting one. But again we find that if this proposal is to throw much light on CTP, more details are required. For there isn't just one form of functionalism in the philosophy of mind, there are different versions of this thesis, each of which might serve as a model for CTP. Consequently, it isn't clear which, if any, of these functionalist accounts CTP could be a generalised form of. Despite the number of different formulations which promise to elucidate CTP, therefore, statements of this thesis still suffer from vagueness. More needs to be said about dispositions, individuation, criteria of identity and functionalism, in order to get a clear grasp of the proposals being put forward. But before turning to this, I shall first look at another issue which divides CTP theorists. This time, the question at issue is this: what is CTP a theory *of*?

1.3 The Scope of CTP

CTP is, of course, a theory about properties, but which properties fall within its scope is a debatable issue. Should CTP theorists say that all properties are subject to their analysis, or only a restricted set? It would be mistaken to think that there is a definitive answer to this question, as the scope of CTP can be taken to be as narrow or wide as is considered appropriate. We could, for instance, restrict CTP's domain to mental properties, or to the fundamental properties postulated by physics, or to some other subset of properties.¹¹ But the boldness of CTP's hypothesis, stems from the fact that

¹¹ The label 'restricted CTP' will be reserved for those theories which claim that only a small subset of properties are subject to the analysis. Functionalists within the philosophy of mind could be called

the scope of the theory is usually taken to be very broad. Shoemaker, for instance, argues that the theory holds of all “genuine properties”, while Swoyer claims that the analysis holds of all those properties discoverable by science.¹²

All CTP theorists, however, recognise the need to place some restrictions on the properties within CTP’s domain. Proponents do not, for instance, want to say that properties of mathematical entities such as *being even* or *being prime* are subject to their analysis, as these properties receive their characterisation from mathematical theories. Similarly, it would seem unwise to include properties and relations which characterise “the form of the world”,¹³ such as causality and identity, into the domain of CTP. For the task of defining these relations in terms of their causal features seems, at best, unpromising. CTP theorists thus require a way of distinguishing between those properties which are subject to CTP’s analysis from those which are not.

How might this be done? One proposal which links in with the issue about mathematical properties, is to say that CTP is an analysis of concrete, rather than abstract, properties. This suggestion recommends itself because it seems to provide a principled way of characterising CTP’s domain. Concrete properties, whatever they are, seem very different things from abstract properties. Thus, the restriction of CTP’s scope to just one of these sets of properties looks *prima facie* reasonable. This proposal may strike some as obviously false, however, as many believe that all properties are abstract entities, so there are no such entities as concrete properties which CTP is supposedly true of. As I have no wish to rule out the view that properties or universals are abstract entities, this is a problem for the proposal. But we can get round it by talking instead about instances or instantiations of properties. We can say that if the instantiations or instances of a property are concrete (such as, the greenness of *this* jumper), then CTP applies to that property. If, on the other hand, a property’s instances are abstract (such as, the evenness of *this* number), then that property falls outside the scope of CTP.

restricted CTP theorists (see chapter two), as could philosophers like Ellis and Lierse (1994). They believe that science will have to postulate some fundamental dispositional and categorical properties. It is these fundamental dispositional properties which, on their view, are subject to CTP’s analysis.

¹² See Shoemaker 1980a, p.297 and Swoyer 1982 p.205.

¹³ Hawthorne 2001 p.373. The form of the world is captured by a world’s structural properties. These structural properties can then be said to structure the ‘matter’ of the world.

Understanding the proposal in this way does not prevent us from thinking of all properties as abstract entities. Moreover, the claim that some property instances are concrete seems far more appealing. It is implausible to claim that property instances, such as the greenness of *this* jumper or the taste of *that* cake, are abstract entities. For they are entities which have a spatiotemporal location, are empirically available to us, are cited in causal statements and do not enjoy necessary existence. Hence, they differ very markedly from the most uncontroversial example of abstract entities, namely, numbers. If, on the other hand, we consider property instances of evenness, say that displayed by the number four, then it does seem plausible to call this property instance abstract. Since it does not seem to be empirically available to us, nor is it cited in causal statements, and so on. But how, more precisely, could we differentiate between concrete and abstract property instances?

Probably the simplest way of distinguishing between the two is to say something like this: abstract property instances are all and only had by abstract objects, while concrete property instances are all and only had by concrete objects. Hale, for instance, states that “those properties and relations are abstract which have application only among abstract objects” (1987, p.45). This suggestion, however, faces two problems. The first arises because of the obscurity surrounding the distinction between abstract and concrete objects. When faced with the variety of suggestions, it is tempting to draw Lewis’s conclusion that “it is not clear at all...what philosophers mean when they speak of ‘concrete’ and ‘abstract’” (1986a, p.81). Let’s put this worry aside here, however, and suppose that Burgess and Rosen are right to suppose that this pessimism is unwarranted because of the “broad consensus” that has emerged regarding the sorts of entities which fit into these categories.¹⁴

The second problem stems from the proposal’s commitment to the claim that all the property instances of abstract objects are abstract, and all the property instances of concrete objects are concrete. For perhaps there are some concrete properties of abstract entities and some abstract properties of concrete entities. At first sight this may seem rather implausible, but there are possible examples. Consider, for instance,

¹⁴ See Burgess and Rosen 1996, p. 13.

a borderline case of an abstract object, namely, the play 'Hamlet'.¹⁵ Is it clear that this play has no concrete properties? Couldn't we argue that it has such concrete properties/relations as having being performed in front of Queen Elizabeth I, or being the most published piece of literature? Maybe all such cases should be dismissed because they are borderline, but there seem to be more compelling examples of concrete objects displaying abstract property instances. Take, for instance, geometric properties such as being square. These figure in pure geometry, but they are also instantiated in everyday objects around us. We may want to say, for instance, that this table's squareness is an abstract property instance which is displayed by a concrete object. If we do say this, however, then we cannot define abstract property instances as those which are solely possessed by abstract objects.

Rather than trying to define a concrete/abstract property instance as one which is had by a concrete/abstract object, we could instead try to directly pick out those distinguishing features of concrete/abstract property instances. To help us, it may be useful to look at the suggested analyses of the distinction between abstract and concrete *objects*. Lewis offers quite a comprehensive list of different ways of distinguishing between concrete and abstract objects.¹⁶ But of the four ways cited, only one seems viable for our purposes. The way of example does not offer a principled way of demarcating the abstract from the concrete. The way of conflation uses pre-existing distinctions between universals and particulars, for instance, or sets and individuals, but it does not offer a way of picking out all those entities which are distinctly abstract or concrete. According to the way of abstraction, abstract entities are abstractions from concrete entities. They are the product of less specific descriptions of concrete entities. This distinction is unsuitable for these purposes, however, as it renders all property instances abstract. If, for instance, we pick out the redness of an object, we are abstracting away from the other features of the object, and thus offering a partial description of that object.¹⁷ Left, then, is Lewis's negative way,

¹⁵ Burgess and Rosen (1996), for example, do take novels to be examples of non-paradigmatic abstract objects.

¹⁶ See Lewis 1986a §1.7.

¹⁷ This sense of abstract captures the meaning of those trope theorists who call tropes or property instances 'abstract particulars'. It seems unsuitable for the purposes of defining a fundamental ontological category, however. While it captures the etymology of the word and is, what Rosen and Burgess call, "the way of history" (1996, p.17), originally, it applied to our representations of entities rather than to the entities themselves. It makes more sense in this context, because we can imagine

which tries to characterise abstract entities by what they're not. It is often said, for instance, that abstract entities are not spatiotemporal entities or do not enter into causal interactions.¹⁸

Of these two suggestions, an analysis along the second line looks the most promising, at least for our purposes, since it seems to capture CTP's concerns. CTP theorists clearly only intend to offer an analysis of those properties which can contribute to the causal powers of particulars. So if we define abstract property instances as those which cannot contribute in any way to the causal nexus of the universe, they will thereby be excluded from CTP's scope. Concrete property instances, on the other hand, can be characterised by the fact that they do contribute to the causal powers of particulars and so to the causal nexus of the universe. Therefore, this type of property is just the sort CTP should be about.

If we demarcate abstract from concrete property instances in this way, then it will be possible for concrete objects to display abstract property instances and abstract objects to display concrete property instances. (Although if the latter is possible, we couldn't define abstract objects as those which never enter into causal interactions.) Unfortunately, however, we still face the problem of what to do about those properties which seem to straddle the concrete/abstract divide. If an object instantiates the property of being square, for instance, then it will clearly be able to enter into a variety of causal interactions in virtue of instantiating this property. Therefore, according to the criterion just offered, it should be subject to CTP. However, the property also appears in mathematical theories and it seems clear that we can identify the property of squareness independent of its causal profile. For we can say that an object instantiates the property of being square iff it is an object which has four equal sides and four right angles. Moreover, it seems much more intuitive to claim that the

abstracting away from the complex details of our experience of an object and homing in on just one of its particular aspects.

¹⁸ Another notable account is offered by Dummett (1973). He argues that we can distinguish between abstract and concrete objects by the way we understand concrete as opposed to abstract singular terms. Concrete objects can be the possible objects of ostension and thus concrete terms get a reference in this way. Abstract objects, on the other hand, have to be referred to as the value of some functional expression, such as 'the shape of X'. It is difficult to see how we could extend this criterion to cover all property instances, however. Is, for instance, Alice's happiness something which is a possible object of ostension or a value of an appropriate function? What of the charm of a quark or the charge of an electron? I shall thus leave this suggestion aside here.

property of being square is individuated by this characterisation rather than by its causal profile.

As restricting the scope of CTP to concrete (understood as causal) property instances accords so well with the concerns of CTP, I think we need to face this problem head on. What should CTP theorists say about these cases? One option would be to argue that only those properties whose instances are all concrete fall within CTP's scope. The idea being that those properties which figure in mathematical theories will have instances which are abstract as well as concrete. So properties like being square will not be subject to CTP. This seems rather unsatisfactory, however. For shapes, being entities which bestow causal powers onto objects, do seem to be the sorts of things which CTP should be true of. That these properties figure in mathematical theories doesn't obviously change this fact.

Perhaps a better response can be formulated if we keep hold of the claim that all properties, which have causally efficacious instances, are subject to CTP. The difficulty arises because some properties are both instantiated by physical objects and appear in mathematical theories. It looks, then, as if these properties fall partly within the scope of CTP (insofar as they are causally efficacious entities of physical objects) and partly outside of it (insofar as they figure in mathematical theories). But CTP theorists can use this fact to explain why it seems plausible to claim that properties, such as being square, can be characterised independently of their causal profile – insofar as they are purely mathematical properties, they can. But this still leaves the question: why should we accept that properties, such as being square, qua instantiated by physical objects, are subject to CTP?

This is a very difficult question, which requires more consideration than can be given here. However, there is some reason to hope that an answer may be forthcoming. Campbell's account of spatial properties (1994), for instance, could perhaps be utilised by CTP theorists to explain why properties like being square, qua instantiated by physical properties, require a causal profile. Campbell argues that pure geometry (a purely formal exercise in mathematical computation) is turned into applied geometry (a body of doctrine about the world in which we live) by connecting spatial properties with physical ones. He writes, "what turns one into the other is the assignment of some physical meaning to the spatial concepts, for example, the

identification of a straight line as the path of a light *in vacuo*” (1994, p.25).¹⁹ A CTP theorist could try offering a similar line of argument to defend their position. They could claim that CTP is the right analysis of physically instantiated shapes because the causal relations which shapes bear to other entities are what give the mathematical entities of pure geometry their physical significance.

On this proposal then, CTP’s scope is said to include properties whose instances are *all* able to enter into causal interactions and exclude properties whose instances are *all* unable to enter into causal interactions. However, some properties, in particular, geometric properties, do not fit this neat distinction, since they have both concrete and abstract applications. In these cases, I have suggested that we should say that the property falls partly within the scope of CTP and partly outside of it. The property qua mathematical theory does not fall under the scope of CTP, as it has an essential non-causal characterisation. But the property qua instantiated by a physical object does fall under the scope of CTP, because the property’s causal profile is what gives the property its physical significance.

Another, less worrying, objection to the proposed scope of CTP is put forward by Rosenberg. He argues that, as it stands, CTP is an extremely exclusive analysis, since it only holds of a very limited number of properties. He writes, “one cannot know that there is a cluster of causal powers associated with any predicate actually in use, because of the inadequacy and incompleteness of contemporary science” (1984, p.84). Rosenberg’s worry is that if CTP only applies to causal properties, i.e. ones which confer genuine causal powers onto their particulars, we cannot be sure that any of our predicates succeed in latching onto these properties. For only those properties which figure in the final inventory of science are able to confer genuine causal powers onto their particulars, and “true total science does not wear this label on its sleeve” (1984, p.82). Consequently, as we can assume that nearly all of our more ordinary properties, such as redness, will not be on that final inventory, CTP will not be true of most properties.

¹⁹ For more on this see Campbell (1995) and Ludwig (1995).

One possible way of developing the proposal put forward here is to maintain, with Armstrong and Mellor,²⁰ that the only properties which exist are those sparse ones postulated by science, as only these are required to account for genuine resemblances between particulars. If this sparse view is accepted, CTP would still be a theory of all concrete properties.²¹ It's just that this category wouldn't include those entities (if there are any) named by more ordinary predicates such as 'is red'. CTP theorists are not committed to this view, however. They may, for instance, believe that there are emergent properties with irreducible casual powers falling outside of science's domain. So in order to account for the genuine resemblances between particulars, we need to postulate properties which are not specifiable by science. Even if CTP theorists deny this, however, they still needn't claim that properties like redness fall outside the scope of CTP. CTP theorists can agree, with Rosenberg, that the property of redness does not confer a fundamental causal power onto its particulars, since its causal efficacy is bestowed upon it by other properties. But nevertheless they can still maintain that redness is a causal property. For redness does have a unique causal profile which identifies it and sets it apart from other properties. So we can say that redness is a causal property, as any object which instantiates it has a unique causal profile in virtue of its being red.

It may be thought that this fails to get to the heart of the matter. Redness isn't a causal property, since it is wrong to say that an object has the power to create such-and-such visual sensations in virtue of its being red. Rather, the object's power is conferred upon it by other properties. This response, however, imposes stringent conditions on causal properties which are not obligatory. No doubt it's true that redness isn't a *fundamental* causal property, since science tells us that its causal power to produce certain visual sensations is due to further properties of light waves, our visual systems etc. But this doesn't mean that redness should be denied the status of a causal property. We can say, perfectly truly, that looking at this *red* post-box caused

²⁰ See Mellor (1991) and Armstrong (1997). Although, unlike Mellor, Armstrong does accept the existence of properties which are conjunctions of sparse properties.

²¹ However, CTP theorists could still say, with Lewis (1983b), that there is another, more abundant, conception of properties. But as these do not ground causal similarities in their objects, but rather have a semantic function, they are not the concern of CTP (for more on this, see §2.2). Therefore, although Rosenberg's claim that CTP is not applicable to all concrete properties is misleading, in one sense, it can be thought of as correct. For CTP theorists could maintain that there is this more abundant conception of properties, and these properties should not be analysed via CTP.

my visual impression of redness. While we may complain that this isn't an informative causal explanation - a better explanation would proceed by appealing to the properties of light waves, our visual system etc. Still, redness can be seen as the cause of my sensation of redness. For we can say that redness inherits the causal powers of those properties which, scientific theory tells us, are responsible for sensations of redness.

What does this 'inheritance' amount to? The idea can be spelt out a little more clearly if we assume that redness is a second-order property (i.e. a property which is characterised by the causal roles of first-order properties), as then we can appeal to Kim's causal inheritance principle. This states that "the causal powers of an instance of a second-order property are identical with (or a subset of) the causal powers of the first-order realiser that is instantiated on that occasion" (1998, p.116). Given this conception of redness as a second-order property, we can say that redness is a causal property, since it inherits the causal powers of those more fundamental properties that realise it. These notions of realisation and second-order properties are quite complex and take us into the subject matter of chapter two. But it is useful to briefly mention this conception here, as it shows how redness can be thought of as a causal property, even once we've granted Rosenberg's claim that it isn't among science's list of ultimate movers and shakers.

By saying that CTP holds of *all* causal properties, therefore, (whether or not those properties are fundamental) we are able to respect the boldness of CTP. Moreover, we can do this without having to say that properties, such as *the evenness of four*, are somehow not genuine. The current proposal also has the advantage of placing the emphasis in the right place. For CTP is a theory about those properties which have causal profiles. Therefore, at least at its broadest, all such causal properties should fall within CTP's scope. In what follows then, unless otherwise stated, I shall assume that CTP is a theory about all causal properties. Now, however, it's time to return to the question: what kind of theory is CTP?

1.4 The Dispositions Debate

Earlier we saw that many philosophers believe there to be a strong connection between CTP and the dispositions debate. Indeed, some have even suggested that CTP can be stated as the view that all properties are dispositional. Campbell, for instance,

characterises Shoemaker's position as "the view that tries to dispense with the idea of appealing to categorical properties...we can conceive of properties...only as complexes of dispositions".²² In this section, I shall examine the relation between the dispositions debate and CTP, in order to show that the connection between the two is much weaker than is often presumed.

Within the dispositions debate, one issue which looms large concerns whether or not there is a conceptual distinction between categorical and dispositional *predicates*. The most obvious way of trying to draw the distinction between dispositional and categorical ascriptions is by saying that dispositional ascriptions, unlike categorical ones, entail a particular conditional (or conditionals).²³ To ascribe the dispositional predicate of solubility to the object X, for instance, is to say that if X is put in liquid, then X will dissolve. The same does not hold true of ascriptions of categorical predicates, however. Although it is true that if we correctly ascribe the categorical predicate of 'is copper' to an object Y, it will be a good conductor of heat, electricity etc, the categorical ascription does not entail the conditional, 'if Y is heated, then Y will transmit heat/electricity easily'.

Problems have been raised for this simple conditional analysis. In particular, Martin offers a damning objection to this view. Consider, for instance, the dispositional predicate 'is live'. If the simple conditional analysis is correct, the meaning of this predicate is given by a conditional or set of conditionals. Let's suppose that the relevant conditional for this predicate is this: 'if X is live, then electrical current will flow from the wire to the conductor'. Martin now asks us to imagine that there is an "electro-fink" (1994 p.5), which detects when a wire is about to be touched by a conductor and immediately either makes the live X dead, or the dead X live. In this situation, it looks like we would still be justified in applying the predicate 'is live' to X, even though the conditional is not true of the object. Therefore, the predicate 'is live' need not always entail the conditional which is supposed to define it.

²² See Campbell 2002a, p.235. See also Shoemaker 1980a p.210 and Armstrong 1997, p.76.

²³ See, for example, Shoemaker (1980a), Jackson, Pargetter and Prior (1982) and Place (1996). Mackie (1973) and Mumford (1998) also offer accounts which have this claim at its core, although both analyses are a little more sophisticated.

This counterexample to the simple conditional analysis has led to a variety of reactions. Many believe that a more sophisticated variant of the standard analysis can be offered, but some think we may need to abandon it altogether, and either search for a new approach or conclude that there is no conceptual distinction between categorical and dispositional predicates after all.²⁴ Whatever the case may be, however, fortunately CTP theorists can bypass the dispute altogether, as they are not committed either way on this issue. A CTP theorist could argue that there is a conceptual distinction between dispositional and categorical predicates, or they could deny this claim.

It is not difficult to see that the second option is consistent with CTP. Granted an unrestricted version of this account, the claim is that *all* properties (with concrete instances) can be identified with or individuated by their causal profile. CTP does not postulate any significant division within this category of properties. Consequently, a proponent of this view is under no obligation to claim that there is a meaningful conceptual distinction between categorical and dispositional predicates, which mirrors this divide.

Given that CTP is an analysis of all properties with concrete instances, however, the fact that CTP theorists can allow that there is a conceptual distinction between categorical and dispositional predicates, may strike some as rather surprising. But such a combination of views is held by Shoemaker - the original CTP theorist. He writes,

Sometimes it belongs to the meaning, or sense, of a predicate that if it is true of a thing then under certain circumstances the thing will undergo certain changes or will produce certain changes in other things. This is true of what are standardly counted as dispositional predicates (1980a, p.210).

Shoemaker goes on to say that this isn't true of categorical predicates, hence there is a distinction between dispositional and categorical predicates. This thesis is open to all CTP theorists, because although they claim that there is a causal profile which individuates or exhausts the nature of a property, this causal profile is not said to be conceptually linked to that property. In other words, no where is it claimed that the

²⁴ Lewis (1997) and Mellor (2000) offer more sophisticated variations of the standard conditional analysis. While Mumford attributes the view that there is no conceptual distinction to Mellor (see Mellor 1974).

causal profile of a property forms part of the meaning of the predicate that picks out that property. That copper conducts electricity, for instance, is something which we discover via empirical methods. It is not something which is drawn out of our concept of copper and so, according to CTP theorists, it is not analytically true that copper conducts electricity.

Some may worry that CTP theorists can only make a distinction between categorical and dispositional predicates given the problematic, simple conditional analysis. This isn't the case: Consider, for instance, Mumford's account of the distinction.²⁵ He writes,

Dispositional ascriptions are ascriptions of properties that occupy a particular functional role as a matter of conceptual necessity and have particular shape or structure characterisations only a posteriori. Categorical ascriptions are ascriptions of shapes and structures which have particular functional roles only a posteriori (1998, p.77).

CTP theorists can endorse this account, because they do not claim that all properties have their "functional role as a matter of conceptual necessity". If a property's causal profile is essential to it, then there is a necessary connection between that property and its causal profile, but the necessity in question is metaphysical not conceptual. Although it's not practical to see whether CTP theorists could accept every account of the conceptual distinction, there is reason to suspect that we can generalise this conclusion to other cases. For CTP is a theory about the metaphysical nature of properties, therefore it would be quite odd if it did dictate answers to these conceptual questions.

The other main issue discussed in the debate about dispositions concerns whether there is an ontological distinction between categorical and dispositional *properties*. As the focus now turns to the nature of properties, this is where we might expect the two debates to collide. This suspicion is strengthened by the fact that it seems possible to view strong and weak CTP as variants of two preexisting theses about dispositions. Strong CTP could be characterised by what Mumford calls "dispositional eliminativism" (1998, p.178). This eliminates categorical properties in favour of dispositional properties, so it claims that "all properties are dispositional"

²⁵ For Mumford's response to Martin's criticism see 1998 §4.7-4.8.

(Mumford, 1998 p.179). The similarity between this and Shoemaker's strong 1980a formulation of CTP is obviously striking. Weak CTP, on the other hand, could be characterised by Martin's two aspect theory (1993, 1996). Two aspect theory states that properties must have both a categorical or "qualitative" side and a dispositional side. A property is thus not exhausted by its dispositional aspect or causal profile, as strong CTP claims. But the categorical or qualitative aspect of a property has to give rise to a particular dispositional aspect or causal profile, so a property's causal profile will be essential to it, as weak CTP claims.²⁶

There does seem, then, to be something in the claim that different versions of CTP can be understood as positions within the dispositions debate. In what follows, however, I shall argue that interpreting CTP as a thesis about dispositions does little to further our understanding of CTP, due to the obscurity surrounding the notions of categorical and dispositional properties. Moreover, at least on some ways of distinguishing categorical from dispositional properties, CTP cannot be stated as a thesis about dispositions without misrepresenting it. So CTP needs to be distanced from the debate about dispositions.

Probably the most perspicuous way of delineating the so-called categorical from the dispositional properties, at least for those who accept that there is a conceptual distinction between categorical and dispositional predicates, is by saying that dispositional properties are those which are picked out by dispositional predicates and categorical properties are those which are picked out by categorical predicates. This proposal leaves a lot unsaid, however, as it doesn't tell us what the properties specified by the dispositional/categorical predicates are like. For instance, are the properties referred to by the categorical properties radically different from those referred to by the dispositional predicates (the position of the property dualist²⁷)? Or, are the properties picked out of the same kind (the position of the property monist²⁸)?

With regards to this issue, the hands of strong unrestricted CTP theorists are tied. For strong CTP claims that *all* properties (within the scope defined) are

²⁶ Martin prefers to talk of the "qualitative" rather than categorical aspects of a property, as he thinks that the latter biases us against dispositions (see, for instance, 1996 p.74). It is not completely clear whether Martin thinks that the categorical or qualitative aspect has to give rise to the dispositional aspect that it has, but this is perhaps the best way of interpreting what he says.

²⁷ See, for example, Place (1996) and Martin (1996).

²⁸ See, for example, Mackie (1973), Mellor (1974), Armstrong (1996) and Mumford (1998).

exhausted by their causal profiles. So there isn't room to claim that there are two fundamentally different types of properties with concrete instances. It may be thought that this is mistaken. In Shoemaker's statement of strong CTP, he allows that there is a "rough correspondence" (1980a, p.211) between dispositional/categorical predicates and powers/properties. As the latter is viewed as a distinction between kinds of properties, it looks as if Shoemaker is postulating an ontological distinction. But closer inspection of his characterisation of powers and properties, reveals that this isn't the case. Properties and powers are interdefinable: properties are characterised as "second-order powers, they are powers to produce first-order powers" (1980a, p.212), while powers are defined in terms of sets of properties. Therefore, although Shoemaker does talk about different kinds of properties, it would be a mistake to think that there are two radically different sorts of properties – powers and properties – which divide reality.

There is room to manoeuvre on this issue, however, if we adopt a version of weak or restricted CTP. It is obvious that the latter view gives us the luxury (if we want it) of restricting CTP's analysis to the properties picked out by the dispositional or categorical predicates. But someone who endorsed weak CTP could also adopt property dualism. For so long as they maintained that the categorical aspect of a property necessarily gives rise to its dispositional aspect or causal profile, they could say, with Martin, that there is a fundamental distinction between the categorical and dispositional elements of reality.²⁹

Let's leave aside these alternatives for now, however, and concentrate on strong unrestricted CTP. This commits us to the claim that there is only one type of property corresponding to the conceptual categorical/dispositional divide, but how should these properties be viewed? It may be thought that the entities picked out by the categorical predicates serve to reduce or eliminate the entities picked out by the dispositional predicates. So the entity referred to by the predicate 'is fragile' is nothing over and above certain other properties picked out by the categorical predicates. This is one way of understanding, what Mumford calls, categorical reductionism or eliminativism, the view that dispositional properties are reducible to, or eliminated by,

²⁹ If we identified weak CTP with Martin's two aspect view, then weak CTP theorists would have to be property dualists. But this, I shall argue presently, isn't correct, so we need to distance the two views.

categorical properties.³⁰ Alternatively, we may think that the entities picked out by dispositional predicates serve to reduce or eliminate the entities picked out by categorical properties. This is one way of interpreting, what Mumford calls, dispositional reductionism or eliminativism, the view that the categorical properties are reducible to, or eliminated by, dispositional properties.³¹ It is this thesis which is often thought to state the position of strong CTP. Shoemaker, in his earlier paper, summarises his account with the slogan “all properties are dispositional properties” (1980a, p.210), so surely strong CTP can be viewed as an instance of dispositional eliminativism?

This is not the case, however, at least given the conception of a categorical/dispositional property outlined above. In order to see this, consider Shoemaker’s statement of strong CTP. Shoemaker claims that there is a rough correspondence between dispositional predicates and powers, on the one hand, and categorical predicates and properties, on the other. The entities which dispositional predicates pick out are thus, by and large, powers rather than properties. Powers, on Shoemaker’s view, rather than being properties themselves, are realised by sets of properties. Consequently, if dispositional properties are defined as those which are specified by dispositional predicates, this account is an instance of categorical reductionism rather than dispositional eliminativism. For Shoemaker allows that there are categorical properties and, moreover, that these categorical properties reduce the dispositional properties, as the latter are realised by sets of the former.

This suffices to show that it is misleading to assimilate strong CTP with dispositional eliminativism, since at least on one reading of the dispositional-categorical divide, strong CTP theorists are not committed to such a thesis. Is the same true of the identification of weak CTP with Martin’s view? Yes, as given this understanding of the categorical/dispositional distinction, weak CTP isn’t committed to property dualism, as two aspect theory is. Weak CTP theorists can endorse property monism and claim that all the properties picked out by categorical predicates eliminate or reduce the properties picked out by the dispositional predicates (or vice versa). For

³⁰ This taxonomy of positions is offered by Mumford 1998, Ch. 8. He classifies Quine as a categorical eliminativist and Armstrong as a categorical reductionist.

³¹ This view is often associated with Mellor’s 1974 account (see, for instance, Mumford 1997).

nothing has been said about the nature of the properties picked out by the categorical/dispositional predicates, so weak CTP theorists are free to say that only the properties picked out by the categorical predicates count as genuine or irreducible properties (or vice versa). Consequently, it is misleading to assimilate weak CTP with Martin's view or strong CTP with dispositional eliminativism.

Are there any other ways of spelling out the contrast between dispositional and categorical properties which forge a closer connection between CTP and the dispositions debate? Answers to this question are hindered by the fact that so many discussions of this distinction are inadequate. The way of example is frequently appealed to, so properties such as squareness or roundness are given as examples of categorical properties and fragility or solubility as examples of dispositional properties. Armstrong, for instance, in one paper writes that the disposition of fragility can be identified with a "purely categorical property of the glass: such things as the molecular structure of the glass" (1996, p.16). Nothing more in the way of elucidation of the notion of a categorical property is offered.

Other attempts to clarify the distinction are also extremely obscure. Martin, for instance, argues that dispositions are those properties which are "pure potency", whereas categorical properties are "potency-free" (1996, p.74). It is difficult to know what to make of this. All non-abstract properties, including paradigmatically categorical ones, have causal effects (as Martin himself notes), and I have to confess I have no grasp on what pure potency amounts to. Place offers a similar notion of a categorical property, writing, "a categorical property is one which consists entirely in what exists at the moment or period of time to which reference is made, to the exclusion of anything that might exist or have existed at some other point in time" (1996, p.21). But, again, on this definition it seems unlikely that any categorical properties exist since, as Mellor (1974) notes, ascriptions of categorical properties can and do justify conditionals which make reference to other possible moments in time.

In another paper, Armstrong (1997) offers a slightly more illuminating suggestion. He argues that categorical properties have "a nature which is self-contained, distinct from the powers that they bestow", while dispositional properties have "a nature that essentially looks beyond the particulars they qualify, outward to potential interactions with further particulars, and where this nature is exhausted by

these potential interactions” (p.69). Again, this fails to provide us with a clear distinction between categorical and dispositional properties. However, it does look as if this conception of a categorical property is at odds with both weak and strong CTP, as these state that all properties (with concrete instances) are characterised by their causal profiles. Consequently, CTP theorists cannot think of categorical properties as having a nature distinct and identifiable independent of the powers they bestow.

Given Armstrong’s characterisation of categorical properties, CTP theorists cannot be thought of as categorical reductionists, as they deny that any subset of properties satisfies Armstrong’s characterisation. However, it is not clear that CTP theorists have to accept the label of dispositional eliminativism, given Armstrong’s description. He views dispositions as those properties whose natures are exhausted by their potential interactions, but it is not clear that a CTP theorist would endorse this conception of properties. This is certainly true of weak CTP theorists, as they only claim that the causal profile individuates or is essential to that property. However, it might also be questioned by strong CTP theorists, as once the claim that the causal profile exhausts the nature of a property has been spelt out, something rather different may be implied than what Armstrong supposes.

A more perspicuous account of the dispositional/categorical divide, which again renders all CTP theorists dispositional eliminativists, is offered by Ellis. He argues that dispositional properties are “properties whose identities depend upon the behavioural dispositions they support” (2001, p.119). Categorical properties, on the other hand, are those whose identities do not depend upon their causal powers. Given this analysis, strong and weak CTP theorists alike end up being dispositional eliminativists, since they think that the identities of all concrete properties depend upon their causal profiles. It should be noted, however, that this account of the distinction will be controversial. It rules out positions such as Campbell’s, since he thinks that shape properties are categorical properties which are nevertheless individuated by their causal roles (manuscript, p.2). It also excludes the interpretation of Martin’s view outlined earlier, which states that the categorical side of a property necessarily gives rise to its dispositional side. Consequently, not everyone will favour Ellis’s way of drawing the distinction.

There is, then, at least one way of interpreting the categorical/dispositional distinction which renders not just strong CTP theorists, but all CTP theorists, dispositional eliminativists. However, it would be extremely misleading to identify these two theses, since there are other ways of drawing the distinction which do not have this consequence. More generally, I think we should say that philosophers are heading in the wrong direction if they utilise the framework of the dispositions debate to interpret CTP. So much of that debate is not relevant to CTP. Moreover, when the concerns do seem to cross, the obscurity surrounding the issue of the ontological distinction between the dispositional and the categorical makes it doubtful whether it is helpful to view CTP in this way. I am convinced by Shoemaker's claim that "the term 'dispositional' is best employed as a predicate of predicates, not of properties" (1980a, p.211).³² When we think about paradigm dispositions like fragility and solubility, they stand out not because of their causal effects, as most properties have these. Rather the difference seems to be that these causal effects are written into the meaning of the predicate. If this is correct, then there will be no overlap at all between the two debates. But even if this claim is rejected, we have seen that there are no simple connections between CTP and the dispositions debate. Therefore, the two areas need to be distanced more than is often the case.

1.5 Individuation

Shoemaker's later papers (1998a, 1999) suggest that CTP should be understood as a theory about property individuation. He writes, "properties are individuated by the contribution they make to the causal powers of their subjects" (1999, p.297). This has been picked up by a number of commentators on CTP. Elder (2001), for instance, interprets CTP as a thesis about the individuation of properties. But if we say that CTP is a thesis about how we should individuate properties, what sort of information about the nature of properties is it purporting to offer?

³² This view is also put forward by Mellor, as he writes, "Dispositionality is a feature not of properties but of predicates" (2000, p.767).

Due to the varied content of individuation theses, there is no clear-cut answer to this question.³³ However, numerous passages from Shoemaker suggest that he is attempting to analyse what it is for one property to be identical to another. For instance, he writes, “properties having the same causal features are identical” (1998a, p.64), and “properties are identical just in case they share the same total sets of casual features” (1998a, p.68). From this, it is reasonable to surmise that Shoemaker is putting forward a criterion of identity for properties. Although he does not use this terminology, there is a close link between accounts of individuation and criteria of identity. For individuation principles are often conceived of as uniquely singling out a particular type of entity from everything else in the world.³⁴ This gives us a criterion of identity, for the individuation principle tells us what constitutes that entity’s identity at a particular time. Hence, anything which satisfies that principle must be identical to that entity.

What is a criterion of identity? Very basically, it is a criterion which specifies the identity conditions for entities of a given kind. So the criterion must state a necessary and sufficient condition for a pair of entities of the kind in question to be identical. Beyond this, however, matters get more complicated. Williamson and Lowe have observed that principles which are known as criteria of identity can be split up into two different types, on the basis of their logical form.³⁵ Williamson has named them one-level and two-level criteria of identity (1990, p.145-6). One-level criteria of identity have this form:

$$1) (\forall x)(\forall y) (\Phi x \wedge \Phi y) \rightarrow (x = y \leftrightarrow Rxy)$$

This states that for all x and for all y, if x and y are of sort Φ , then x is identical with y iff x bears an equivalence relation (i.e. one which is reflexive, symmetric and

³³ Fortunately, a number of individuation theses can be quickly excluded, because they are exclusively concerned with individuals or particulars. Thiel, for instance, writes that individuation theses are about “the identification of the principles or causes which are responsible for the individuality of individuals” (1997, p.355). Similarly, Wiggins (1980) develops “a theory of the individuation of continuants, including living substances and other substances” (p.1). As properties can exist at different places simultaneously, they cannot be thought of as individuals, substances or particulars.

³⁴ See, for example, Evans 1982 p.107 and Denkel 1996, §3.2.

³⁵ Williamson and Lowe have done much to try to elucidate the central notion of a criterion of identity. See Williamson 1986, 1990 and Lowe 1989a, 1989b, 1998.

transitive) to y . The most well-used, uncontroversial example of a criterion of identity of this sort is the Axiom of Extensionality in set theory:

$$(\forall x)(\forall y) ((\text{set}(x) \wedge \text{set}(y)) \rightarrow (x = y \leftrightarrow (\forall z)(z \in x \leftrightarrow z \in y)))$$

This states that if x and y are sets, then x is identical to y iff x and y have the same members. This is a one-level criterion, because the variables x and y range over the entities for which the criterion is a criterion. So, here, x and y range over sets and it is the identity of sets which this principle is a criterion for, as it says what it is for x to be identical to y , *where x and y are both sets*.

Two-level criteria of identity have this form:

$$2) (\forall x)(\forall y) (f(x) = f(y) \leftrightarrow Rxy)$$

This differs from 1) because it does not tell us when x and y are identical, rather it states that the function f of x is identical to the function f of y iff there is an equivalence relation between the entities over which variables x and y range. So there is a domain of entities over which variables x and y range and over which the equivalence relation R is defined. These entities are then mapped onto a distinct domain by the functional term $f(x)$. Consider, for instance, Frege's famous criterion of identity for directions:

$$(\forall x)(\forall y) (d(x) = d(y) \leftrightarrow \text{parallel}(x, y))$$

This states that the direction of x is identical to the direction of y iff lines x and y are parallel. Here we see that the variables x and y range over lines. But the identity of lines is not in question, for two lines that are parallel could fail to be identical if, for instance, they were of different lengths. However, by defining an equivalence relation that holds between lines, Frege provides a necessary and sufficient condition for the identity of the objects picked out by the functional terms, namely, the directions.

Due to the significant differences between one and two-level criteria of identity, the question arises: what sort of criteria of identity – one-level or two-level – does CTP offer? Let's begin by considering the hypothesis that CTP should be understood as a one-level criterion of identity. The first thing that can be said in its

favour is that it is at least possible to formulate such a criterion. One potential formulation of CTP can be stated as follows (where z ranges over causal features and R stands for the relation of ‘...has...’):

$$(\forall x)(\forall y) (\text{property } (x) \wedge \text{property } (y)) \rightarrow (x = y \leftrightarrow (\forall z)(Rxz \leftrightarrow Ryz))$$

This says that if x and y are properties, then x is identical with y iff x and y have all the same causal features. This seems a reasonable way of interpreting Shoemaker’s thesis that “properties having the same causal features are identical” (1998a, p.64). Although, if we are to capture his further claim that the causal features of a property are essential to the identity of a property, the criterion above should be prefixed with a necessity operator.³⁶ Then it would read, necessarily, if x and y are properties, then x is identical with y iff x and y have all the same causal features.

Interpreting CTP as a one-level criterion then, does appear to capture the content of Shoemaker’s later thesis. But what function are one-level criteria supposed to have? Lombard and Durham argue that one-level criteria are required for the metaphysical categories, i.e. the “broadest, most general sorts of things there are”, entities such as “physical objects, sets, propositions, properties, events and the like” (p.25).³⁷ They are required because they provide “a crucial part of any metaphysically serious reason for thinking that there are such entities” (p.45). A criterion of identity for a kind ϕ should state “our beliefs about what it is to be an ϕ : such a thing will tell us what kind of thing an ϕ is” (p.32). If we do not have a criterion for an entity of a certain sort, we do not know what kind of thing it is. Hence, we could have little incentive to postulate entities of that sort.

It is noticeable that this characterisation of the function of one-level criteria does not fit happily with Shoemaker’s concerns. Although Shoemaker and Lombard both want these criteria to be seen as metaphysical theses, Shoemaker never considers what distinguishes properties from the other metaphysical categories. Instead, he discusses what makes a property the property of F-ness, rather than, say, the property

³⁶ See, for example, Shoemaker 1998a, p.66. There he writes “the causal features of properties are essential to them”.

³⁷ Every quote in this paragraph is taken from Lombard 1986, but Durham explicitly accepts this characterisation (see, for example, 2002, p.16). Neither Lombard nor Durham distinguish between one-level and two-level criteria of identity, all the criteria they consider are one-level.

of G-ness or H-ness. This difference in emphasis is exploited in Durham's criticisms of Shoemaker's account. Durham argues that Shoemaker's criterion is inadequate for two reasons: first, CTP fails to articulate the essence of universals, because non-universals (such as tropes) can also satisfy the role which is supposed to be definitive of universals. Second, Shoemaker explicitly excludes some properties (e.g. mathematical ones) from his analysis, hence the account fails to capture what is common to *all* properties. Both of these criticisms of Shoemaker's account, however, seem to be based on a misunderstanding of his intentions. Shoemaker is interested in offering an analysis of what makes a property the property it is, he is not telling us what it is to be a property rather than an object, event, etc. So Shoemaker is not aiming to do what Durham criticises him for failing to do.

This conclusion leaves us with two options. Either we could keep hold of the hypothesis that CTP is a one-level criterion of identity and challenge Lombard and Durham's conception of what role such criterion should play. Or, we could allow that Lombard and Durham's conception of one-level criteria is correct, and so give up the claim that CTP is such a criterion. There is independent reason to investigate the second line of response, due to an inherent weakness in all one-level criteria of identity. Earlier we saw that the variables in one-level criteria range over the entities whose identity is in question. If we are trying to form a criterion of identity for sets, for instance, then the variables x and y will range over sets. This makes the identity criteria impredicative, as the variables of quantification range over a domain which include the very entities whose identity is in question. So in order to interpret the equivalence relation (Rxy) as a one-level criterion, we need to be able to identify individual ϕ s as the values of the variables x and y , since these variables range over the entities in question. One-level criteria are thus open to the charge of circularity, as it looks like we must already be in possession of an account of what constitutes ϕ 's identity.

A debate rages concerning whether this circularity is excusable. Lowe (1998), for instance, argues that it is on the grounds that the circularity in question is not vicious. I am not going to engage in this debate here, however, since all that is required for my purposes is the rather weak and plausible claim that a criterion which is not vulnerable to this circularity is preferable to one that is. If we grant this, we have

motivation enough to see whether we can formulate another CTP criterion, which avoids such circularity.

I think we can, if we are prepared to expand CTP's usual apparatus a little and formulate a two-level criterion. In order to fit CTP into this two-level mould, however, we need to find values for the variables which could be used to provide the identity conditions for properties. What could these be? By taking property instances to be the values of the variables, we can formulate this two-level criterion for properties:

$$(\forall x)(\forall y) (\text{property } (x) = \text{property } (y)) \leftrightarrow (\forall z)(R_{xz} \leftrightarrow R_{yz})$$

(where x and y range over property instances, z over causal features and R stands for the relation of '...has...'). This states that the property of property instance x is identical to the property of property instance y iff x and y have all the same causal features. So the identity conditions for properties are given via the causal features of their instances.

Some may object that this two-level criterion is too costly, as it involves introducing property instances. Later, however, I shall argue that every ontology needs to distinguish properties from their instances. Therefore, as no contentious metaphysical analysis of property instances is being assumed, there is no reason why CTP theorists shouldn't avail themselves of this two-level criterion and avoid the circularity inherent in one-level criteria.³⁸ On two-level criteria, the items whose identity the criterion is for are distinct from the entities that are related by the equivalence relation. Consequently, as we don't have to identify properties as the values of the variables x and y , we avoid the criticism that we must already be in possession of an account of what constitutes the identity of properties.³⁹

There is, however, an important issue that remains outstanding. Earlier we saw that Shoemaker's criterion of identity doesn't cohere well with at least Lombard and

³⁸ I also think that there are other reasons for preferring this two-level criterion. I argue later, that we need to focus on a property's instances when analysing causal powers and causal relations.

³⁹ We may worry that as our criterion of identity for properties talks about property instances having the *same* causal features, this implicitly involves reference to properties or universals. So properties are being defined in terms of other properties. I think that this problem can be overcome, however. In the next chapter, I shall argue for a slightly different functional variation on this two-level criterion of identity, which makes it clear that the relevant similarity need not be analysed in terms of properties, but rather complex formulas holding of particulars. Moreover, while this complex formula does render properties' predicates interdefinable, we can suppose that they are holistically defined.

Durham's conception of one-level criteria of identity. But what kind of information do these two-level criteria purportedly provide? Prima facie at least, the resulting two-level CTP criterion fits far better with Shoemaker's intention of analysing what the identity of a property consists in, as it does not tell us what distinguishes properties from other basic metaphysical categories. Instead, it tells us something about what constitutes the identity of properties, since it claims that two property instances are instances of *the same property* iff they share all their causal features. But do two-level criteria of identity really entitle us to such metaphysical claims about the entities cited in them?

Admittedly, there are different ways of interpreting these two-level criteria of identity. We could, for instance, interpret them as epistemological principles, which tell us something about our capacity to know about the identity and distinctness of entities of certain types. Understood in this way, the question which CTP is addressing is something like this: when is one presented property x (presented to a subject and thus instantiated at a particular point in space-time) the same property as another presented property y (presented to a subject and thus instantiated at a particular point in space-time but at a different space-time location from x)? It may well be that CTP's criterion, understood as an answer to this question, has metaphysical dividends. For it is reasonable to expect the reliability of a recognitional process to approximate the conditions which are in fact necessary and sufficient for different presentations of the entity in question to be presentations of the same entity.⁴⁰ However, there is no guarantee that our recognitional processes will latch onto what, metaphysically speaking, distinguishes a property from other properties. Therefore, read in this way, the criterion is not a metaphysical principle which discloses what the identity of properties consists in.

Nevertheless, the two most prevalent ways of reading two-level criteria of identity (sometimes referred to as abstraction principles) do result in metaphysical commitments. Moreover, these two readings of two-level criteria offer a way of understanding the dispute between strong and weak CTP theorists. The first is what Wright, Williamson and Dummett call the reductionist view.⁴¹ This identifies the

⁴⁰ This conception of two-level criteria of identity is put forward by Williamson 1990.

⁴¹ See Wright 1983, Williamson 1990 and Dummett 1991.

entities referred to by the functional terms with the equivalence class of entities under the specified relation. So, in the case of properties, the values of the functional term, the properties, are identified with the equivalence class of property instances under the relation, *sameness of causal features*. The result of this being that a property is nothing over and above a set of property instances, all of which have the same causal features.

This interpretation of two-level criteria, properly understood, is endorsed by Dummett. He argues that the reductionist approach doesn't commit us to saying things like 'directions don't exist', since these criteria explain what it means to say things like 'there is a direction orthogonal to those lines'.⁴² However, directions should not be taken to refer to abstract objects in the world, since the sense of the term 'direction' is grasped through an equivalent sentence not containing it. There is thus no need for the reference of the term 'direction' to be thought of as referring to anything over and above that which is referred to by its equivalent sentence. Hence, directions are nothing more than lines under the relation of parallelism.

Wright vigorously opposes this reading of the two-level criteria, opting instead for the non-reductionist alternative. Frege wanted his criterion of identity for directions to be read non-reductively, i.e. as genuinely referring to directions, understood as abstract objects. On this interpretation, a two-level criterion of identity states that the entity picked out by the functional term is distinct from the equivalence class of entities under relation R. However, the identity of that entity referred to by the functional term is nevertheless necessarily correlated with the identity of that class. So, in the case of properties, a property's identity is necessarily correlated with an equivalence class of property instances, under the relation *sameness of causal features*.

It is not difficult to see how this relates to the distinction between strong and weak CTP. We can understand strong CTP as endorsing the reductionist reading of the two-level criterion of identity, and weak CTP as maintaining the non-reductionist interpretation. This coheres well with the contrast drawn earlier between strong and weak CTP. The strong CTP's claim that the causal profile of a property is all there is to a property implies identity conditions for properties. And on this reading it does, because we are committed to the claim that the property of property instance x is the

⁴² See Dummett 1991, p.191.

same property of property instance y iff x and y have the same causal features. However, this does not exhaust the content of strong CTP, because it also commits us to the further ontological claim that properties are nothing over and above equivalence classes of property instances under the relation, *sameness of causal features*. Therefore, this reading offers a way of spelling out strong CTP's interesting but obscure claim that properties are nothing more than their causal features.

In contrast to strong CTP, weak CTP does not just result in a thesis about property individuation, it is a thesis about property individuation. This coheres with the non-reductionist reading, because all this says is that there is a necessary correlation between a property and an equivalence set of property instances under the relation, *sameness of causal features*. So we can account for the identity of properties in terms of their instances. Furthermore, like weak CTP, this non-reductionist reading views properties as something over and above sets of property instances. Interpreting weak and strong CTP as non-reductionist and reductionist readings of the two-level criterion thus captures the essence of both positions. It is difficult to see what is at issue between strong and weak CTP, so one may suspect that they boil down to the same thing. But this reading illustrates that this needn't be the case. Strong CTP can be understood as the claim that properties are nothing over and above sets of property instances, while weak CTP can be thought of as denying this claim.

Which reading, the reductionist or non-reductionist, should be adopted? I don't think that either interpretation is forced upon us. Wright convincingly argues that reductionism about directions cannot be established from the mere fact that every statement explicitly about directions can be systematically replaced by one which refers only explicitly to lines. For if we are prepared to accept that the translated statements (i.e. ones not involving directions) do not refer to directions, on the grounds that their surface grammar doesn't mention them, we have reason to believe that the untranslated statements refer to directions, on the grounds that their surface grammar does mention them. Without further argument, therefore, there is no reason to think that the entities referred to by the translated entities are ontologically prior to those referred to by the untranslated sentences. However, Wright only shows that we cannot just assume a reductionist reading of these two-level criterion, he doesn't establish that such a reading is always inappropriate. Whether or not a reductionist

reading is fitting will, I suspect, depend upon the entities in question. Williamson, for instance, rejects a reductionist reading of the two-level criterion he offers for persons. He argues that the identity conditions for a person should be given in terms of spatiotemporal locations but, he writes, “a person is not a class of spatio-temporal regions” (1990, p.150). So we cannot accept a reductionist reading of this two-level criterion.

The claim that a property just is a class of property instances, however, seems far more plausible. Those who are trope theorists take it upon themselves to argue for the thesis that properties are sets of *sui generis* property instances. Consequently some CTP theorists may be perfectly willing to accept this reductionist reading of the criterion. The same will not be true of those CTP theorists who are firm believers in universals (understood as *sui generis* entities), however. They will want to hold onto the claim that universals or properties are something over and above their instances, and so the non-reductionist reading of the criterion will be far more acceptable to them.

Which of these two positions is ultimately most attractive depends, therefore, upon far-reaching metaphysical issues. Initially at least, the reductionist reading recommends itself on the grounds of ontological economy. Moreover, if properties are just equivalence classes of property instances under the relation, sameness of causal features, this explains why there is a necessary correlation between the two. It may turn out, however, that *sui generis* universals have important theoretical work to do in our metaphysical theories, so a more abundant ontology meets our needs better. As these are enormous issues, which reading of CTP’s criterion is preferable will not be decided upon here. What I do think we can say, however, is that the two-level criterion outlined captures the sorts of claims CTP theorists have made. So in one form or another, it should be considered central to CTP.

2. Generalising Functionalism

In his 1981 paper, Shoemaker suggests that functionalist theories within the philosophy of mind bear an important relation to CTP, because the latter can be thought of as a generalised form of functionalism about mental properties. In other words, rather than just saying that functionalism holds true of *mental* properties, instead CTP can be understood as the thesis that functionalism holds true of *all* (non-abstract) properties. It seems evident that there is a close relationship between functionalism within the philosophy of mind and CTP. Roughly speaking, functionalist theories analyse mental properties via their functional or causal roles, i.e. by what causes them and what they themselves cause. While CTP theorists claim that the causal features of a property exhaust its nature or at least are essential to it. This close connection has lead philosophers to identify CTP as a generalised form of functionalism. Campbell, for instance, in his recent book, classifies Shoemaker's account as "a functionalist theory".¹ In this chapter, I intend to examine this claim that CTP can be understood as a form of functionalism. I think that once we have spelt out the details, matters become more problematic than they initially seem. Indeed, I shall argue that generalising functionalism results in a number of serious problems. Before we can embrace this theory, therefore, some resolution to these problems must first be found.

2.1 Functionalism Outlined

In order to assess the claim that CTP is a form of generalised functionalism, we first need to understand what it is to say that functionalism holds true of all properties. The discussion will thus begin with a more detailed outline of functionalism.

Since Lewis, it has become standard practice for functionalists within the philosophy of mind to employ the Ramsey-Carnap-Lewis procedure to explicate their theory.² Shoemaker (1981) is one such philosopher who utilises this procedure for mental properties. What is innovative about his characterisation, however, is his

¹ See Campbell 2002a, p.236.

² See Lewis (1970, 1972). I shall call this the Ramsey-Carnap-Lewis procedure (RCL for short).

suggestion that we can extend the use of the RCL procedure to cover all properties. He writes,

If we could specify all of the causal features of a property in a set of propositions of finite length, then using that set of propositions as our ‘theory’ we could use the Ramsey-Lewis technique to construct a functional predicate which is true of a thing...just in case it has that property (1981, p.274).

What, then, is this ‘Ramsey-Lewis’ (or RCL) technique? It provides us with a way of formulating a functional definition for any predicate, be that mental, theoretical or just plain old physical. We start by stating the theory for the predicate(s) in question (the T-terms). This theory should state how the referents of the predicates interrelate with the referents of other, already understood, terms (the O-terms). Let’s begin by taking an example from the philosophy of mind. Suppose that the mental predicate ‘is irritable’ is defined by this very simple theory:

(T) If an animal is hungry and tired, then it will be *irritable*.

Or, $\forall x[(Hx \wedge Tx) \rightarrow Ix]$

(where x ranges over animals, H stands for ‘is hungry’, T for ‘is tired’ and I for ‘is irritable’.)

This theory, according to Lewis, provides us with an implicit definition of the predicate ‘is irritable’, but by utilising the RCL, we can turn it into an explicit definition of this term. First, we write the Ramsey sentence of the theory by replacing the T-term with a variable and prefixing an existential quantifier to the theory. (If the theory is being used to define more than one T-term, then a different variable must be used for each T-term.) In our theory, F_3 will replace the T-term ‘is irritable’, so we get:

$\exists F_3 \forall x [(Hx \wedge Tx) \rightarrow F_3x]$

This states that there is some F_3 such that for all animals x, if x is tired and x is hungry then x is F_3 . This does not rule out the possibility of there being more than one state of which everything the theory says is true of it. So if we accept Lewis’s (1972) claim that a term only succeeds in denoting something if it is true of one entity, then the Ramsey sentence should be modified. We can do this by stating that there is only one thing which the theory is true of. Let’s write that as $\exists!F_3 \forall x [(Hx \wedge Tx) \rightarrow F_3x]$. We

can then form the Carnap sentence of the theory by stating the conditional of the Ramsey sentence. So we say, *if* there is an F_3 such that everything the theory says is true of it, *then* that property of being F_3 will be the property of being irritable, i.e. $\exists!F_3 \forall x [(Hx \wedge Tx) \rightarrow F_3x] \rightarrow (\lambda x F_3x = \lambda x Ix)$.

It is not difficult to see why Shoemaker thinks that the RCL technique could be used to provide functional definitions for all properties' predicates. Take any non-mental predicate like 'is one degree celsius'. This can be given a theory T, which states what the causal profile of the referent of this predicate is, in other words, it states what its causes are and how it causally interacts with other entities. Suppose that the predicate 'is one degree celsius' is defined by this very simple theory:

(T) If an object is made out of ice and *is one degree celsius* then it will begin to melt.

Or, $\forall x[(Ix \wedge Tx) \rightarrow Mx]$

(where x ranges over objects, I stands for 'is ice, T for 'is one degree celsius' and M for 'is melting').

Then we can offer this modified Ramsey sentence: $\exists!F_2 \forall x[(Ix \wedge F_2x) \rightarrow Mx]$, which states that there is exactly one F_2 such that for all objects x , if x is made out of ice and x is F_2 then x will melt. Finally, we can state the Carnap sentence for the theory, by saying that if $\exists!F_2 \forall x[(Ix \wedge F_2) \rightarrow Mx]$ then F_2 will be the property of being one degree celsius.

Why have philosophers been so keen to utilise the RCL technique? Its popularity, to a large extent, can be put down to its ability to combat circularity. Behavioural definitions of mental states are notoriously vulnerable to circularities, because in an attempt to define what someone who, for instance, desires a drink will be disposed to do, reference will have to be made to other mental states of theirs, like their beliefs and other desires that they possess.³ Shoemaker's proposal looks vulnerable to a similar criticism. Take, for instance, the definition offered of the predicate 'is one degree celsius'. This involves reference to other predicates, namely, 'is ice' and 'is melting'. Now, according to Shoemaker, all the predicates of properties

³ See, for instance, Chisholm (1957) and Geach (1957).

which have a causal profile should be defined via the RCL technique. The predicates ‘is ice’ and ‘is melting’ that appeared in the definition of ‘is one degree celsius’, therefore, should also be defined in the manner that ‘is one degree celsius’ was. But then our definition of ‘is ice’ will have to make reference to the predicate ‘is one degree celsius’ and ‘is melting’, and our definition of ‘is melting’ will have to make reference to the predicate ‘is ice’ and ‘is one degree celsius’. The definitions thus appear circular, because although the predicate being defined will not appear in its definition, the definitions of the predicates used to define the predicate in question will make reference to the predicate they supposedly define.

By utilising the RCL technique, however, Shoemaker can bypass this objection, in the same way that functionalists such as Lewis and Shoemaker himself avoid the circularity complaint. It is true that a property’s predicates will be interdefinable if CTP theorists decide to utilise the RCL technique to define all the predicates which pick out properties. Nevertheless, this does not show that the definitions on offer are circular, because it makes sense to suppose that a theory can simultaneously define several items. Take, for instance, our theory for ‘is one degree celsius’ and suppose that this also serves as the theory for ‘is ice’ and ‘is melting’. In my initial definition of ‘is one degree celsius’, the other two predicates were seen as O-terms and so they were not replaced by variables. However, if we treat them all as T-terms, then we can rewrite the Ramsey sentence as follows: $\exists!F_1 \exists!F_2 \exists!F_3 \forall x[(F_1x \wedge F_2x) \rightarrow F_3x]$. Then the predicate ‘is one degree celsius’ can be defined as the property P such that $\exists!F_1 \exists!F_2 \exists!F_3 \forall x[(F_1x \wedge F_2x) \rightarrow F_3x]$ and $\lambda x F_2x = \lambda x Px$, similarly with the predicates ‘is ice’ and ‘is melting’. The idea is that because all of these predicates get their designations concurrently, all the definitions succeed in picking out their references. So the RCL technique provides us with definitions of which it makes sense to suppose that several terms get defined simultaneously. If we formulate a mammoth theory in which all the predicates of properties appear and RCL that theory, therefore, it looks like the generalised form of functionalism which Shoemaker proposes can avoid the circularity objection, like its sister account in the philosophy of mind.

The RCL technique is, then, an important tool for functionalists of all varieties. But what do the resulting RCL definitions of predicates tell us about the nature of their

referents? Answers to this question vary, depending upon how the RCL definitions are interpreted. In the philosophy of mind, however, two different interpretations of the RCL definitions stand out, as they have led to a fundamental division between functionalist accounts.⁴ One way of interpreting the definition accords with functional realiser theory.⁵ This takes the property, say of being one degree celsius, to be *the property* which, when instantiated by an object alongside the property of being ice, will result in the object's melting. This way of interpreting the RCL was suggested by my exposition of it. Take the Ramsey sentence for the theory $\exists!F_1 \exists!F_2 \exists!F_3 \forall x[(F_1x \wedge F_2x) \rightarrow F_3x]$. The predicate 'is one degree celsius' was defined as the property P such that $\exists!F_1 \exists!F_2 \exists!F_3 \forall x[(F_1x \wedge F_2x) \rightarrow F_3x]$ and $\lambda x Px = \lambda x F_2x$. Therefore, the entity which the predicate picks out was identified with the state which occupies or realises F_2 .

The second way of interpreting the RCL definition conforms with functional role theory.⁶ This doesn't claim that the referent of the predicate 'is one degree celsius' is the occupier of F_2 . Instead it argues that it is a higher-order property - it is the property of possessing a property which, if instantiated by an object alongside the property of being ice, will result in the object's melting. Or, in other words, the predicate 'is one degree celsius' refers to the property of having a property which realises or occupies the functional role of F_2 . The functional role theorist's interpretation of the RCL, therefore, states that the property of being one degree celsius is instantiated by an object if it does or can display all of the relations specified by the RCL definition for its predicate. Functional realisers, on the other hand, claim that the property of being one degree celsius is picked out by the definition which describes its interrelations with other entities. However, the property itself is identified

⁴ Within the philosophy of mind, there are other versions of functionalism than those I outline above. Philosophers, for instance, are divided upon the issue of what sort of theories should introduce mental predicates. Should it be a scientific theory (psychofunctionalism) or our folk theory about the mind (analytic functionalism)? For our present purposes, however, the two different kinds of functionalism I shall outline are the most relevant to the discussion.

⁵ This is a popular position in the philosophy of mind, endorsed by Lewis (1966, 1972, 1995) Armstrong (1970) and Kim (1972, 1998). (The view is also sometimes called 'functional specification theory' see Block 1980.) The label 'functional realiser theory' here, however, shall be used to refer to this way of interpreting the RCL technique, rather than to any specific theory of mind. The same goes too for the soon to be mentioned functional role theory (also called 'functional state identity theory', see Block 1980).

⁶ See, for example, Putnam (1967), Block (1980) and Shoemaker (1981).

with the placeholder of F_2 , not with the abstract causal role specified by that definition. Both of these interpretations of the RCL incorporate various metaphysical assumptions into the procedure. Before examining what these definitions tell us about the nature of properties, therefore, I shall first spell out, in a little more detail, what metaphysical claims are being fed into these interpretations.

One thing immediately evident about these two interpretations of the RCL definition is that they both make the transition from talk of words (predicates) to talk of entities in the world (properties). This aspect can be seen as stemming from the technique itself, because when ramsifying the definition we are presuming that there are entities which ‘satisfy’ or ‘realise’ the formula. However, the conception of realisation differs between the two sorts of functionalism. Functional role theorists talk about first-order physical properties realising second-order mental properties. It is thus viewed as a relation between properties which, due to their distinctness thesis, is weaker than identity. The functional realiser’s notion of realisation differs from this, because instead of talking about one sort of property realising another sort, they speak of n-tuples of entities realising *theories*. The thought is that if a theory is realised by a certain set of entities, then it is true of or applicable to those entities.

The functional realiser’s notion of realisation, then, is just the relatively well understood relation of identity. To say that the predicate ‘is in pain’ is realised by C-fibres firing, for instance, is just to assert that there exists a relation of identity between the entity named by the predicate ‘is in pain’ and the entity named by the predicate ‘C-fibres firing’.⁷ So if a theory is true of a certain (unique) set of entities, then the entity which fills the role marked by the predicate or T-term in question will be the metaphysical correlate of that predicate. What is noticeable about this way of proceeding is that it avoids having to assume that the predicate will latch onto any property. The predicate may pick out a very complex state of affairs, composed of entities like compound or negative states of affairs which, some have argued, cannot be properties.⁸ Or, alternatively, there might not be a unique property which occupies this causal role, and so we will not be able to identify any single property with this

⁷ Although like ‘the Prime Minister of Britain’, the predicate ‘is in pain’, on this view, is a (disguised) definite description, so it may name other entities in different possible worlds. If we go along with Kripke, therefore, we are committed to the claim that the predicate ‘is in pain’ is not a rigid designator.

⁸ See, for example, Mellor 1991.

predicate. Whatever the case may be, because functional realisers don't claim that the predicate has to pick out a property, this interpretation of the RCL definition accords with the realist thesis that not all predicates need name properties.

This thesis, however, appears threatened by the functional role theorist's interpretation. They reject the functional realiser's claim that the predicate (if it names any property) picks out the property which occupies the role specified by the definition. Instead, they argue that the predicate refers to a new, second-order property, not named in the domain quantified over. This renders their interpretation of the technique ontologically creative. Why? The quantifiers in the RCL definitions are said to range over first-order properties, or unproblematic properties which are agreed to exist. So in the philosophy of mind, the quantifiers are said to range over respectable physical properties. However, the property which the predicate names, on this view, isn't one of the properties quantified over. Instead, it is a new, higher-order property, which is instantiated just in case some first-order property satisfies the functional role definitive of the predicate in question. So this interpretation of the RCL definitions is metaphysically inflationary, because by quantifying over first-order properties, new, second-order properties are created, which cannot be identified with those properties originally quantified over.

This interpretation of the RCL definitions raises a substantial issue: why say that the predicate specifies a new, second-order property rather than either one of the pre-existing properties named in the domain, or no property at all? It may be thought that the functional role theorist's use of the term 'second-order property' points to an answer to these questions. Originally, talk of different orders of properties belonged to Russell's ramified theory of types. So it will help us to understand the present proposal if we digress, for a moment, on Russell's theory. There, every *type* of entity was further split into orders of entity. Take, for instance, any type 1 property (i.e. those that apply significantly to individuals).⁹ We can split this type further into orders. The first-order of properties will be those properties of individuals expressed by predicates which either contain no quantifiers, or only quantifiers whose domain is individuals.

⁹ This draws on Russell's 'simple theory of types'. On this view, individuals belong to the lowest type in the hierarchy – type 0. Type 1 consists of those properties which can only be significantly applied to individuals. Type 2 consists of those properties which can only be significantly applied to type 1 properties, etc.

The second-order will be all those properties expressed by predicates whose definitions require quantification over first-order properties. The third-order will be all those properties expressed by predicates whose definitions require quantification over second-order properties, and so on.

The motivation for Russell's theory of types and orders sprang from his desire to avoid logical and semantic paradoxes. Russell argued that in order to sidestep these paradoxes, we need to impose restrictions on our language. For instance, with second-order predicates of type 1, the quantifiers must range over entities which are distinct from those referred to by second-order predicates. If they do not, the vicious-circle principle will be violated, as the entity defined will be among those quantified over.¹⁰ But why should we accept that the entities specified by these definitions couldn't name one of the entities quantified over? After all, as Quine, Ramsey and many others have pointed out,¹¹ the definite description 'the tallest man in the group', identifies a person via the group of people to which he belongs and there seems to be no vicious circularity here.

The answer to this lies in Russell's understanding of the properties specified. Sainsbury, for instance, claims that the vicious-circle principle holds because Russell accepted

an anti-realist theory of properties...a property is properly specified only by a predicate with a certain structure, we can regard this structure as woven into the nature of the property. By contrast with objects which can in principle be named, and thus specified by a structureless expression, in the case of properties we lack a full distinction between our mode of specifying them and their nature (1979, p.285).

So the idea is that because the definitions are seen as introducing new properties, whose very essence is given "by our mode of specifying them", they cannot be among the entities quantified over as these properties have different structures.¹²

¹⁰ There are many different formulations of Russell's vicious-circle principle, but here is one Russell employs, "whatever involves all of a collection must not be among one of the collection" (1925, p.37).

¹¹ See Ramsey (1931) and Quine (1963). The example is taken from Ramsey.

¹² As Russell also famously held a Platonist theory of universals, the claim that he was an anti-realist about properties may sound odd. However, this can be explained by the fact that we should probably distinguish between Russell's theory of universals and his theory of propositional functions (see Linsky 1999 for a convincing defence of this interpretation). Although Russell was a realist about universals, he takes a constructivist approach to propositional functions (see above for more on this. By 'propositional functions' I just mean to refer to those propositions where one or more of the arguments

But while this view of properties certainly suffices to explain why the definitions are conceived of as picking out different entities from the ones quantified over, such strong anti-realism is not necessary. We can think of these definitions as specifying entities not quantified over in the definitions, without thereby viewing them as somehow creations of our human language. Linsky, for instance, claims that Russell thought of these properties as “constructions out of constructions. They are constructions of, and thus dependent upon propositions, which are in turn constructed from particulars and universals” (1999, p.28). The definitions of higher-order predicates indicate how the properties specified are constructed out of those named by lower-order predicates. But this doesn’t mean that these properties are our creations. While they are not ultimate constituents of reality, if we conceive of a construction as a metaphysical relation, independent of us, then the entities thus constructed will be dependent upon the particulars and universals they are constructed from, but not upon us.

When we turn to the philosophy of mind, there are noticeable similarities and dissimilarities with Russell’s project. Their theory is similar in that functional role theorists wish to postulate hierarchies of properties, which are ontologically creative in the same way that Russell’s are. But, clearly, the motivation behind functional role theory is very different from Russell’s, as they are not interested in trying to solve logical or semantic paradoxes. Furthermore, functionalists within the philosophy of mind do not display any obvious adherence to the rules imposed by the ramified theory of types. So what is it that motivates functional role theorists to postulate these hierarchies of properties? And why do they think that these newly defined predicates serve to pick out properties not quantified over in their definitions?

It is the multi-realizability argument that is supposed to show that mental properties cannot be straightforwardly identified with physical properties. This argument states that the property of being in pain, for instance, cannot be identified with any one physical property which realises the property of being in pain within a system, since many different properties could (and perhaps do) realise that property. Consequently, it is unacceptably chauvinistic to assert the identity of mental properties

are removed, like ‘x is a man’). This distinction is loosely reflected in Russell’s preferred terminology, as he tends to refer to universals as ‘qualities’ and propositional functions as ‘properties’.

with physical properties, as functional realisers do. For this means denying that creatures with different physical make-ups can instantiate the very same mental properties.

The claimed multi-realizability of these RCL definitions, however, poses a problem for the proposed RCL technique. For if different physical properties can realise this functional role, the uniqueness of the RCL definition is threatened, as no one entity is specified by it. This means that functional role theorists have to reinterpret the RCL definitions in light of multi-realisation. How? Consider a certain property's mental theory, which specifies how that property causally interrelates with other entities. This theory (be it scientific or folk) needs to be localised to all those systems where the same properties realise the functional role stated in it. Take, for instance, the aforementioned property of being irritable. If the F_3 role in $\exists F_3 \forall x [(Hx \wedge Tx) \rightarrow F_3x]$ is always realised by the property of I-fibres firing in system S, then that counts as one concrete realisation of the property's RCL theory. For restricted to just this system, there is a unique physical property vouchsafing every predicate variable.

This concrete realisation of the mental theory, however, cannot be identified with the mental property of being irritable, because different properties can realise this theory. In dogs, for instance, K-fibres firing might realise this theory, so no single physical property can be identified with that property. In order to get round this, we need to take all the concrete realisations of the mental theory. Then we can abstract (in the Fregean sense¹³) to the functional role in common across the realisations, by considering the equivalence class of concrete systems that exemplify the form of the mental theory. This abstracted functional role property, which is common to all the various concrete realisations, is then identified with the mental property. By starting with uniqueness relative to a given realisation, therefore, uniqueness can be secured again by taking the set of systems which are functionally similar.

But why should we take multi-realizability to show that the predicate in question denotes a functional property, rather than taking it to be a trait common to all those predicates which don't specify any property? The answer to this stems from the functional role theorist's commitment to causally efficacious mental properties. They

¹³ See §1.5.

argue that multi-realizability shows that no physical property has a causal profile which is suitably wide enough to occupy the RCL theory. So, because of the numerous and significant causal generalisations mental predicates such as 'being in pain' are part of, these functional mental properties need to be recognised as part of our ontology. The functional role approach, therefore, does not just give us a complex way of talking about disjunctions of physical properties. It isolates an independent, functional property whose essence is captured by the functional role expressed in its RCL definition.

Despite this proposed ontological addition, however, functional role theorists are still keen to preserve the claim that all entities are, in some sense, physical. This is where the aforementioned notions of hierarchies and constructions come in. Functional role theorists utilise Russell's notion of a hierarchy in order to explain why their postulation of mental properties distinct from physical properties doesn't result in an ontology like the substance dualist. For while functional role theorists are willing to relax the physicalist's ontology a little and say that there are mental properties as well as physical properties, they do not want to claim that mental properties are of a radically different kind to physical properties. So the logical apparatus employed by functional role theorists is there to show how mental properties are constructed out of physical properties. Once we've understood how mental properties are derived from their physical realisers, we'll see that mental properties do not require new substances or ontologically independent laws.

I hope that this brief foray into functionalism has illustrated that the RCL technique, and the resulting definitions, have no definitive readings. Although the RCL is useful in providing a focal point with which to understand and interpret functionalist theories, as various different metaphysical assumptions can be incorporated into this procedure, different versions of functionalism can be extracted from it. The lesson to take away from this is that logical techniques such as the RCL cannot do the meaty metaphysical work for us. In order to get substantial conclusions about the nature of properties (or a certain set of properties) from the RCL, we need to put substantial claims about the nature of properties into the RCL. This, we discovered, is precisely what functional realiser theory and functional role theory do. Functional realiser theorists incorporate the metaphysical thesis that the referent of the

predicate is identical to the occupier of the role specified in their interpretation of the RCL definition. Whereas functional role theorists incorporate the contrary thesis into their RCL definition, since they claim that the occupier of the role specified is not the referent of the predicate. Of the two, the first metaphysical thesis is probably the least problematic, but both should be recognised for what they are – theses which *supplement* the RCL technique.

This examination of functionalist theories within the philosophy of mind has also illustrated that it is not enough to say that CTP is a generalised form of functionalism. For there are different kinds of functionalism, therefore, we need to know which one is supposed to serve as a model for CTP. Unfortunately, this issue is not addressed by Shoemaker, Campbell or anybody else, so the content of CTP is left undetermined by these expositions. In what follows, I hope to remedy this, by looking to see whether some form of functionalism can capture a plausible-sounding CTP. I shall begin by considering whether CTP can be understood as a generalised form of functional realiser theory.

2.2 CTP as Functional Realiser Theory

If we generalise functional realiser theory to all properties with concrete instances, then all the predicates that pick out causally efficacious properties will be defined via the functional realiser interpretation of the RCL. This commits us to the claim that all the predicates of concrete properties have functional definitions, but what does it tell us about the nature of the properties thus defined? Take, for instance, the predicate ‘is red’. What ontological analysis of this property and its instances does the resulting definition commit us to? The answer is, none at all. We cannot assume that the predicate picks out a property, since it may be that no one entity satisfies the predicate. But even if one does, the definition does not commit us to any analysis of what these properties are. They could be universals, sets of tropes, sets of possible particulars, etc. This interpretation of the RCL leaves the whole range of theories wide open, thus, no ontological analysis is given of properties.

Does the definition tell us something about what kind of property redness is? Again, I think the answer is no. The definition does not give any indication of what kind of property the T-term ‘is red’ refers to, since the domain over which the

existential quantifier ranges in the RCL definition is not specified. It may range over scientifically respectable properties, irreducible colour properties, mental images in the mind, and so on. A functional realiser can decide to identify the domain, they may, for instance, decide to restrict it to scientifically respectable properties. Then, granted of course that they're right, we can say that the predicate 'is red' refers to one of these scientific properties. But the RCL definition does not provide us with this information – we only get this metaphysical claim out of the definition by putting it in.

Despite these limitations, however, it may be thought that these RCL definitions do at least tell us something about the nature of redness, as they result in an account of property individuation. It looks as if the property which the RCL definition picks out (granted it succeeds in picking out a property) must have the causal relations specified by the definition. For, if it didn't, the predicate wouldn't be applicable to it, so it wouldn't be that property. In order for an object to instantiate the property of redness, therefore, it must instantiate a property which occupies the causal role specified by the functional definition for redness. So functional realiser theory seems to commit us to the following two-level criterion of identity for properties: two instantiations or property instances of red are instantiations of the same property, namely redness, iff they both realise the causal role specified by the functional definition for redness.

The generalised form of functional realiser theory is thus beginning to look like a CTP. However, the claim that functional realiser theory results in a two-level criterion for properties can be contested, by questioning whether the property which the definition picks out *must* have the causal relations specified by that property's functional definition. We can argue that it does not follow from the fact that predicates are defined by the causal role specified in their RCL definitions, that the properties picked out by these definitions must always realise these causal roles. For we can utilise the RCL procedure in order to identify the property in question, after which, we can treat the property as an entity independent from the causal role that picked it out. Lewis (1980), for instance, argues that, in special cases, a mental property need not occupy the role specified by its RCL definition. He asks us to suppose that nearly all tests show that C-fibres firing satisfies the causal role specified by the predicate 'is in pain' in humans. However, there exists a mad human who displays none of the normal

behaviour we associate with pain, but whom instantiates the property of C-fibres firing when inflicted with injuries which would cause pain in the rest of us. In this case, Lewis argues, we would still have grounds for ascribing the property of being in pain to this person, as he instantiates the property which *typically* occupies the causal role of pain. Lewis's response to this situation shows that a criterion of identity for properties does not follow from the functional realiser's reading of the RCL definitions. While we discover what pain is via the causal relations described in its RCL definition, once we have picked out that property, it stands as an independent entity, detachable from the causal relations that identified it.

Do matters change once we have generalised the account to cover all concrete properties? It may initially appear so. Lewis's example seemed plausible because while pain was defined functionally, the same was not supposed of C-fibres firing. But now we are assuming that the realiser C-fibres firing is also defined via a RCL functional definition. Presumably, C-fibres firing's functional definition must at least overlap with pain's, if the two can be considered identical. So we may reason that if C-fibres firing didn't bring about the behaviour definitive of pain, in accordance with its functional definition, then it wouldn't be C-fibres firing. This response, however, misses the point. If we accept this reading of the RCL technique under consideration, then while C-fibres firing will be identified as the occupant of a certain causal role, once this identification has taken place, the property will be able to stand independently of the causal role that identified it. This would then allow pain (which is identical to C-fibres firing) to sometimes deviate from its typical causal role in the way suggested by Lewis's example. To say that C-fibres firing always has to produce the behaviour definitive of pain, therefore, ignores this way of interpreting the RCL.

How does this reading of the RCL definition allow for such circumstances to arise? In Lewis's hands, the RCL definitions serve an epistemological, not metaphysical, purpose. For he is trying to show that we can identify a certain, problematic subset of properties, with a domain of, what are taken to be, acceptable properties. The initial causal characterisation given by the RCL definitions, then, isn't intended to characterise the nature of the property. Rather its purpose is to identify the properties in question. This allows properties to come apart from their causal profiles. For once a property has been picked out, there is no reason to think we must rely on its

causal characterisation. There may be some independent, non-causal description of that property, which could then be used to identify that property. If we adopt this reading, therefore, the RCL functional definitions will not result in two-level criteria of identity.

There are other functional realiser readings of the RCL technique, however. If we adopt Kim's (1998) functional realiser account, for instance, there will be a tighter link between a property and its causal role. Kim argues that in order for a creature to instantiate a mental property, such as pain, that animal must display the causal relations specified in its functional definition *and* instantiate the typical realiser of pain for this sort of system with these laws of nature. On this interpretation of the RCL definition then, the property cannot be separated from its causal characterisation. Even if someone does instantiate the typical realiser of pain, the property cannot be attributed to them if it does not give rise to the behaviour definitive of pain. This reading does seem to commit us to the claim that two instantiations are instantiations of the same property, namely pain, if they both realise the causal role specified by the functional definition for pain. We are not yet at a two-level criterion of identity for properties, however, as this only states a necessary condition for property identity. It is not sufficient because the creature must also instantiate the property which typically realises that role for its sort of system.

If we functionalise all the way down, however, and say that every property has to realise the causal profile specified by its functional definition, then it looks like the two-level criterion might be forthcoming after all. For C-fibres firing, which we can suppose occupies the pain role in humans, will also be defined functionally. In order to instantiate the property of C-fibres firing, a creature will have to instantiate the property which realises the causal role specified by its functional definition. That will mean that the creature will display the behaviour definitive of pain, as this is at least part of C-fibres firing's causal role, and so the property of pain will not be able to come apart from its functional role. This functional realiser reading of the RCL definitions, therefore, does seem to result in the aforementioned two-level criterion of identity.

Closer inspection of this theory, however, shows that this conclusion can be challenged. Kim thinks that we can ascribe the predicate 'is in pain' to an animal, if it

displays the causal role definitive of pain. In different systems, however, the property which occupies this role can vary. Imagine, then, that C-fibres firing is the occupier of the pain role in systems like humans, and D-fibres firing the occupant of the role in systems like dogs. It looks like a realist about properties will want to say that C-fibres firing is a different property from D-fibres firing. But if this is the case, we cannot say that two instantiations or property instances of pain are instantiations of the same property iff they both realise the causal role specified by the functional definition for pain. These will not be two instances of the same property at all – one will be an instance of the property of C-fibres firing and the other an instance of the property of D-fibres firing.

It may be thought that this objection doesn't apply once all concrete properties are defined in Kim's way. But this isn't so. If the RCL definition for C-fibres firing is interpreted in Kim's way, the predicate will only be true of properties which can occupy this causal role. But this doesn't mean that the predicate won't apply to different properties in different systems. Consequently, even once we've generalised the account, there will still be no guarantee of a property which the predicate's causal characterisation is true of. If a functional realiser did want their account to result in the two-level criterion of identity outlined above, they would need to say that an RCL definition serves to pick out a new property. But this path leads to problems. The functional realiser could treat the property which the predicate picks out as a new, higher-order property. But then our functional realiser goes the way of the functional role theorist. Alternatively, they could say that the property picked out by the predicate 'is in pain' is a disjunctive, gerrymandered affair which is composed of all the different occupiers of this pain role. But now the account we are advancing will raise the hackles of the realists. Why should these gerrymandered and disjunctive affairs count as properties?¹⁴

Functional realisers' interpretations of the RCL definitions, therefore, do not produce commitments that look particularly CTP-like. Even when we have generalised

¹⁴ If we endorse Lewis's account of properties (1983b), then we need two different types of entities to satisfy the roles associated with our notion of a property. The first of these does count distinctly disjunctive and gerrymandered sets of possibilities as properties. But these sorts of properties are assigned a semantic role, and it is clear that the notion of a property CTP theorists are intending to analyse is the second that Lewis outlines, namely, the entities which ground genuine resemblances between objects and can thus be used in accounts of causation and laws.

the account to include all concrete properties, metaphysical claims about the nature of properties are still not forthcoming. This doesn't mean that CTP theorists shouldn't utilise the RCL procedure. After all, it is very useful for avoiding the circularity objection. In order to make the account recognisably a causal theory of properties, however, a CTP theorist would need to incorporate different metaphysical assumptions into the RCL procedure. In particular, a CTP theorist would need to augment it with the claim that a property's causal role is essential to it and that no other property could realise the same causal role. This would rule out the possibility of there being different occupiers of the causal role defined, thus making it possible for realists about properties to adopt the two-level criterion of identity for properties.

I think we must conclude, however, that something which can be thought of as a CTP doesn't just fall out of the functional realiser's interpretation of the RCL procedure (whether that be Lewis's or Kim's). In view of this, it is questionable how useful a reading this is, since it does not yield any of the metaphysical claims definitive of CTP. What went wrong? Lewis's and Kim's interpretations of the RCL definitions make good sense within the philosophy of mind. For, there, the RCL technique is employed to show how mental properties can be identified with physical properties. But the generalisation of this strategy to *all* properties doesn't seem to make sense. For there needs to be an uncontroversial domain of properties, with which the RCL-defined properties can be identified. So a generalised form of functional realiser theory in the philosophy of mind cannot serve as a model for CTP.

2.3 CTP as Functional Role Theory

The inappropriateness of interpreting CTP as a form of functional realiser theory was, perhaps, to be expected. For in the philosophy of mind, Shoemaker (1981) adopts the functional role interpretation of the RCL. So when he talks about CTP being a generalised form of functionalism, it is natural to think that he has his own version of functionalism in mind. In this section, therefore, I shall consider what metaphysical claims the functional role reading commits us to.

According to functional role theorists, the referents of RCL definitions are higher-order properties. These higher-order properties display all of the causal features specified by their predicate's definition. Take, for instance, the property of being cold.

If we generalise the functional role theory, and so regard the property of being cold as a functional role property, then it is exhaustively characterised by the sorts of circumstances that cause it, and what an entity is disposed to do in virtue of instantiating that property. What does this tell us about what properties, ontologically speaking, are? Again, it seems nothing. To say that a property is exhaustively characterised by its causal role, doesn't seem to tell us anything about the ontology of properties - whether, for instance, a property is a universal, a set of tropes, a set of possible particulars, etc. Does it tell us something about the sort of property picked out by the RCL definition, i.e. whether the property is physical, irreducibly mental, etc? If we generalise the account, then all properties are said to be functional. These are properties which, in Block's words, "consist in the having of some properties or other...that have certain causal relations to one another" (1980, p.155). But this doesn't tell us what sort of property realises the causal role of the functional property, unless we specify what the quantifiers can range over. So the realiser of the functional property pain, for instance, may be a physical property or mental property. No insight into these sorts of issues is provided.

Nevertheless, it does look as if generalising functional role theory offers an informative account of the metaphysical nature of properties. Since it is saying, in effect, that properties are what they do. This results in metaphysical commitments which are in line with what CTP theorists have wanted to say. First, functional role properties have their causal features essentially. On their account, the RCL definition refers, if it refers at all, to a new, higher-order property which is exhaustively characterised by the causal role described. This higher-order property has to have all the causal features stated in the RCL definition, in every possible world where it exists, since that property can only be instantiated in a possible world, if something realises the causal features definitive of it. Why does this commitment arise from the functional role, but not the functional realiser, interpretation of the RCL? Because functional role theorists claim that the predicate introduces a new, higher-order property, whose sole characterisation is given by the RCL definition. These new higher-order properties have their causal features essentially; they could not exist without them. On the functional realiser theory, on the other hand, the RCL definition picks out a preexisting property, which can vary from system to system. There is thus

no guarantee that any of these preexisting properties which are picked out by the definition have to instantiate the causal features they do in certain systems. For the property which occupies the causal role in the RCL definition can vary from system to system and world to world.

As well as resulting in the thesis that properties have their causal features essentially, the functional role theorist's interpretation of the RCL technique also commits us to the aforementioned two-level criterion of identity for properties, which states that property instance *x* (at a particular time/place/possible world) instantiates the same property as property instance *y* (at a different time/place/possible world) iff they realise the same functional role.¹⁵ This follows from the account because the nature of a functional property is exhaustively characterised by its functional definition. So an object instantiates such a property just in case it has some entity capable of realising its functional role.

The functional role reading of the RCL technique, therefore, does result in claims which render it recognisably a causal theory of properties. For it commits us to both a causal criterion of identity for properties and the claim that the causal features of properties are essential to them. Furthermore, as functional role properties are wholly characterised by the causal relations they can enter into, we have a reading of CTP which spells out Shoemaker's intriguing claim that "properties are causal powers" (1980a, p.210). Unfortunately, however, there are large problems on the horizon.

¹⁵ This criterion is slightly different from the one outlined in §1.5. There, the following CTP criterion was put forward: one property instance is an instance of the same property as another property instance iff they have all the same causal features. This may be a more illuminating statement of the criterion that results from generalising functional role theory, but this depends upon whether we accept the claim that property instances have to realise all *and only* the causal features stated in the RCL definition. If we do, then adopting the functional role reading of CTP does result in the criterion outlined in §1.5. But some may argue that, due to considerations involving determinates and determinables, we shouldn't accept this restriction, because a property instance could realise the causal features stated in the RCL definition of a property, and then some. We may want to say, for example, that property instance *x* realises the causal features definitive of redness and then also those causal features definitive of scarlet. Accepting this commits us to the disputed metaphysical thesis that two properties can be instantiated in a single instance. I shall return to the issues surrounding this thesis in §6.3. For now, however, I shall just utilise the criterion which states that two property instances are instances of the same property iff they realise the same functional role. While more needs to be said about what exactly this commits us to, this criterion is definitely a commitment of the functional role reading of CTP, and it is neutral between the interpretations outlined here.

2.4. A Problem

We've seen that the functional realiser interpretation of these RCL definitions fails to provide an illuminating account of the nature of concrete properties. At best, it asserts additional constraints on properties, but without clear justification. The functional role interpretation, by contrast, looks far more promising, as it seems to offer an informative and principled analysis of properties. So far, however, I have just assumed that we can generalise the functional role theorist's account to all properties. It was clear that all (non-abstract) properties could be given a functional realiser's analysis, as these properties contribute to the causal powers of particulars. But it is not at all obvious that the same can be said of the functional role interpretation. In order to see why, it will help to reiterate the functional role theorist's strategy in the philosophy of mind.

We have seen that functional role theorists believe that physical properties act on behalf of the mental properties they exemplify. These mental properties are only made manifest by attending to the shared functional roles of these physical properties and perhaps other properties besides. For mental properties are characterised by equivalence classes of concrete systems that exemplify their RCL theories. The result is an informative account of the nature of mental properties. Not only does the functional role interpretation of the RCL definitions render theses about the identity and essential nature of these mental properties, the properties analysed are also shown to be logical constructions out of their realisers.

When we turn to a generalised form of this theory, however, it is not clear that the functional role theorist's method can be implemented. For in the philosophy of mind, it is assumed that there is a distinct domain of physical properties which realise the mental properties (i.e. which implement the functional roles stated in their RCL theories), in particular concrete realisations of these theories. But once we've generalised the account to all properties, then we are explicitly denying that there is this separate domain of realiser properties. Without these, it is difficult to see how the account could work. For, if there are no properties which can realise the functional roles stated in these RCL theories, no functional properties can be abstracted (in the Fregean sense) from their concrete realisations.

To this, the following response might be offered: although, if we generalise the account, non-functional properties cannot realise the roles of functional properties, we can suppose that the RCL theory for a functional role property is realised by a distinct functional role property. So we get an infinite series of functional properties, each realising the order of property higher than itself. I don't think that this removes the problem, however.¹⁶ Why not? What is distinctive about functional role theory in the philosophy of mind is its constructive character. Functional role theorists deny the functional realiser's claim that mental properties can be identified with their realisers, arguing instead that mental properties are new, higher-order properties which are constructed out of their realisers. It is important that this aspect of the account is incorporated into CTP. For we've seen that if we opt for the functional realiser interpretation, these RCL definitions just identify certain properties with such-and-such functional roles, and nothing exciting follows from this about the nature of properties. It is only when properties are thought of as being constructed out of the realisers of these RCL roles, that these RCL definitions can be thought of as specifying the very essence of properties. If CTP's metaphysical ambitions are going to be fulfilled, therefore, a generalised form of functional role theory must parallel the constructive character of functional role accounts in the philosophy of mind.

But why does this exclude an analysis where the realisers are themselves functional role properties? In order for CTP to have the constructive character of functional role theory, there must be a number of properties which are not analysed via this theory. For functional role analyses work by showing how functional properties can be viewed as logical constructions out of their realisers. So if the realisers are of the same kind as those entities realised, the account will not tell us anything about the kind of entity being analysed. To illustrate, recall Russell's hierarchy of types and orders. If the properties specified by higher-order predicates are regarded as logical constructions out of those specified by lower-order predicates, in the way suggested by Linsky's interpretation, then the hierarchy must accord with the vicious-circle principle. For the kind of entity in question will not be illuminated if it is among the collection we quantify over in its predicate's definition.

¹⁶ Although I think that it is effective against a slightly different objection (see §2.5).

Now although functional role theorists in the philosophy of mind do not specifically appeal to Russell's hierarchy, something very like his vicious-circle principle must hold of their hierarchy. Why? If we say that the entity specified by a mental predicate's RCL definition names one of the entities quantified over, then mental properties have to be identified with the realisers of the RCL theories. So (granted the domain is taken to be physical properties) we end up with functional realiser theory, as mental properties are not distinct from physical properties. In order to avoid just getting generalised functional realiser theory, therefore, we must say that the realisers quantified over in functional role predicate's definitions are distinct from the functional role properties thus specified. For only then will we get the CTP commitments outlined in §2.3.

Unfortunately, however, this constraint cannot be met if we generalise the account to all properties. Even if we say that the realiser of a functional role property is another functional role property of a lower-order, and so on to infinity, this will still not help. For properties will be constructed out of a totality of entities which includes those entities whose very nature we are intending to analyse. So generalising functional role theory to all properties isn't workable, because we cannot preserve what is distinctive about this kind of analysis given unrestricted usage.¹⁷

Therefore, despite the initial advantages of this functional role reading of CTP, as it stands it is simply not sustainable. Without a range of properties which are distinct from functional role properties, no properties can be analysed via functional role theory. This makes the possibility of developing a coherent form of generalised functionalism seem remote. Shoemaker and Campbell fail to notice this, because they

¹⁷ The account also conflicts with Russell's ramified theory of types, since if all the predicates that specify properties are given RCL definitions, all the predicates become of an order second or above, as they are all defined by the fact that some other property occupies such-and-such a functional role. Consequently, the predicate quantifiers in the RCL definitions of the second-order predicates have nothing to range over. For according to the restrictions laid down by the ramified theory of types, the definition of a second-order predicate cannot quantify over entities picked out by second-order predicates. The account's failure to meet the constraints imposed by Russell's theory, raises wider issues regarding the legitimacy of the use of this notion of a hierarchy. For it is not clear how the assumption that there is any such ordering of properties is justified, once we have postulated an infinite hierarchy of properties. Why? An infinite hierarchy of properties will not be well-founded, as there won't be a distinct base of fundamental properties, which higher-orders of properties in the hierarchy can then be constructed from. Without this, it is difficult to see how we can defend the claim that properties form a hierarchy, rather than just a collection of properties, all of which are analysed in terms of one another.

do not spell out what consequences follow from generalising functionalism. But once this is done, it seems clear that the strategies employed by functional realiser and functional role theorists in the philosophy mind cannot be extended to all (concrete) properties. What's next? Soon, I shall attempt to challenge the conclusion reached here. But before I do this, I first want to mention another objection to functional role theory. This objection is clearly related to the one here, but rather than pointing out a problem in trying to generalise the functional role theorist's method, instead it argues that extending this theory to all properties commits us to an unsatisfactory metaphysical picture.

2.5 A Metaphysical Worry

Although the strategic difficulty engendered by unrestricted use of functional role RCL definitions has gone unnoticed, the strangeness of the resulting metaphysical picture has not. Not surprisingly, a number of philosophers have expressed some concern over this idea that properties could be "functional all the way down".¹⁸ For if we generalise functional role theory, and so say that all properties are functional role properties, then we're claiming that all properties consist in the having of *some other property* which occupies such-and-such a functional role. But of course if *all* properties consist in the having of some other properties which occupy such-and-such functional roles, then there won't be any properties left to stand in such-and-such functional roles.

This problem is closely related to the first, as both draw our attention to the fact that we seem to require a base of non-functional properties. However, rather than focusing upon how generalised functional role theory is supposed to illuminate what properties are positively, this draws our attention to the peculiarity of the view that properties are functional all the way down. In the previous section, we saw that the functional role theorist's strategy could not be extended to all properties, even granted an infinite number of properties. But here we find that with an infinite number of properties, it does at least make sense to say that properties are "functional all the

¹⁸ For the quote see Block 1990 p.166. Also see Blackburn 1991 and 1993 for a discussion of this idea.

way down”.¹⁹ For the characterisation of functional properties allows for the possibility of their being realised by further functional properties. Therefore, the problem is avoided if there is always another functional role property there to occupy such-and-such a functional role.

Unfortunately, as this hierarchy of functional properties can never be terminated, we are forced to postulate an infinite number of properties. This commitment is highly undesirable, since the issue of how many (concrete) properties there are seems an empirical one. If our best scientific theories inform us that there are an infinite number of such properties, then fine. But to decide that matter a priori on the basis of a certain metaphysical theory of properties seems dubious. This is especially so when we remember that generalised functional role theory isn’t just committed to an infinite number of properties, it is committed to an infinite number of *instantiated* properties. Why? In order for the functional property of, say, being red to be instantiated in this world, there has to be an infinite number of instantiations of other functional properties. For there has to be a functional property which stands in red’s functional role, call that P, another to stand in P’s functional role, call that Q, another to stand in Q’s functional role, and so on ad infinitum. As a result, we have to say that the universe is infinite and this, I suggest, is a rather hefty and surprising commitment to get from our theory of properties.

Moreover, postulating an infinite number of properties doesn’t eradicate the strangeness of the view. For there is another source of resistance to the claim that properties are “functional all the way down” which, historically, has been very influential, cropping up in objections to phenomenalism and behaviourism, as well as CTP.²⁰ The worry, although admittedly rather vague, can be put something like this: if all properties just consist of what other properties will do given certain circumstances, there seems to be nothing in the universe actually doing the causal work. For everything is relying upon something else and so on ad infinitum. As Blackburn puts it, “We can head toward the engine room, perhaps, but never get there” (1993, p.229).

¹⁹ Although, as I said in the last section, it is not clear that the assumption that there is such an infinite ordering of properties is justified.

²⁰ See, for instance, Berlin 1950 (phenomenalism), Geach 1957 (behaviourism) and Armstrong 1999a (CTP).

We could respond to this by pointing out that although properties no longer seem able to do anything, nevertheless, objects and events can at least be thought of as causally active. But CTP claims that the causal efficacy of objects at least is bestowed upon them by their properties. Given this, we can restate the worry as follows: if all properties just consist of what other properties will do given certain circumstances, there seems to be nothing *in* the object responsible for its causal powers. In the next chapter, we'll see that this thought underlies an important cluster of objections to CTP, where CTP is understood in a way less specific than as generalised functional role theory.²¹ I shall thus postpone the task of trying to pinpoint more precisely the content of the intuition being appealed to here, until this objection to CTP has been outlined.

²¹ See §1.2.

3. Grounding Causal Powers

3.1 The Grounding Intuition

Every causal transaction, according to Shoemaker, is a matter of things with certain causal potentialities bringing it about that these or other things have further potentialities, because properties are analysed as nothing but potentialities. In Scholastic language, we never get beyond potency to act. Act, so far as it goes, is just a shifting around of potencies. And is this a believable story? ‘Where’s the bloody horse?’ as the poet Roy Campbell might have said (Armstrong, 1999a p.31).

Here, Armstrong expresses a prevalent and deep concern against CTP. Most discussions of CTP mention it in one form or another,¹ but it is quite difficult to see exactly what the objection is all about. The basic idea seems to be this: if all instantiating a property involves is that a particular will do X in C circumstances, Y in D circumstances etc, properties are rendered mere promises of what will happen in the right circumstances. But then it looks like there is nothing ‘in’ the particular responsible for making it behave in the way that it does. Take, for instance, the property of being composed of copper. If we say that properties are clusters of potentialities or conditional powers, then we’re supposing that this property can be fully characterised by the fact that objects which instantiate it will be able to conduct heat and electricity (in certain circumstances), they will be resistant to certain kinds of pressure (in the right circumstances), they will be malleable...and so on. This analysis, however, seems counter-intuitive. For we tend to ascribe conditional causal powers to objects on the basis of what properties the object is said to instantiate. These properties are thought of as the ‘truthmakers’, ‘basis’ or ‘grounds’ of causal power ascriptions. Consequently, if properties are analysed as clusters of conditional causal powers, there seems nothing *in* the object - no categorical ground - responsible for an object’s causal powers.

There certainly does seem something counter-intuitive about this claim that properties are mere powers or potentialities, solely characterised by facts about what an object would do in such-and-such circumstances. But, if this is the case, what

¹ See, for example, Swinburne (1980), Armstrong (1996, 1997 and 1999a), Fales (1990), Jackson (1998), Elder (2001) and Campbell (2002a).

plausible-sounding thesis is CTP contravening? It is difficult to pinpoint a well-worked out thesis, but there does seem to be a deep-seated metaphysical intuition, which CTP appears to offend. This intuition, which I shall call ‘the grounding intuition’, appears in various places and guises within philosophy. Here are a few examples which bear witness to it:

what happens should be explicable in terms of the God-given nature of things. Natural laws are not as arbitrary and groundless as many think (Leibniz, 1988, p.205).

Dummett is not offering what Berlin was missing, namely a ground of the second kind for these subjunctive conditionals – a relatively abiding property of an object or place which could be used to explain his experiences...Berlin certainly put his finger upon a deep conceptual prejudice of ours that is offended by dispositional properties without categorical grounds (Evans, 1980 p.276).

It is the unique state, realising state, or array of magnitudes or tropes or instances of properties of points, that causes. It is here that the ‘making it happen’ happens: how, then, can we identify the cause by citing the relational dispositional or role-given properties with which physical thinking leaves us? (Blackburn 1991, p.238-9).

Evans and Blackburn talk about dispositional properties failing to have ‘categorical grounds’ or ‘realising states’, while Leibniz claims that what happens should be explicable in terms of the nature of things. The grounding intuition thus centres around the analysis of dispositions or causal powers. The idea, manifest in all of these passages, is that there must be something about the actual object, its intrinsic nature, for instance, which realises or provides the categorical ground for these dispositions.

In order to root out this intuition further, consider Ryle’s analysis of dispositional ascriptions, as he is one of the few who rejects the grounding intuition. He writes,

To say that this lump of sugar is soluble is to say that it would dissolve, if submerged anywhere, at any time and in any parcel of water (1949, p.119).

On this analysis, then, when we ascribe the disposition of solubility to a particular we are saying, ‘if x is put in water, then x will dissolve’. A particular is soluble just if this conditional is true of it. So far so good, it may seem, for when we ascribe dispositions to an object it looks as if we are interested in saying what the object will do in certain

circumstances, rather than in picking out any particular properties or states of the object which account for its behaviour. Nothing, as yet, prevents us from claiming that these conditionals are true of the object because of its intrinsic nature. However, Ryle's empiricism leads him to reject the idea that there is anything about the object which makes the conditionals true of it. According to him, dispositions are just complexes of conditionals, which inform us what sorts of events tend to happen in which circumstances.

It is this claim which many find difficult to accept. Surely there must be something about the object which accounts for the fact that the conditionals hold true of it? If we allow for the possibility, as Ryle's account does, that one object can have a disposition to X, while its duplicate has a disposition to not-X, then it looks as if dispositions are randomly imposed upon the object, because they have no grounding in its intrinsic nature. Many philosophers have found Ryle's account of dispositions unsatisfactory for just such reasons.² Some even talk of the account contravening a deep-seated intuition. Geach, for instance, writes,

When Ryle explains a statement of an actual difference between two men's mental states as really asserting only that there are circumstances in which one would act differently from the other, and apparently holds that this could be all the difference, he is running counter to a very deep-rooted way of thinking. When two agents differ in their behaviour, we look for some actual not merely hypothetical difference between them to account for this (1957, p.5).

The thought here is that the dispositional ascription should be made true by something actual, or existent in the object. This diagnosis is also echoed in Mumford's discussion, he writes,

Dispositions are actual, intrinsic states or properties rather than 'bare potentialities' and that to say something is now soluble is to say something about what it is like actually rather than something about possible future events (1998, p.74).

The general tenor of dissatisfaction with Ryle's view stems, I suspect, from this fact that it contravenes what I have been calling the grounding intuition. The

² See, for example, Armstrong (1996, 1997), Martin (1996), Jackson (1998), Prior (1985) and Mumford (1998).

negative train of thought which lies behind this intuition is the old adage – you can't get something for nothing. All these dispositions or conditionals we ascribe to objects can't be based on mere potentiality, something has to be 'making it happen'. The positive claim which goes along with this is the idea that there must be some "actual, intrinsic states" which make these conditionals or dispositions true of an object. Here is my, admittedly unclear, shot at characterising the grounding intuition:

The causal powers or dispositions of an object are grounded in or determined by the categorical or intrinsic properties of that object.

From this, it is clear that the grounding intuition is a metaphysical thesis about the nature of the causal powers or dispositions of particulars.³ It can be understood as a thesis about truthmakers, as it tells us what *makes true* those causal power ascriptions we ascribe to objects.

What motivates this thesis? One factor is the continuing attribution of causal powers to objects when they are unmanifested. Consider, for instance, Mellor's forceful example,

The safety precautions at our nuclear power station are intended to prevent an explosion by making impossible the conditions in which the fuel would explode. It is ridiculous to say that their success robs the fuel of its explosive dispositions and thus the precautions of their point (1974, p.116).

It would indeed be foolish in the extreme to forget that causal powers are persisting features of objects. But if we say that the explosive power of the fuel is just, as Ryle claims, the fact that certain conditionals are true of this substance, what reality does this power have when it is not being displayed? What justifies our continued attribution of its causal power when it is unmanifested?

Ryle's answer is this: an object's causal power persists over a period of time just in case a counterfactual or set of counterfactuals is true of that object over the specified period of time. This response, however, raises another, even more troubling question: what makes these counterfactuals true of an object? If we say, with Ryle, 'nothing - it is just a fact that these counterfactuals hold true of an object', then we

³ In what follows, however, I shall avoid using the terminology of dispositions for reasons outlined in §1.4.

have to reject a plausible metaphysical thesis, namely, Lewis's claim that "all contingent matters supervene on what there is, together with the pattern of instantiation of the fundamental properties and relations" (2001, p.614). For the counterfactual truths at a world will not supervene upon what there is in that world, plus its properties and their arrangements. So accepting Ryle's analysis means giving up the intuitive, and I would suggest fundamental, metaphysical thesis that truth supervenes on being.

This, of course, isn't anything like a decisive objection to Ryle's view. For anyone sympathetic to his position would, I presume, be happy to ditch this truthmaking principle. But I think it highlights the counter-intuitiveness of this analysis, thereby placing the grounding intuition in a positive light. Although much more should be said in support of this thesis, the rest of this section will be devoted to elucidating the grounding intuition further. For, ultimately, I wish to argue that CTP does not commit us to a counterintuitive account of causal powers. So in order to do this, I must first expound a way of viewing causal powers which is in line with the grounding intuition and CTP.

According to the grounding intuition, the truthmakers of causal power ascriptions are the 'categorical or intrinsic properties' of a particular. But what kind of properties are these? I don't think that we should pursue this notion of a categorical property because, as I argued in §1.4, I doubt that there is an important metaphysical distinction between dispositional and categorical properties. The notion of a property which is intrinsic to its object, however, looks more promising. For it seems to encapsulate the idea that causal powers are somehow grounded 'in' their objects. Can, then, the grounding intuition be understood in terms of intrinsicity?

Unfortunately, contrasting analyses of intrinsic properties muddy the water a little. The gloss usually given on an intrinsic property is one which an object can instantiate regardless of what is going on outside it. But once we try to spell out this idea further, the cracks begin to show. One very well known analysis of intrinsicity is put forward in Lewis's 1983b paper. There, an elite set of properties, the perfectly natural properties, are employed in an analysis of duplication. Two objects are said to be duplicates iff they have exactly the same perfectly natural properties. Intrinsicity is then defined in terms of this: an intrinsic property is said to be one which a

particular shares with all its duplicates. This view is problematic, however. For although intrinsic properties can be conjunctive, disjunctive etc, all the perfectly natural properties have to be intrinsic. Even if this does turn out to be the case, it doesn't look as if it should be made true by fiat.

Perhaps in light of this objection, Lewis and Langton (1998) have offered another analysis of intrinsicity which doesn't require such a strong assumption. This develops Kim's analysis (1982), which sticks closer to the original gloss on an intrinsic property. Kim had tried to define an intrinsic property as one which is compatible with loneliness, i.e. one which does not imply any contingent object wholly distinct from itself. As Lewis pointed out, however, this is not extensionally adequate, because loneliness is a property which meets this criteria but is not an intrinsic property.⁴ Lewis and Langton (1998) thus propose a modification of Kim's definition, arguing that a (basic) intrinsic property is one which a lonely and accompanied thing can either instantiate or not instantiate. But this analysis still makes essential use of the idea of duplication, as in order to define what a non-basic intrinsic property is, they appeal again to this idea that it is one which is had by a particular's duplicates (i.e. those particulars with exactly the same basic intrinsic properties).

This analysis of an intrinsic property as one that is sharable by its duplicates has come under attack from Dunn (1990). He argues that the property of being identical with α , for instance, is intuitively an intrinsic property. For it is one that "the object has in virtue of itself, depending on no other thing" (p.178). But it doesn't count as an intrinsic property on Lewis's accounts, because it isn't a property which is had by α 's duplicates. Similarly, the property of being a duplicate of α doesn't seem to be an intrinsic property, because this depends upon α 's relations to things external to itself. But this does count as an intrinsic property given the duplication analysis, because if χ is a duplicate of β and β has this property of being a duplicate of α , as duplication is transitive, χ too will be a duplicate of α .

Dunn's criticisms here draw on remarks made by Moore. Moore notes that there are two senses of intrinsic. He writes,

⁴ See Lewis 1983a p.113-114.

Two patches of colour may be exactly alike, in spite of the fact that each possesses a certain constituent which the other does not possess, provided only that their two constituents are exactly alike. And yet, in a certain sense, it is obvious that the fact that each has a constituent, which the other has not got, does constitute an intrinsic difference between them, and implies that each has an intrinsic property which the other has not got. And even when the two things are simple the mere fact that they are *numerically* different does in a sense constitute an intrinsic difference between them, and each will have at least one intrinsic property which the other had not got – namely that of being identical with itself. It is obvious therefore, that the phrases ‘intrinsically different’ and ‘having different intrinsic properties’ are ambiguous (1922, p.262).

One sense of ‘intrinsically different’ then, refers to all those entities which are not qualitatively alike. This is tracked by Lewis’s analysis, for by claiming that an object’s intrinsic properties are those that are instantiated by its duplicate, he ensures that intrinsic properties can be had by more than one particular. Hence, properties such as *being identical to α* fall by the way-side. The second sense of ‘intrinsically different’, in contrast, concentrates on differences in the ‘constituents’ of a particular. Any differences in the entities which lie entirely within the confines of the particular, count as changes in its intrinsic nature. This sense of intrinsic is best captured by Dunn and Humberstone’s notion of an intrinsic or interior property. The intuitive idea is that an intrinsic property is one whose existence and nature has been, in Humberstone’s words, “entirely determined by what is the case within the confines of the would-be possessor” (1996, p.242). Therefore, no object distinct from its instantiator can affect the property in any way.

Humberstone argues, and I agree, that both these notions of intrinsicality are legitimate. But due to their different extensions, we need to keep them apart. I think that Humberstone fails to realise how different these two notions are, however. He treats the category of interior properties as more inclusive than Lewis’s intrinsic properties.⁵ For it includes non-qualitative properties, such as the property of having d as a part, as well as qualitative properties. But the difference between the two seems to go far deeper than this, for Lewis’s analyses exclude the possibility of intrinsic property instances. Let me explain. A property instance, as I mentioned earlier, is a single instantiation of a property. As such, it cannot be had by more than one object.

⁵ Although there is the aforementioned exception of the property of being a duplicate of α .

So if an object has a property instance of F-ness, because this property instance is unique to that object, it cannot be had by a duplicate of that object. Similarly, Humberstone's notion of an interior property, seems to exclude intrinsic universals. For an interior property is characterised as one which is wholly internal to its object, but a universal exists (or could exist) at different places at the same time in the universe. So we shouldn't say that it is the *universal* of redness which is wholly present within the confines of an object, since this universal is present in all sorts of different objects. What we should say instead is that it is an *instance* of this universal which is interior to that object.⁶

What characterisation of an intrinsic property best fits the intuitive gloss given earlier? Perhaps surprisingly, we find that it is this notion of an interior property instance. Take, for instance, Vallentyne's (1996) way of spelling out this idea.⁷ He begins by stating that a contraction of a world is a world which is obtainable from the original one by removing objects from it. A maximal contraction of a world (an x-t-contraction), on the other hand, is one which has contracted as much as is possible while still leaving a specified object at a specified time. Intrinsic properties are then defined as follows:

P is intrinsic = for any world w, any time t, and any object x: (a) if Px at t in w, then Px at t in each x-t-contraction of w, and (b) likewise for ¬P (p.212).

This definition of an intrinsic property not only includes those non-qualitative intrinsic properties excluded by Lewis's account, it also seems more appropriate to say that it tracks intrinsic property instances rather than universals. Why? If we contracted a world until it contained just one red apple then, arguably, we would be left with a property instance of red, rather than the universal of redness. This claim is qualified because on Armstrong's theory of universals, this wouldn't be the case. For the universal of redness, as well as the property instance of redness, would be present in the object.⁸ But if universals are thought of as abstract universals, or sets of tropes, or sets of possible particulars, then the contractions would rid us of the abstract realm of

⁶ For more on this distinction between universals and their instances, see §3.3.

⁷ Other ways of spelling out this idea are given by Kim (1982) and Yablo (1999).

⁸ See Armstrong 1997. In §5.4, however, I shall argue that Armstrong's account of universals is problematic.

Platonic universals, or the sets of tropes/possible particulars. So at least given some analyses of properties, the standard gloss on intrinsicity seems to be tracking property instances, not universals.

The key question for our present purposes, however, is this: what sense of intrinsicity is being appealed to in the grounding intuition? Earlier, I claimed that the notion of an intrinsic property was promising, because it seemed to encapsulate the idea that the causal powers of an object depend solely upon the nature of that object. So an object's powers don't rely upon facts about what's going on in other possible worlds, or in other times and places in this world. Now this notion of a property which is internal to its possessor, isn't well captured by Lewis's analysis of intrinsic properties. The universal of redness, for instance, isn't something which is wholly confined to its possessor, since it is present at the same time in all sorts of different objects across the universe. In order to get at this idea of an entity which is internal to one object, then, we need instead to appeal to this notion of an interior or intrinsic property instance. Since this manages to capture the idea of an entity which is wholly present in its object. By making these intrinsic property instances the truthmakers of causal power ascriptions, therefore, we preserve the intuition that the causal powers of objects depend solely upon the natures of those objects.

Although much more would need to be said in order to develop a complete analysis of intrinsicity as interiority, I shall take this notion of a property which is internal to its object, and hence will remain a feature of that object in the absence of all other entities, to be the kind of entity which the grounding intuition invokes. Our task is not yet done, however. For even more problematic than the notion of an intrinsic property, is the relation of grounding or determination which the grounding intuition utilises. In other areas of philosophy, the relation of supervenience is often employed to capture the idea of one set of entities determining another. So we can equally apply it here: take a set of intrinsic properties P_{1-n} , and a causal power Z , both of which are instantiated by an object. We can say that the causal power Z supervenes upon the intrinsic properties P_{1-n} if (given the laws of nature⁹) the set of properties P_{1-n} are sufficient but not necessary for causal power Z .

⁹ This wouldn't be required, however, granted the thesis that the causal powers are essential to properties (see chapter seven).

This supervenience thesis seems fitting for the realms of properties and causal powers. Small variations in the properties of an object can result in a change in that object's causal powers. Conversely, instantiating similar properties results in similar causal abilities, so the sufficiency condition looks plausible. The same goes for the not-necessary condition. Shoemaker, for instance, cites an example in which two substances are both poisonous, but one kills by injuring the heart, and the other by damaging the nervous system.¹⁰ In such a case, the causal power of being poisonous will be realised by two different sets of properties. So an object can display this causal power in the absence of any particular set of properties.

Despite the plausibility of the claim that there is a supervenience relation between properties and causal powers, this relation cannot fully capture the content of what is meant by the grounding or determination relation. For the supervenience relation could hold in the absence of the determination relation. Consider, for instance, a world in which God, in the style suggested by act occasionalists, directly intervenes on every occasion where the right circumstances are present, to ensure that certain sets of intrinsic properties are always correlated with certain effects. God's acts would guarantee the supervenience of causal powers on intrinsic properties in this world, but they wouldn't preserve the intuition that the intrinsic properties of objects *determine* their causal powers. For the properties are not what are doing the fixing, the will of God is fulfilling this role. The problem with supervenience is that it is not, as Kim writes, "a metaphysically deep relation" (1998, p.14). It only records how two sets of entities covary, it does not explain why such correlations hold. At best, therefore, supervenience can be thought of as symptomatic of the determination relation, it cannot provide us with an analysis of what this relation involves.

Are there any other ways of spelling out the relation of determination or grounding? One notion which might at least help convey the kind of relation required is that of constitution. This relation provides a useful analogy, because it manages to capture the asymmetry of the grounding relation. If these words constitute this thesis, for instance, we cannot say that this thesis constitutes these words. Yablo offers this analysis of the constitution relation. He writes,

¹⁰ See Shoemaker, 1980a p.211.

- x constitutes y iff
- (a) x coincides with y,
 - (b) any part of x essential to it has parts that are not essential to y,
 - (c) no part of y essential to it fails to have parts that are essential to x (1999, p.491-2).

The basic idea then, is that the relation of constitution holds when two entities have all their parts in common, but whilst one of them could exist without some of its parts, the other could not. Take, for instance, the famous ship of Theseus. We can suppose that at the start of its life, it was made up of a number of different planks of wood. Each of these planks was essential to what we could call the aggregate of planks, because any change in its parts would result in a different aggregate of planks. But not all the planks were essential to the ship, because it survived many changes to its planks. So, in Yablo's words, the aggregate of planks "hugs" its parts more closely than the ship because, unlike the ship, all its parts are essential to it.

How might this help illuminate the relation of grounding? We can liken the relation of grounding to that of constitution. When a collection of an object's properties, call this collection P, ground or constitute an object's causal power C, we can suppose that the parts of the collection P and the parts of the causal power C overlap completely. This may sound very strange – although it's okay to think of a collection of properties having parts, it doesn't seem sensible to say that a causal power has parts. However, if rather than focusing on the detachable parts of a chair or table, we think instead about the parts or ingredients of a cake, the analogy still proves useful. For we tend to say that sugar, eggs, flour etc. constitute or make up a cake, even though they are unlike the legs of a chair as they combine to form a unified whole, in which the different parts/ingredients are no longer separable or distinguishable.¹¹ Similarly, then, we might think of the properties of an object as like

¹¹ Another analogy is found in the debate about the compositionality of sense. There, Dummett argues that "the sense of a complex expression has to be regarded as made up of the senses of the constituent words" (1973, p.25). He thus treats the senses of words as the proper parts of a sentence. But Geach argues that while Frege does talk about thoughts being built up out of its components parts, "this way of thinking should to my mind be charitably expounded, not imitated" (1975, p.149). For some of the senses of thoughts are not independent objects which stand alone. At least some of them must be "unsaturated...otherwise they would not hold together" (1975, p.149). On Geach's view, then, we need to reject this idea of senses as the independent, complete parts of a thought, in order for these senses to form "a unity", in which "the sense of names and simple predicates stick together to form a thought" (1975, p.150). Instead, the senses of thoughts are much more like ingredients. For they combine to form a unified whole, in which the parts are not wholly independent and complete in themselves.

the ingredients or parts of a causal power. Certain properties of an object combine to form one of its causal powers, an entity whose parts or ingredients are no longer distinguishable.

Once we have this in place, the rest of Yablo's characterisation goes through. Every part or property in the collection P is essential to it because we can suppose that if any part of it changed it would result in a different collection. c) is thus met because if a part/property is essential to the object's causal power C, it will also be essential to the collection P. We can also claim that a slight variation in the collection P need not lead to any change in the object's causal power C. For so long as the object continues to display the features characteristic of C, we can suppose that it still possesses causal power C. So b) is met because some parts/properties are essential to the collection P but not to the object's causal power C. Although this offering remains a picturesque suggestion rather than a proper analysis of the grounding relation, I think that the analogy between the relation of constitution and the grounding relation proves helpful. For unlike the relation of supervenience, that of constitution does at least convey the notion of determination we're after, as it captures the thought that these entities somehow *make-up* these other entities.

This suggestion, however, still leaves an important issue outstanding: does this relation of grounding involve the laws of nature or not? On the basis of what has been said, we could formulate these two theses:

- (a) The weak grounding relation: given the laws of nature, the intrinsic property instances of an object constitute its causal powers.
- (b) The strong grounding relation: the intrinsic property instances of an object constitute its causal powers.

Depending on what analysis of the laws of nature we adopt, the weak grounding relation may temper the force of the grounding intuition somewhat. For the causal powers of an object will only depend solely upon that object, if the laws governing it are made true by its intrinsic nature.¹² Otherwise, these causal powers will also depend upon how we can best codify the regularities in this world, or upon relations between

¹² For more on this, see chapter five.

Platonic universals, etc.¹³ The strong grounding relation is, then, more in line with the idea that the causal powers of an object depend solely upon that object. So any account which incorporates (b) will certainly have preserved the intuitive force of the grounding intuition. However, as (b) will no doubt be too strong for many tastes,¹⁴ I shall assume that the original grounding intuition would be satisfied if this weak grounding relation holds between the intrinsic property instances of objects and their causal powers.

The grounding intuition, therefore, does encompass different theses about the relation that holds between these intrinsic properties of objects and their causal powers. But the main point to take away from this section is not the specific detail about how we should analyse the grounding intuition. Instead, what I want to stress is that the grounding intuition captures, in Geach's words, "a very deep rooted way of thinking". So if Armstrong and others are right to say that CTP contravenes this intuition, then the plausibility of this theory is undermined.

3.2 CTP in Trouble Again

Earlier, I touched upon why CTP seems to jar with the grounding intuition. The worry was that if properties are analysed solely in terms of how they make their objects behave, the causal powers of an object will not be determined by entities intrinsic to that object, as properties will be nothing more than their relational aspects. So the entities grounding the causal powers will not be intrinsic to their object, since they are characterised by their relations with entities outside the object. This conclusion is too quick, however. For we've seen that CTP cannot be treated as a unified theory, so we have to look at whether this objection has any weight against all, some or none of the proposed CTP analyses.

First, does CTP, understood as an individuation thesis, contravene the grounding intuition? No. The fact that properties are individuated by their causal features, in no way shows that these causal features exhaust the nature of properties. To say this, would be like claiming that because the spatiotemporal properties of

¹³ See, for instance, Lewis's best systems analysis (1983b) and Tooley's analysis of laws (1987). More will be said about these accounts and the issues they raise later.

¹⁴ See, for instance, Armstrong 1996.

physical objects individuate these entities, there is nothing more to these physical objects than their spatiotemporal properties. This is clearly a mistake. We can say that a property is individuated by how it relates to other entities, while nevertheless maintaining that properties are something more than their relational aspects. But doesn't the mere fact that properties have these relational aspects prove that they are not intrinsic to their objects? No, no more than the fact that a proper part of an object, like the leg of a table, stands in relations to entities outside that table shows that this isn't interior to that object. So the view that properties are individuated by their causal roles doesn't render them unsuitable for the purposes of grounding causal powers.

Do matters change at all when we say that these causal features, as well as individuating properties, are also essential to them? Again, no. Let's suppose that we accept Kripke's claim that in order to be me, my parents had to be Bob and Kate Whittle. So in all possible worlds where I exist, my parents are Bob and Kate Whittle.¹⁵ This clearly does not demonstrate that this is the only property that I instantiate, and thus that it exhaustively characterises me. Just because certain features are essential to some entity, doesn't mean that this is all there is to an entity. Without this claim, CTP won't contravene the grounding intuition. For we can always say that there are other aspects to a property, which make it appropriate to characterise them as intrinsic to the objects that instantiate them. So Armstrong's objection has no force against weak CTP, as this should be understood as a transworld individuation thesis for properties.

Does CTP, understood as a generalised form of functional realiser theory, contravene the grounding intuition? It should be pretty clear that it doesn't, as functional realisers do not claim that properties are exhausted by their causal characterisations. All they claim is that properties can be identified and tracked by their causal features. This leaves plenty of room for there being further characteristics to properties, which render them suitable for the task of grounding object's causal powers. Moreover, if we go on to supplement the functional realiser's RCL definitions with claims about how properties are individuated and what features are essential to them, in order to make the thesis more characteristically CTP-like, the resulting theory

¹⁵ See Kripke 1972, lecture III.

will still be immune to Armstrong's grounding objection. For this just makes CTP a transworld individuation thesis. Therefore, for reasons just given, these variations on the generalised functional realiser theory will not jar with the grounding intuition.

What of CTP understood as a generalised form of functional role theory? Does this contravene the grounding intuition? If we ignore, for a moment, doubts about the coherence of this interpretation, then this way of understanding CTP seems a likely target for Armstrong's attack. For functional properties are properties which consist in *other* properties implementing a certain causal role. But if we generalise this account and make all properties functional then, granted we don't postulate an infinite number of properties, there will be no properties left which can realise the causal relations definitive of functional properties.¹⁶ As a result, there will be no properties, intrinsic or otherwise, left to stand in these causal roles. So there will be no properties of the object which could be seen as grounding its causal powers (and functional properties).

The grounding objection raised by Armstrong, therefore, is clearly another statement of the one made earlier against functional role theory. The worry there was that if all properties consist in what other properties do, there will be nothing in the universe which is actually doing the causal work or grounding objects' causal powers. Functional role theory, however, isn't alone in being susceptible to the grounding objection. It also pertains more generally to strong CTP, understood as the view that the causal profile of a property is all there is to that property. Why? If a property just consists of the fact that the object which possesses it will do x in certain circumstances, y in certain other circumstances etc, then properties are rendered, in Armstrong's words, "congealed hypothetical facts or states of affairs" (1997 p.79). We are left with a view like Ryle's. For while we can still claim that an object instantiates a certain property just in case certain counterfactuals are true of it, these won't be intrinsic states of that object which can ground its causal powers.

In light of this and the other objections raised here, it may appear sensible to cut our losses and reject strong CTP of any form. As weak CTP is not affected by these difficulties, some variant of this theory can be embraced instead. This conclusion, however, is not forced upon us, since the reductive two-level criterion

¹⁶ See §2.5.

outlined earlier is still standing. Moreover, in the next sections, I hope to show that the serious objections which beset other formulations of strong CTP can be overcome. I'll begin by looking at the strategic problem for functional role theory, since this has to be resolved for the analysis to work.

3.3 A Solution

Earlier we saw that if we try to generalise the use of functional role RCL definitions to all properties, then the functional role theorist's strategy cannot be implemented. For in order to preserve what is distinctive about this view, there needs to be a number of properties which aren't analysed via functional role theory. This makes the prospects of a generalised form of functional role theory seem dire. For if we're forced to postulate this domain of non-functional role properties, this means abandoning the idea of extending this theory to all properties.

Is there any other way out? I think that there is, but in order to see this solution, it will help to draw our attention to the aforementioned distinction between properties and their instances. It has already been noted that the term 'property' slides between two usages. Sometimes talk of properties refers to universals. These are unifying entities which can be instantiated in different particulars at the same time. It is, for instance, the universal of redness or roundness that makes different particulars alike in certain respects. In other contexts, the term 'property' is used to talk about instantiations of properties or property instances. Used in this way, the term does not refer to a universal which can be instantiated in different particulars at the same time, rather it picks out a single instantiation of that universal. These instances are particulars because they cannot be instantiated in more than one place at the same time. They are thus inextricably tied to the particular that instantiates it – a property instance is *that* cable's weakness or *that* ball's redness.

Regardless of what stance we adopt towards the ontology of properties and their instances, attention needs to be paid to this distinction between two senses of the term 'property', as they name different kinds of entities.¹⁷ Take, for instance, two

¹⁷ In what follows, unless otherwise stated, I shall use the term 'universal' or 'property' to refer to the first sense of property and the term 'property instance' or just 'instance' for the second. I do not suppose anything further by these terms. In particular, realism about universals is not implied, nor is

chocolate cakes which are both the same shade of brown. These cakes are alike in that they each instantiate the property of being that particular shade of brown, but is the colour of one cake identical to that of the other? If we think of properties in the first way, then the answer is yes, because they both share the same universal (whether universals are understood realistically or are reduced to something else). If, on the other hand, we think of properties in the second way, then the answer is no, because we have two instantiations of the universal of this shade of brown and, hence, two property instances of it.

This distinction between properties and their instances gives us more room to manoeuvre. Earlier we saw that functional role theory demands both realising properties and functional role properties. As these functional properties are logical constructions out of their realisers, these realisers have to be distinct from the functional properties they display. Consequently, once we've extended the functional role analysis to all properties, there is no longer a domain of properties which can act as realisers. Now, however, we can see that this conclusion is too hasty, as it overlooks the distinction between properties and their instances. Once this distinction is noted, a possible solution reveals itself: Properties qua universals can be thought of as functional role properties, while instances of these properties can be the entities which realise these RCL theories.

This proposal looks promising, as we've seen that CTP is a theory about (concrete) properties qua universals. So what matters is whether functional role theory can be generalised to all (concrete) universals. But the suggestion needs spelling out further. For if property instances are made the realisers of RCL roles, there will be many concrete realisations of these theories. If, for instance, *this* knife's sharpness can realise a particular functional role, the sharpness of this other qualitatively indistinguishable knife will also be able to. The multi-realizability of these functional roles threatens the uniqueness required for these implicit RCL definitions to work. But this can be regained by making the same move as functional role theorists in the philosophy of mind. So by taking all the concrete realisations of a particular RCL theory, we can abstract what is common to all of them by considering the equivalence

realism about tropes. Property instances are meant to be neutral between instantiations of universals, tropes, members of sets, etc.

class of concrete realisations under the relation, sameness of functional role. These abstracted functional properties can then ensure the uniqueness of the RCL definitions.

If we embrace this account, then property instances are the entities which realise these functional roles, while properties qua universals are identified with abstracted functional entities, whose very essences are specified by these RCL theories. So far so good, but there is still a piece of the solution missing. For something more must be said about the nature of these property instances, in order for the analysis to have the constructive character definitive of functional role theories. Why? If the realisers of the RCL definitions are not distinct from the properties they analyse, functional role theory will collapse into a form of functional realiser theory.¹⁸ It is only when the functional role properties are distinct from their realisers, that it makes sense to suppose that the entities being defined by the RCL definitions are logical constructions out of the realisers of these definitions.

What consequence does this have for a generalised form of this theory? As the realisers of the RCL definitions must be entirely distinct from the entities being defined, this significantly restricts the analysis we can offer of property instances. They cannot, for instance, be analysed as instantiations of universals. For, on this view, property instances are constructed out of complexes of particulars and properties. So it would be circular to then attempt to construct properties out of property instances. Another position which is ruled out is the view that property instances are members of sets of possible particulars. Given this analysis, property instances cannot be thought of as something independent of the sets of possible particulars of which they're part. For this makes property instances far too coarse-grained, as the members of these sets are possible particulars, and possible particulars clearly instantiate numerous properties. In order for these entities to count as property instances and not just particulars, therefore, they have to be thought of as part of a particular set or universal. This leaves us with the same difficulty: If property instances are not entirely distinct from the properties which they are instances of, we cannot construct properties out of them without circularity.

¹⁸ See §2.4 for an explanation of why this is.

In order to avoid this problem, what we require is an account of property instances which views them as self-standing entities, independent of the properties of which they're instances. Does any account fit this criterion? If we think of property instances as *sui generis* entities, in the way conceived of by trope theorists,¹⁹ then property instances are not analysable in terms of anything more basic, like instantiations of universals or members of sets of possible particulars. Property instances or tropes, on this view, are very fine-grained entities, which exist independent of the universals of which they're instances. Indeed, most trope theorists believe that tropes are ontologically more basic than universals, since universals are analysed as sets of tropes. So it looks like these tropes could be the realisers we're looking for.

But doesn't this claim that universals are sets of tropes generate the same difficulty? Since if tropes are parts of universals, the view seems analogous to the claim that property instances are members of sets of possible particulars. These two positions, however, are not alike in the relevant respect. Although tropes can be thought of as parts of universals, they nevertheless exist independent of the universal of which they're part. For tropes are not analysed in terms of these universals, rather the reverse is the case – universals are analysed in terms of tropes. On the sets of possible particulars view, however, property instances do not exist independent of the set of which they're part, since it is being part of a particular set/universal that makes them property instances. So property instances are analysed in terms of universals, rather than vice versa. This is what makes the difference. Unlike property instances *qua* tropes, property instances *qua* possible particulars cannot be thought of as independent of the universals or sets of which they're part – hence the problem – these entities are not wholly distinct from the entities which they are supposed to construct.

With an ontology of tropes, however, there is a way of making sense of generalised functional role theory. Tropes are not partly composed of universals, nor are tropes analysed in terms of them. Moreover, CTP theorists do not claim, of these entities, that they are analysed via functional role theory. So tropes can be the entities which realise the functional roles stated in the RCL definitions, and properties can be abstracted from classes of tropes, all of which realise the same functional role.

¹⁹ See, for instance, Stout (1921) Williams (1953) and Campbell (1981, 1990).

If we accept this theory, then universals become logical constructions out of tropes. The universal of F-ness, for instance, is identified with the set of tropes which realise the functional role specified in F-ness' RCL definition. The parallel between this unrestricted form of functionalism and functional role theory within the philosophy of mind should now be clear. In the philosophy of mind, functional role theorists postulate first-order physical properties, which mental properties are taken to be logical constructions of. A mental property, on this view, is instantiated just in case there is some physical property occupying the functional role specified by that mental property's RCL theory. Similarly, here we find that tropes are analogous to the physical properties, as they are not logical constructions out of any other entities. The universals, by contrast, are comparable to the mental properties, as they are logical constructions out of tropes. So the generalised RCL definitions are ontologically creative, like they are in the philosophy of mind, as they pick out new entities (the universals), which are constructed out of the entities that realise these RCL roles (the tropes), but which are nevertheless distinct from them.

We can, then, generalise functional role theory to all properties, if we are prepared to endorse an ontology of tropes. For tropes can be the entities which realise the RCL functional roles – problem solved. The question remains, of course, why anyone would want to endorse such a reading of CTP. I hope that an answer to this will emerge in the course of the discussion. But, as a first shot, we can say this in its favour. First, this reading of CTP offers an ontological analysis of properties, for properties are identified with sets of tropes which meet certain conditions specified in their predicates' RCL definitions. Second, it commits us to this two-level criterion of identity for properties: two tropes are instances of the same property iff they both occupy the functional role specified in that property's RCL definition.²⁰ Third, we are given a novel account of how tropes are sorted into genuine kinds. Although everything is similar to everything else in indefinitely many ways, we can distinguish between genuine and non-genuine similarity via the tropes of objects. Those tropes which stand in the causal relations definitive of a property form a set, which grounds genuine resemblances between objects. Finally, in the next section, I hope to show that

²⁰ See §1.5 and §2.3.

this account isn't subject to any of the metaphysical or grounding worries previously outlined.

3.4 More Solutions

With the distinction between universals and property instances in view, I think that the other problems which strong CTP appears vulnerable to can be overcome. Let's begin by considering the metaphysical worry generated by the claim that properties are "functional all the way down". Such a thesis appears to require an infinite number of properties to realise the functional roles outlined in these RCL definitions. But an ontology of tropes relieves us of this commitment. For functional role theory doesn't claim, of these entities, that they consist in other entities realising a certain causal role. This is only said of properties qua universals. So tropes can be the entities which stand in the causal relations definitive of a particular universal.

In order to avoid this particular problem, however, an ontology of tropes isn't strictly required. We can make do with less than this, since property instances, no matter how they are analysed, would do the job as well. This is an academic point really, since in the previous section we saw that generalised functional role theory requires tropes for the functional role strategy to work. But it is perhaps worth noting that property instances, no matter their stripe, would suffice here. For so long as there is some entity which can stand in the causal roles specified in the RCL definitions, there will be something which can do the causal work in our metaphysical picture. It doesn't matter if the entity standing in that role isn't entirely distinct from the functional universal since, granted it is not identical with that universal, it can satisfy the condition of being the entity which realises the causal roles definitive of a functional universal. So there is no need to postulate an infinite number of functional universals.

This distinction between universals and property instances also gives CTP theorists the resources to deal with the grounding objection. Earlier we saw that the notion of intrinsicity, which captures the intuition that the causal powers of objects depend solely upon the nature of those objects, leads us to the idea of an intrinsic or

interior property *instance*.²¹ For an intrinsic property instance is one which is wholly internal to its object, but a property qua universal exists (or could exist) at different places simultaneously in the universe. So we shouldn't say that it is the *universal* of being twenty stone which grounds this object's causal power to crush fragile things, as this universal is present in all sorts of different objects and so it is not intrinsic (in the sense we're interested in) to any particular one. What we should say instead is that it is an *instance* of this universal which is intrinsic to that object, and hence grounds its ability to crush fragile things. The grounding intuition should thus be rewritten as follows:

The causal powers or dispositions of an object are grounded in or determined by the intrinsic property instances of that particular.

This puts a slightly different slant on matters, since strong CTP and the functional role reading of CTP is a theory about properties *qua universals*. It states that the nature of a property or universal is exhausted by its causal characterisation, it is just the property of having some other entity which satisfies such-and-such a causal role. All that is said about property instances is what it is to be an instance of the property F-ness, rather than the property G-ness or H-ness. But this provides us with information about how we should individuate properties, not their instances.

Again we find, then, that by disambiguating between properties and their instances, CTP theorists can respond to the grounding objection. Strong CTP and functional role theory claim that universals are exhausted by their causal powers/causal role characterisations, but the same is not said of property instances. We are thus free to adopt another account of them. The grounding intuition, on the other hand, properly understood, says that it is the property *instances* of an object which ground or determine its causal powers. The fact that properties qua universals are not the sort of entities which can play this grounding role, therefore, is not a problem. So long as strong CTP and functional role theory do not exclude an analysis of property instances which can ground causal powers, they need not be at odds with the grounding intuition.

²¹ See §3.1.

Property instances, therefore, do offer a way out of the grounding objection. But is any particular analysis of property instances required? I think that more needs to be said about the nature of these property instances, in order to show that they are able to ground causal powers in a way which accords with the grounding intuition. If we hold a Platonist view of universals, for instance, and think that property instances are instantiations of these universals, it is somewhat strained to say that property instances are intrinsic to their objects, in the way suggested here, since they partly consist of an entity which is not present in space and time.²² Tropes, however, are not the only kind of property instances which would be suitable for this role of grounding an object's causal powers. A theory of universals and their instances like Armstrong's would also suffice. For on this view, universals are thought of as spatiotemporal entities. So an instance of a universal will be internal or interior to the object which instantiates it.

Adopting functional role theory or strong CTP, therefore, does not commit us to a counter-intuitive view of causal powers. For the analysis of causal powers suggested earlier, namely that causal powers are constituted out of the intrinsic property instances of objects, is consistent with strong CTP, functional role theory and the grounding intuition. Instead of having bare potentialities, in the way that Armstrong suggests, CTP theorists can maintain that these powers/potentialities are grounded in the intrinsic property instances of objects. These property instances are what justify our continued ascription of powers to objects when they are not being manifested, and so are responsible for the truth of the causal counterfactuals we ascribe to objects. On the view being proposed here then, if two objects differ in causal powers, and so in what counterfactuals we can ascribe to these objects, then their intrinsic property instances will also differ. Consequently, causal counterfactual truths will supervene upon what properties there are at a world and the way those properties are arranged, in accordance with Lewis's truthmaking principle.

So the functional role reading of CTP is not hopeless after all. For all the objections which were levelled against this account can be overcome. The problem for strong CTP has also been resolved, but because of the unclarity which surrounds the claim that properties are exhausted by their causal roles or properties are causal

²² The same also seems true of property instances viewed as members of sets of possible particulars. But I'll say more about this in chapter five.

powers, more needs to be said about the content of this thesis. From the preceding discussion, however, these central strains of CTP can be identified:

Strong CTP: Functional role reading - This combines the reductionist two-level criterion of identity for properties with the RCL technique. So properties are nothing over and above equivalence classes of tropes under the relation, *sameness of functional role*.

Weak CTP: Supplemented functional realiser reading – This combines the non-reductionist two-level criterion of identity with the RCL technique. So there is a necessary correlation between the identity of a property and its instances under the relation, *sameness of functional role*.

Henceforth, unless otherwise stated, weak CTP and strong CTP will refer to the theses stated here. There will be a bias towards strong CTP in the discussion, however, because I think that functional role theory offers a more ambitious and informative analysis than the other readings looked at. Moreover, it provides a context in which to understand the theory, while resulting in commitments which are central to CTP.

What's next? My aim is to investigate CTP further. It is clear that CTP implies much for an account of causation, as it does for an account of properties. Here, for instance, I've tried to show that CTP is not committed to a counter-intuitive analysis of causal powers. In what follows, I shall continue looking at CTP's ramifications for a theory of causation. My hope is to show that these ramifications are both defensible and plausible, thus adding weight to CTP.

4. The Relata of Causation

Fire...does not warm, because it is a body, but because it is hot; nor does one body put forward another body because it is a body, but because it moved into the place of that other body. The cause, therefore, of all effects consists in certain accidents both in agents and in the patients...CAUSE simply, or an entire cause, is the aggregate of all the accidents both of the agents how many soever they be, and of the patient, put together.¹

In this chapter, I shall defend the claim that property instances are the fundamental causes and effects.² They are the entities which the relation of causation relates. But before I do this, I first want to explain why I think it is important to support this claim here. This, I hope, will also serve to clarify the thesis being proposed.

4.1 CTP's Causes and Effects

A defence of a causal ontology of property instances finds its way into the discussion here because of its links with CTP. According to CTP, an object's causal powers are bestowed upon it by its property instances. Put another way, we can say that the range of possible causal interactions an object can stand in is determined by its property instances. This means that objects are causes *in virtue of* their property instances, as their ability to function as causes is bestowed upon them by their property instances. So they should be treated as non-fundamental causal relata.

Why should events be viewed as non-fundamental causal relata?³ There isn't a direct link from CTP to this view, but commitments implicit in CTP do lead us most of the way there. Nearly all the events which are causal relata involve objects.⁴ Now,

¹ Hobbes (1839, p.121-122). Clatterbaugh (1999) notes that during the causation debate in modern philosophy, the scholastic's claim that causes and effects are substances, was gradually replaced with the view that causes and effects are accidents or properties of substances.

² By property instances here, I do not mean to presuppose any particular ontological analysis of them. So in order to defend strong CTP, we would need to go further and argue that property instances qua tropes are the causal relata. This task, however, shall be postponed until chapter five.

³ Unless those events are Kim's exemplifications of properties at times (1973, 1976), for then it looks like they are just property instances. Matters are complicated by the fact that there are different theories of events, but I am talking broadly here. The suggestion is that events are non-fundamental causal relata whether you endorse the Aristotelian view of events as changes in objects, the Quine-Lewis view of events as regions of space-time, or the Davidsonian view of events as concrete, basic particulars.

⁴ Some philosophers have argued that events must always involve objects, as they are changes in the properties of objects (see, for instance, Lombard 1986 and Lowe 2002). I see no reason to exclude the possibility of subjectless events, however, and so the argument above will not encompass them. But in

according to CTP, the way an object can behave is determined by the properties it instantiates, for these are what determine the object's range of possible causal interactions. Take, for instance, the event of dropping the sugar-cube in water. CTP claims that the properties of the sugar-cube determine its range of possible interactions, so they are what make it true that if this sugar-cube is dropped in water, it will interact with the water and dissolve. The event's efficacy in bringing about the dissolving of the sugar-cube is thus dependent upon the property instances of the object it involves.⁵ For even given the event of dropping the sugar-cube in water, the sugar-cube would not have dissolved if it had instantiated different properties (properties that couldn't ground the sugar-cube's solubility).

This isn't the whole story of course, as which of an object's possible causal interactions are actualised depends upon what happens in the world extraneous to that object. Events are frequently cited as 'the cause' because they are usually complex entities, which encompass changing relations between objects and changing conditions in their environments. But, still, there's reason to think that an event's efficacy depends upon the persistent and changing properties and relations of the objects and environments they involve, plus those properties which we ascribe to the event rather than to any subjects of those events.⁶ At any rate, such a position is strongly suggested by CTP, because CTP claims that property instances are the locus of causal activity, at least when it comes to causation involving objects. So as most events concern objects, whose behaviour is determined by their property instances, there is reason to think that events too are causes and effects in virtue of their (and their subjects) property instances.

the interests of simplicity, it seems sensible to treat subjectless events in the same way as other events if possible.

⁵ For more on this notion of involvement, see Lowe (2002). It is not difficult to understand what is intuitively meant by it, however. We can say that in the marriage of Joel and Ann, for instance, both subjects, Joel and Ann, were involved in this event.

⁶ We need to distinguish between the properties of events and those properties of entities involved in events. In 'the door's slamming *violently*', for instance, the property *violently* is not being ascribed to the door, rather it is characterising the way in which the event occurs. In the case of subjectless events, such as *the thunderstorm*, causal efficacy can be attributed to the persistent and changing properties of spatiotemporal regions.

What does it mean, however, to say that objects and events are causes and effects *in virtue* of their property instances? There seem to be two possible ways of spelling out this thesis further:

(A) The fact that e_1 causes e_2 supervenes upon e_1 and e_2 's property instances.

(B) The fact that e_1 causes e_2 supervenes upon the causal relations taking place between the property instances of e_1 and e_2 .

Now I take it that, given what has been said, CTP is at least committed to (A). CTP claims that property instances determine an object's range of possible causal interactions, so if an object's causal powers change, that object's property instances must also change. Moreover, due to the nature of events and their connection to objects, we are led to the view that an event's causal efficacy depends upon its (and its subjects) property instances. Therefore, as supervenience is at least symptomatic of a dependence relation, we should claim that there could be no change in the fact that e_1 caused e_2 , without some change in the property instances of e_1 and/or e_2 . But there could be a change in the property instances of e_1 and/or e_2 , without this changing the fact that e_1 caused e_2 .

Thesis (B), however, is stronger than (A). For we might think that although event causation supervenes upon the property instances of events, nevertheless, property instances are not themselves causes and effects. So we cannot say that the fact that e_1 causes e_2 supervenes upon the property instance causation occurring between the property instances of e_1 and e_2 . It is the stronger thesis (B) that I intend to defend, however. For earlier I claimed that property instances are the (fundamental) causes and effects, and only supervenience thesis (B), not (A), states that property instances are causes and effects. CTP is not obviously committed to this stronger thesis, as it states that the property instances of objects and events bestow the causal power or efficacy onto these objects and events. So it looks like CTP theorist could endorse (A) but not (B).

There is, however, an extremely close link between (A) and (B). If we think that property instances are the fundamental causes and effects, we have an explanation of why (A) holds, since (B) implies (A). There couldn't be a change in the property instance causation occurring between e_1 and e_2 , without some change in the property

instances of e_1 and e_2 . But there could be a change in the property instance causation occurring between e_1 and e_2 without this altering the fact that e_1 caused e_2 . So there is no reason to suppose that a change in the property instances of e_1 and e_2 need change the fact that e_1 caused e_2 . If, on the other hand, we deny (B), it's not clear why anyone would accept that causation between objects and events supervenes upon the property instances of those relata. For if property instances are not themselves causally efficacious entities which can stand in the relation of cause and effect, it is unclear what could account for the fact that property instances have such a special position in the causal relation.

Even more importantly, however, I think that the stronger thesis (B) is implied by a generalised form of either supplemented functional realiser theory or functional role theory. For then we get the following two-level criterion of identity: two property instances/tropes are instantiations of the same property P iff they both realise the causal role specified by P's RCL definition. Implicit in this, is the claim that property instances can be causes, as they have to implement certain nexuses of causal relations. Facts about event and object causation can thus be viewed as supervening upon the causal relations between complex aggregates of particulars' property instances, as Hobbes suggests. So whilst I accept that it may be possible to develop a version of CTP which only endorses (A), as we have developed it, CTP is committed to (B). Moreover, (B) is a satisfying development of (A), as it provides us with a rationale for accepting it.

According to the thesis being proposed here, then, property instances are the fundamental causal relata - they are the entities which are doing the causal work. Although there is causation between other types of entities, namely the more coarse-grained objects and events, their ability to function as causes and effects is bestowed upon them by their property instances. What this means, is spelt out further by supervenience thesis (B). To illustrate (B), consider the causal statement 'dropping the sugar-cube in water caused it to dissolve'. The event of dropping the sugar-cube in water is a cause, but it so in virtue of its property instances. This means that there could only be a change in the fact that dropping this sugar-cube in water caused it to dissolve, if there had been some change in the underlying property instance causation. For instance, we can suppose that if the water had instantiated the property of being

saturated, the sugar-cube wouldn't have dissolved, since the property instance causation occurring between the water and the sugar-cube would have been different. However, a change in the underlying property instance causation needn't have made a difference to the fact that dropping this sugar-cube in water caused it to dissolve. If the sugar-cube had been placed in slightly hotter water, for instance, the dissolving might have happened quicker, but this wouldn't alter the fact that dropping the sugar-cube in water caused it to dissolve.

4.2 Davidson's Rival Ontology

The view that property instances are the (fundamental) relata of causation, obviously lays heavy emphasis on the role played by properties in causation. This emphasis is found in most other accounts of the causal relata.⁷ But there is an alternative viewpoint, put forward by Davidson, which doesn't allow that properties have any place in the causal ontology. So in this section, I shall try to motivate the claim that properties should be part of our causal ontology.

Davidson argues that events are the sole relata of causation. These events are *sui generis* occurrences or happenings, such as *Joel and Ann's wedding, the explosion, World War Two* etc. Like objects, events are concrete particulars. The event of Joel and Ann's wedding, for instance, has an unrepeatable location in space-time. Another feature shared by events and objects is that they are both multi-faceted or multi-propertied entities. They have, what Steward calls, "a secret life" (1997, p.35), as they can be redescribed in many different ways. 'Joel's and Ann's wedding', for instance, can be redescribed as 'Joel's coming to be a husband', or 'the wedding of Bob and Kate's daughter', and so on. The description of an event may be partial or unhelpful, but this doesn't matter so long as it serves to latch onto the right event.

Despite these similarities between objects and events, Davidson is keen to maintain the distinction between the two. He writes, "One is an object which remains the same object through changes, the other a change in an object or objects" (1985, p.176). Intuitively, events do not persist through time, like objects do. One way of capturing this difference is by saying that events (unless instantaneous) are not wholly

⁷ See, for instance, Armstrong (1997), Kim (1976), Mellor (1995) and Menzies (1989). Even Lewis (1986b) allows that we need to appeal to properties in order to specify the essences of events (see §5.7).

present at any one time. Rather they are spread out over time, like a carpet is spread out over different spatial points.⁸ On Davidson's view, then, it is only these *sui generis* events that can be causes and effects. Objects and, in particular, properties or property instances cannot stand in this relation. This makes Davidson's causal ontology coarse-grained, as his causes and effects can remain the same despite numerous changes to their properties.

Why does Davidson want to say that causal statements, such as 'the bridge's weakness caused its collapse', or 'Bill's anger caused Jessica's fear, or 'the brightness of the light caused the picture to fade', do not report causal relations between property instances? Although Davidson generally avoids denying the existence of anything, he frequently says that we don't need to posit such entities as properties or property instances to explain anything.⁹ Events and objects are all that are required to make sense of what we say and hold to be true. Now, clearly, every CTP theorist is going to have to oppose this kind of nominalism. But I will not engage with this position here, since my intention is to try to say why CTP is a good theory, *given* that the need for properties has been granted. What I do want to argue, however, is that if we endorse an ontology which includes property instances, then we have good reason to claim that they can partake in causal interactions. For their fine-grainedness enables us to distinguish between good and bad causal explanations. In what follows, I shall begin by outlining a well-known argument against Davidson's view. In the course of examining Davidson's response to this argument, I hope to show why accepting a causal ontology of property instances provides a better account of the distinction between good and bad causal explanations.

The charge against Davidson's view is basically this: his causal relata are not fine-grained enough to capture the subtleties involved in questions concerning what causes what. For there are occasions where we want to say that some particular aspect of the event, not just the event *per se*, was the cause of the effect in question. In order to make this objection more concrete consider this (we'll suppose true) causal

⁸ If you hold a perdurantist view of objects rather than an endurantist view, then objects won't persist through time in the way suggested above (see Lewis 1986a for a characterisation of these positions). But we can still say that our intuitive conception of an object, as opposed to an event, is of something that endures through change. So this characterisation still helps us to see the kind of entities Davidson has in mind when he talks of events.

⁹ See, for instance, Davidson 1977.

statement, 'Don's gripping *lightly* caused his fall'. If we think that this causal statement picks out a causal relation between two concrete events, it doesn't matter how we describe them so long as the same events are picked out. So we can assume that 'Don's gripping', although only a partial description of the cause, refers to the same event as 'Don's gripping lightly'. Now, as Davidson thinks we can substitute co-referring event descriptions without this affecting the truth-value of a causal statement,¹⁰ we should be able to substitute the event description, 'Don's gripping lightly', for the other, 'Don's gripping', without this affecting the truth-value of the resulting causal statement. So the causal statement 'Don's gripping caused his fall' should come out true. The problem should now be evident. It seems reasonable to maintain that the causal statement 'Don's gripping caused his fall' is false, despite the fact that 'Don's gripping lightly caused his fall' is true.

Davidson, of course, disagrees. He argues that we should accept that 'Don's gripping caused his fall' is a true causal statement. Any appearances to the contrary can be explained away by distinguishing between true causal statements and causal explanations. Commenting on a different case, he writes,

'The cause of this match's lighting is that it was struck. – Yes, but that was only part of the cause; it had to be a dry match, there had to be adequate oxygen in the atmosphere, it had to be struck hard enough, etc.' We ought now to appreciate that the 'Yes, but' comment does not have the force we thought. It cannot be that the striking of this match was only part of the cause, for this match was in fact dry, in adequate oxygen, and the striking was hard enough. What is partial in the sentence, 'The cause of this match's lighting is that it was struck' is the *description* of the cause; as we add to the description of the cause, we may approach the point where we can deduce, from the description and laws, that an effect of the kind described would follow (1967b, p.155-6).

Davidson's point here, is that if the striking of the match in actual fact took place in the right conditions (namely, in the presence of oxygen, when it was dry etc.), we can say that this event of striking was the complete cause of this match lighting. What leads us to talk of it being a mere 'part' of the cause, is considering what conditions generally have to pertain in order for the match to light on striking. But while all this information would provide us with a fuller causal explanation of the effect, these fuller

¹⁰ See Davidson 1967b. The claim that "events can be expressed...by an ordinary two-place predicate in an ordinary, extensional first order language" (p.161) is central to Davidson's position.

descriptions of the event do not add anything to the cause, as the cause is just that concrete happening – in all its complicated detail.

If we apply this response to the Don's gripping case, we see that things don't look that bad for Davidson. 'Don's gripping' is, in actual fact, 'Don's gripping lightly'. Although it's true that our description of the event doesn't mention the fact that it is a *light* gripping, and so this description does not count as a good explanation of the effect, nevertheless, the event the description picks out *is* a light gripping. Therefore, we still succeed in latching on to the right cause by employing the event description 'Don's gripping'. To think that we need our event description to mention the fact that it is a *light* gripping is, according to Davidson, to make a simple error. It is to suppose that "every deletion from the description of an event represents something deleted from the event described" (1969, p.157). Once we have realised that the cause-event doesn't alter just because our ways of describing it do, we will no longer be tempted to insist that it is 'Don's light gripping' not just 'Don's gripping' that is the cause, because both descriptions are referring to the same thing.

A number of philosophers have tried to undermine Davidson's account by simply dismissing his proposed distinction between causation and causal explanation. Mellor, for instance, writes,

Davidson needs more than a distinction between causation and causal explanation: he needs a dichotomy...in the sense in which a cause might be an explanation, explanations must be facts, not particulars; and so must what they explain. But this entails that if all causes had to be particulars, none could ever explain its effects – or anything else – because explanations, being facts, could not be causes. Similarly, if all effects had to be particulars, they could never be explained, by causes or by anything else, because only facts can be explained (1995, p.131).

This certainly seems a surprising and implausible result. Does Davidson's distinction between cause and causal explanation really lead to such a conclusion? No, as Mellor's claim that only facts can explain and be explained is highly questionable. 'Because', as Mellor rightly points out, is a sentential connective and so requires true sentences on each of its sides. So if we think that sentences stand for facts, as many do, then 'because' explanations cite facts rather than particulars. Even if we are prepared to grant this, however (which of course Davidson wouldn't, as he does not think that there are such things as facts to which sentences refer) we're still a far cry

from the claim that facts are the *only* thing that can stand in the explanatory relation. Take, for instance, the connective ‘...explains...’. In English, this is not a sentential connective, so it can link terms which, plausibly, stand for things other than facts. Statements such as ‘the explosion explains the mess’, or ‘the cold in winter explains the hedgehog’s hibernation’, have terms which are standardly taken to refer to events and properties. So there’s no reason to deny that entities other than facts can explain.

I don’t think that we will be able to undermine Davidson’s position by simply dismissing his proposed distinction between causation and causal explanation. Although Davidson does talk about distinguishing causal statements from causal explanations, it would be extremely uncharitable to interpret him as saying that true causal statements reporting causal relations between events aren’t causal explanations – of course they are. The claim is only that *some* causal explanations do not report relations between events. So there is no reason to suppose that causal statements reporting events do not offer causal explanations, as Mellor suggests. Davidson’s reply, therefore, and his stance more generally, seems perfectly consistent. I hope to show, however, that despite this, his account is problematic. For Davidson’s response to the fine-grainedness argument crucially appeals to this notion of a good/bad causal explanation. But his causal ontology of concrete events isn’t rich enough to properly capture this distinction.

We can all recognise that ‘Don’s gripping *lightly* caused his fall’ is an illuminating causal statement, unlike ‘Don’s gripping caused his fall’. But, according to Davidson, properties of an event make no special causal contribution to the event-cause. So these two statements are causally on a par, as they both accurately report the cause of the fall. In order to explain away appearances to the contrary, we’ve seen that Davidson’s distinction between good and bad causal explanations plays a pivotal role. For he argues that although ‘Don’s gripping caused his fall’ is a true causal statement, because it is a bad causal explanation we are inclined to think that it is false. So if this response is going to be persuasive, we require an account of why this is a bad causal explanation, unlike ‘Don’s gripping *lightly* caused his fall’.

Davidson’s analysis of what makes good and bad causal explanations is not entirely clear. In his 1967b paper, however, he seems to endorse a very permissive

form of the covering-law model of explanation.¹¹ For good causal explanations are said to be those where the causes and effects are “characterised in such a way that we can deduce, or otherwise infer, from laws or other causal lore, that the relation was causal” (1967b, p.155). So the idea is that some descriptions of the cause and effect will be such that it will be clear what law (or other causal generalisation) the events in question are subsumable under. In the Don’s gripping case, for instance, it seems very unlikely that there is going to be a law (strict or otherwise) under which grippings and fallings are subsumable. However, it does seem likely that there will be a general causal correlation between the strength of a gripping and the likelihood of a fall. So ‘Don’s gripping lightly’ turns out to be a good causal explanation, unlike ‘Don’s gripping’.

This account, then, can handle the gripping case, but there are other causal explanations which do not seem to fare so well. In particular, consider those causal statements which, on Davidson’s account, are non-canonical because they do not explicitly report relations between events. In order to meet Davidson’s criterion for a good causal explanation, the events cited in these explanations must be characterised in such a way so that we can subsume them under a law or causal lore of some kind. But non-canonical causal explanations, such as ‘the bridge’s weakness caused its collapse’, or ‘the brightness of the light caused the picture to fade’, do not explicitly refer to any events. So even if we are prepared to grant the contentious assumption that these explanations do implicitly refer to some events, we nevertheless have to admit that the causal explanations, as they stand, give us no indication of what those events might be. In the case of the collapsing bridge, for instance, it might have been a bolt’s snapping, or a piece of metal dissolving, etc. This means that Davidson cannot employ his account of good causal explanations. We cannot say that ‘the bridge’s weakness caused its collapse’ is a good causal explanation, for the inferred event(s) is

¹¹ For the classic statement of this view, see Hempel (1965). Davidson can’t just be proposing this, however, as he argues that “action explanations are a species of causal explanation” (1963, p.3), while asserting the anomalism of the mental (see, for instance, 1970 p.208). So causal explanations such as ‘I turned on the light because I wanted to illuminate the room’ are not good explanations through being subsumable under some law. We thus have to appeal to the fact that this explanation is subsumable under some “causal lore”. I think that the vagueness of this proposal is cause for some concern. What kind of “causal lore” makes these causal explanations good explanations? But I shall leave this class of causal statements aside here, for I think that there are other problematic causal explanations, which Davidson’s account needs to deal with.

not characterised at all by the causal explanation in question. So we cannot say that the event(s) has a characterisation which falls under a law or causal lore.

In cases where there isn't any obvious event cited, therefore, Davidson's analysis of what counts as a good causal explanation cannot be utilised. This undermines his response to the fine-grainedness problem, as that problem can be raised again for non-canonical causal statements. Take, for instance, the good causal explanation, 'the *brightness* of the light caused the picture to fade'. We can assume that the light had to be bright in order for the effect to occur. So, at the very least, we should say that 'the light caused the picture to fade' is a bad causal explanation. Davidson, however, hasn't got the resources to make this distinction in this kind of case. For we can't say that the causal explanation does or doesn't display which law/lore the causal relation is subsumed under, as the inferred event is not characterised at all by the causal explanation in question.

Davidson does still have some recourse in these kinds of cases, since he can appeal to the subjective aspects of explanation to account for the differences between these good and bad causal explanations. So he can say that 'the light caused the picture to fade' is a bad causal explanation because it doesn't provide people with the sort of information they want, it is scanty, misleading, and so on.¹² But this analysis of the situation seems deficient. Imagine, for instance, that the picture in question was painted with a special kind of paint, one which is resistant to fading under all kinds of light except extremely bright light. In this kind of case, it would indeed be misleading to say that the light was the cause of the faded picture, but the reason for this isn't primarily because it fails to meet the concerns of its listeners. For no matter what these concerns are, given the story just told, we can all recognise that 'the brightness of the light' is a better causal explanation *per se* of the picture's fading than 'the light'.

What, then, should be said about this and many other cases? It seems extremely intuitive to say that 'the brightness of the light caused the picture to fade' is a good causal explanation because it accurately reflects the objective causal structure of that situation. Whereas 'the light caused the picture to fade' is a bad causal

¹² Lewis takes this route. In 1986c, p.226-7, he provides a list which gives general guidance as to what makes a good explanation. He argues that an explanation is bad if the explanatory information is scanty, disorganised, unjustified from the perspective of the informer, adds little or nothing to the information the recipient already possesses, or fails to provide the recipient with the sort of information they want.

explanation because it fails to do this. But, given Davidson's sparse ontology, we cannot avail ourselves of this simple but plausible response. For there are no such things as property instances which figure in any way in the causal interactions that take place. So we are left having to appeal to the subjective aspects of explanation in order to explain why one is better than the other.

Does a causal ontology of property instances fare any better? Yes, their fine-grainedness gives us the resources to account for the differences between these explanations. If we allow property instances into our causal ontology, then we can say that the explanation, 'the *brightness* of the light caused the picture to fade' is a better causal explanation than 'the light caused the picture to fade' because, in this situation, it was the *brightness* of the light, not just the light, that was the cause. Similarly, in the Don's gripping case, we can say that the cause of the event is not Don's gripping, but rather the property instance of the event, namely its being a *light* gripping. Alternatively, we could agree, with Davidson, that 'the gripping caused the fall' is a true causal statement. Although it is a misleading description of the causal relation, it does pick out the right event and so the right cause. If we say this, however, we still have the advantage over Davidson, as we can give a better account of why it is such a misleading causal explanation. For although Don's gripping is the cause, on the account proposed here, it is a supervenient cause. Moreover, this characterisation of the supervenient cause fails to convey the kind of property instance causation it supervenes upon. For that the event instantiates the property of being a light gripping is something which the event's ability to function as the cause depends on. So if we omit this detail, the causal explanation will not accurately reflect the objective causal structure of the situation.

If we say that property instances figure in causal relations, therefore, we can capture these seemingly objective differences between good and bad causal explanations. Although context-relative features like how much the recipient already knows, what sort of information they want, what their capabilities for understanding are, etc. no doubt play a part in what a good causal explanation is, this isn't the whole story. What went wrong with Davidson's account? The problem doesn't stem from any overarching conception of what a good causal explanation is. Davidson can maintain, as I have done here, that a good causal explanation is one which, primarily,

accurately reports the objective causal structure out there in the world. The problem rather is caused by trying to combine this plausible, albeit sketchy, conception of a good causal explanation, with a coarse-grained ontology of events. Given these causal relata, much of the information we would normally take ourselves to have about the causal workings of a situation, cannot be said to be part of its objective structure. For so much of that information concerns the part that properties play in these causal interactions (it was the *hot* stove that burned her, it was the *weight* of the car that crushed him, it was the *violence* of his outburst that surprised her...). Consequently, if we deny that these entities can figure in causal interactions, a lot of the information we would standardly take to be about the objective causal structure of those situations, has to be put down to the subjective aspects of explanation.

Once we've allowed that property instances figure in causal interactions, this implausible result is avoided. We can say that causal explanations, such as 'the *brightness* of the light caused the picture to fade', report the objective causal structure of those situations. For property instances, like the *brightness* of the light, have causal influence and thus a bearing on the causal relations which take place. Due to the fine-grainedness of these relata, causal interactions will have a more complex structure than they will given Davidson's ontology. So the desire for comprehensive causal explanations will need to be tempered by the subjective requirements on a good explanation, as too much detail will often prove detrimental and unhelpful. But as a general rule of thumb, we can say that a good causal explanation is one which describes its cause in such a way so that everything that is highly pertinent to the occurrence of the effect is mentioned. So 'the brightness of the light caused the picture to fade' is a better causal explanation than 'the light caused the picture to fade', because the latter fails to mention a feature of the cause which was essential to the occurrence of the effect, namely, that it was a *bright* light.

By adopting a fine-grained account of the causal relata, therefore, we are able to capture the subtleties involved in questions concerning what causes what. This results in a better account of causal explanation, because good/bad causal explanations can be said to be tracking features of reality, rather than just subjective concerns about what makes an explanation appropriate for the particular recipient in question. So if we are prepared to endorse an ontology of property instances, there is every reason to

claim that these entities can engage in causal interactions. At least granted, that is, this claim isn't subject to any powerful objections. It is to this issue we now turn.

4.3 The Slingshot

If we give up Davidson's austere ontology of objects and events, a number of options confront the causal theorists. Probably the most popular of these, however, is the thesis that facts are the relata of causation. Facts go by a variety of different names, and different types of entities bear this title. On the Fregean view, for instance, facts are true propositions, which are necessarily existent entities outside space and time. When philosophers talk about facts as the relata of causation, however, they tend to have something more akin to Russell's view in mind. He states that facts are "the sorts of things that you express by a sentence, and...these, just as much as particular chairs and tables, are part of the real world" (1918, p.182-3). Facts are thought of as 'complexes' of objects and properties. The fact corresponding to 'the jumper is red', for instance, contains the jumper and its redness as components. Despite some differences in detail, Mellor, Bennett, Armstrong and Menzies all argue that these Russellian facts are the relata of causation.¹³

Whatever our conception of facts may be, however, Davidson thinks he has a powerful argument against them, namely, the slingshot. This purports to show that all facts collapse into one, for there can only be one fact which all true sentences stand for. There are a number of different versions of the slingshot. Davidson's slingshot depends upon the contentious principle of substitutivity for logical equivalents.¹⁴ However, Neale's version of Gödel's slingshot requires less controversial assumptions. His argument is rather complex, but I shall try to faithfully represent it here.¹⁵ It relies on two basic rules of inference:

(I) *The Principle of Substitutivity for Definite Descriptions (t-SUBS)*: "if a predicate F applies to exactly one object (i.e. if it has exactly one thing in its extension), in

¹³ See Mellor (1995, 2003), Bennett (1998), Menzies (1989) and Armstrong (1997), although Menzies calls them "situations" and Armstrong names them "states of affairs". Another conception of facts is put forward by Austin (1950). He thinks that facts are entities in the spatiotemporal world, but he does not think that they are structured entities, as Russell does.

¹⁴ See Davidson 1967b, p.152-3.

¹⁵ The following is taken from Neale 2001, chapter 9.

truth-functional contexts the description ιxFx can be treated as if it were a singular term for derivational purposes” (Neale, 2001 p.159).

(II) *Iota-Conversion* (ι -CONV): From Fa we can derive that a is identical to the x such that x is identical to a and x is F , and vice versa.

The argument works by supposing that a particular sentence connective supports both these rules of inference. In the version of the argument I shall outline, the sentence connective shall be ‘the fact that Bill is angry causes it to be the case that (...)’. The argument can then be stated as follows:

1. ‘Jessica is scared’. (Premise)
2. ‘Jessica is not identical to Bob’. (Premise)
3. ‘Bob is chubby’. (Premise)
4. By ι -CONV, ‘Jessica is identical to the x such that x is identical to Jessica and x is scared’. (From 1)
5. By ι -CONV, ‘Jessica is identical to the x such that x is identical to Jessica and x is not identical to Bob’. (From 2)
6. By ι -CONV, ‘Bob is identical to the x such that x is identical to Bob and x is not identical to Jessica’. (From 2)
7. By ι -CONV, ‘Bob is identical to the x such that x is identical to Bob and x is chubby’. (From 3)
8. By ι -SUBS, ‘the x such that x is identical to Jessica and x is scared is identical to the x such that x is identical to Jessica and x is not identical to Bob’. (From 4, 5)
9. By ι -SUBS, ‘the x such that x is identical to Bob and x is chubby is identical to the x such that x is identical to Bob and x is not identical to Jessica’. (From 6, 7)
10. ‘The fact that Bill is angry causes it to be the case that Jessica is scared’. (Premise)
11. By ι -CONV, ‘The fact that Bill is angry causes it to be the case that Jessica is identical to the x such that x is identical to Jessica and x is scared’. (From 10)
12. By ι -SUBS, ‘The fact that Bill is angry causes it be to the case that Jessica is identical to the x such that x is identical to Jessica and x is not identical to Bob’. (From 11, 8)

13. By ι -CONV, ‘The fact that Bill is angry causes it to be the case that Jessica is not identical to Bob’. (From 12)
14. By ι -CONV, ‘The fact that Bill is angry causes it to be the case that Bob is identical to the x such that x is identical to Bob and x is not identical to Jessica’. (From 13)
15. By ι -SUBS, ‘The fact that Bill is angry causes it to be the case that Bob is identical to the x such that x is identical to Bob and x is chubby’. (From 14, 9)
16. By ι -CONV, ‘The fact that Bill is angry causes it to be the case that Bob is chubby’. (From 15)

The moral of all this is that if we allow that our sentence connective, in this case ‘the fact that Bill is angry caused it to be the case that (...)’, licenses the ι -SUBS and ι -CONV rules of inference, then we can prove that any true sentence can be placed in (...).

This argument, however, as Neale rightly points out, is not decisive against facts. For it can be avoided by endorsing Russell’s theory of descriptions.¹⁶ Why does this help? A sentence connective such as, ‘the fact that Bill is angry caused it to be the case that (...)’, will not support the ι -SUBS and ι -CONV rules of inferences, given Russell’s theory. Why? Two definite descriptions of the same object do not usually contribute the same descriptive properties to a fact. For instance, ‘the prime minister of Britain’ and ‘the husband of Cherie Blair’ while referring to the same object, do not stand for the same fact. For different properties are involved in these complexes and, given a Russellian view of facts, the properties they introduce are components of these facts. Consequently, we cannot accept the ι -SUBS rule of inference. The ι -CONV rule is also undermined, for definite descriptions receive a quantificational analysis, so they introduce properties not present before the iota-conversion. In the sentence ‘Jessica is scared, for instance, no relation of identity is to be found. So the ι -CONV rule again changes the nature of the fact being referred to.

Fact theorists, therefore, have a way of escaping even this powerful version of the slingshot. But answering the argument comes at a cost, since it means taking a

¹⁶ Davidson’s slingshot can also be circumvented in this way.

stand on the semantics of definite descriptions. Is this price also incurred by the view that property instances are the relata of causation? It may be thought so, for some have argued that there isn't a significant difference between property instances and facts.¹⁷ Hence, as the slingshot applies to facts, property instances won't remain unscathed. I think that this view is mistaken, however. The thesis that property instances are the relata of causation needs to be firmly distinguished from the claim that facts are the causal relata. Why? Central to all characterisations of facts is this claim that facts correspond, are the referents of or stand for true sentences. Mellor, for instance, says that facts "correspond to true statements" (1995, p.8), Menzies claims that they are "worldly correlates of true sentences" (1989, p.67) and Armstrong writes, "there is something in the world that corresponds to a true proposition", namely, a state of affairs (1997, p.128). But those who just assert the existence of property instances do not commit themselves to any such thesis. The sentence 'Bob is chubby' does not correspond or refer to a property instance. Although property instances, such as Bob's chubbiness, will be involved in the truth of such sentences, this is very different from saying that property instances are the entities which sentences stand for.

It is this latter claim, however, which is crucial to the slingshot against facts. For these arguments attempt to show that the fact referred to by one true sentence is the same fact as that referred to by another true sentence. The version of the slingshot just outlined, for instance, works by supposing that the ι -SUBS and ι -CONV inference rules apply to the sentence connective 'the fact that Bill is angry causes it to be the case that (...)'. Given this, we can then go on to prove that any true sentence whatsoever can fill the blanks. So, granted the assumption that all these sentences correspond to some fact, we get the conclusion that all true sentences stand for the same fact. If we endorse a causal ontology of property instances instead, however, then the argument cannot be formulated. For there are no property instances corresponding to sentences. Consequently, we can't prove that all property instances collapse into one.

Despite the fine-grainedness of both property instances and facts then, there is an important difference between a causal ontology of property instances and a causal ontology of facts. Earlier I urged that property instances should be restricted to those

¹⁷ See, for instance, Steward 1997, chapters 1 and 5.

entities which ground genuine resemblances in the world, they should not be thought of as fulfilling a semantic role.¹⁸ So property instances appear out of the slingshot's firing range, as they shouldn't be viewed as entities which correspond to or are the referents of true sentences.

There is, however, a more indirect way of trying to wield the slingshot against properties and hence property instances. Davidson argues that we are only committed to those kinds of entities we need to posit in order to make sense of what we say and hold to be true.¹⁹ If this Quinean view is accepted, then an anti-nominalist has to find contexts in which more than the extensions of the terms (i.e. the objects and events) contribute to the semantic value of the sentence.²⁰ So in order to show that something mentioned in a sentence is a property, we need to connect it with the sentence in some non-extensional way. Earlier, for instance, I suggested that 'Don's gripping *lightly*' contributes to the semantic value of the sentence, because we can suppose that while 'Don's gripping *lightly*' and 'Don's gripping' have the same extension, substituting one for the other can turn this true sentence into a false one. So this seems to count against Davidson's nominalism, since it looks like the former sentence isn't just referring to the event of Don's gripping.

This, however, is where Davidson's slingshot comes in. For if we want to say that we are talking about a property or property instance, we need a non-extensional context, i.e. one where more than the extensions of a sentence's terms contribute to its semantic value. But, granted we still want to allow substitution of co-valued terms for the relevant kinds of entity, Davidson can then wield his slingshot against us. So the idea is that in order to preserve valid inferences, we need to allow for the substitution of co-valued terms. For instance, if Don's gripping lightly caused his fall, and Don is Quine's best student, then we want to say that Quine's best student's gripping lightly caused his fall. But, then, our (allegedly) non-extensional context, by use of the slingshot, becomes provably extensional.

Davidson's slingshot, combined with his wider philosophical apparatus, therefore, does pose a challenge to those who wish to posit an ontology richer than

¹⁸ See §2.2.

¹⁹ See, for instance, Davidson 1977.

²⁰ Unless, that is, we can find some sentences which involve reference or quantification over entities other than objects and events, and which can't be paraphrased.

objects and events. This takes us into the battle against nominalism which, as I said earlier, I cannot engage in here. But there are two ways we could challenge the conclusions Davidson reaches. First, we could try to argue against the background philosophical theory it presupposes. So we could reject the claim that we are only committed to those kinds of entities we need to make sense of what we say. Instead we could insist that if there are good metaphysical reasons to posit property instances, then we should allow these entities into our ontology. Second, if we do want to endorse the Quinean view (and no sentences which a nominalist can't paraphrase can be found), then we could reject the slingshot by employing the aforementioned response.

Since there is good reason for wanting to block the slingshot anyway, we should endorse the second of these responses. Why? Earlier I argued that causal explanations should be taken to reflect the causal structure of the world and to be made true or false on the basis of this. If the slingshot works, however, then causal explanatory contexts have to be extensional. So if Don's fall is the first because he has the weakest rope,²¹ and Don's rope is the weakest rope, then, given the slingshot, Don's fall is first because he has Don's rope. Thus, although the slingshot is not a deadly objection against a causal ontology of property instances, given the wider perspective we still need some way of dodging it. But as a plausible semantic theory is armour enough, there is no cause for concern here.²²

4.4 Omissions

'The void caused Billy's blood to boil', 'Hamlet's failure to act led to his downfall' 'Alice's precautions prevented the explosion from ensuing', 'the absence of rain caused the food shortage'. All these causal statements seem to report either an absence causing something, or something causing an absence, or even an absence causing an absence. So, if we take these statements at face value, we have to say that absences can be causes and effects, as it is there *not* being any rain that causes the food shortage, or there *not* being any forces which causes Billy's blood to boil. These

²¹ See Mellor 1995, p.116.

²² See Neale 1990 and 2001. There, Neale makes a persuasive case for Russell's theory of definite descriptions.

causal statements pose a problem for the view that the relata of causation are property instances. Because the property instances in question are those entities which are postulated by CTP. So they have to be the sort of entities which contribute to the causal powers of objects and efficacy of events. Given this analysis, it seems implausible to posit negative property instances. For we don't want to say that it is an entity's not being 100°C which bestows a certain set of causal powers onto a particular. Rather it is the fact that it instantiates a positive property instance, say of being 10°C, which gives it those powers.

The frequency and importance of causal statements involving omissions indicates that absences should be included in our causal relata. But an ontology of negative objects or negative events looks as equally unpalatable as that of negative property instances.²³ So what could the relata of causation be when causal statements cite omissions? One option is to simply reify absences and say that there are such things as the absence of rain or the failure to act, which can stand in the relation of cause and effect. But few have been tempted by this. Another option is to claim that facts are the relata of causation in cases involving omissions. For facts are defined as those entities which correspond to true sentences. So sentences like 'the ball is not red' or 'this chicken is not 10 kg' will stand for a fact, no less than the sentence 'this chicken is 2 kg'.

Despite a number of fact theorists using omissions to argue that their ontology has the edge over its competitors,²⁴ it is difficult to avoid the feeling that this solution is just another way of reifying nothings. If we claim that the sentence, 'the ball is not red' stands for a Russellian fact, for instance, then we still have to say that there are such things as the property of not being red. Allowing these entities into our ontology isn't only unparsimonious, it also seems unmotivated beyond our concerns here. For we can say that what makes a sentence like 'the ball is not red' true, isn't a complex of an object and a negative property, rather it is the complex of an object and a positive property, such as its being blue. This approach is sensibly endorsed by a number of fact theorists. Armstrong, for instance, denies the existence of negative states of

²³ Davidson's account of adverbial modification commits him to denying that there are such things as negative events (see Davidson 1967a).

²⁴ See Menzies (1989, 2003) and Steward (1997 §4.1). Mellor (1995) also initially presents this as an argument in favour of his account over Davidson's, but later he backtracks.

affairs.²⁵ While Mellor argues that although our definition of a fact guarantees the existence of an entity which stands for a true negative existential statement, this doesn't mean that what makes the negative existential statement true, is this negative fact.²⁶

The problem of omissions or, as Lewis calls it, “the problem of the missing relatum”²⁷ is, therefore, a worry for anyone who wishes to avoid reifying negative entities, such as the absence of rain, or his failure to act. This diffuses the objection against a causal ontology of property instances somewhat. For if there are no negative objects, no negative events and, plausibly, no negative facts which can do the causing, then at least the claim that there are no negative property instances puts the view on a par with its competitors with respect to this issue. Nevertheless, there is still a question that needs facing here, for if we don't want to reify these absences, what should we say about this seeming causation by omissions?

Broadly speaking, there are two options open to us. First, we could keep hold of the commonsense claim that there is causation by absences, while avoiding any reification of them. All this comes at a price, however. For “a relation requires relata”,²⁸ and we haven't got these relata if either the cause or the effect is an absence. So we are forced to deny that causation is a relation.²⁹ Second, we could maintain that causation is a relation and instead sacrifice, or at least qualify, the claim that absences can be causes and effects. Noordhof, for instance, argues that there may be “various positive events each of which – given the laws which hold – would make the negative causal statement true and one of which, in fact, did make it true in the circumstances” (1998, p.858). If such positive instances of causation could always be found, then the truth of causal statements involving absences could be preserved. But this shouldn't obscure the fact that there is no real causation by absences on this view.³⁰

²⁵ See Armstrong 1997, §8.7. There he argues that as well as postulating positive states of affairs to be the truthmakers of negative sentences, we also have to postulate higher-order states of affairs, whose content is that such-and-such are all the entities of some selected sort. This is because states of affairs are meant to guarantee the truth of certain sentences. So if those states of affairs exist at a world, then that truth has to hold at that world.

²⁶ See Mellor, 1995 p.157.

²⁷ See Lewis 2003b.

²⁸ Lewis 2003b, p.6.

²⁹ See, for instance, Mellor 1995, chapter 13 and Lewis 2003b.

³⁰ Dowe (2000) and Beebe (2003) both defend this kind of view. They follow Noordhof in maintaining that there is no causation by absences. However, they do not take on the burden of claiming that causal

Given what has been said so far, both of these options are open to someone who claims that property instances are the relata of causation. In the next chapter, however, I shall argue for a thesis which is committed to the claim that causation is a relation. So anyone wishing to accept the conclusions there, will have to say that there is no genuine causation by absences. Due to the limits of space, I won't try to defend this thesis here. But, fortunately, much has already been done to show that it is defensible.³¹ With this commitment noted, the wider issue concerning the right analysis of causal statements involving omissions will have to put to one side. But a very limited conclusion can be drawn from the discussion here. For we have seen that a causal ontology of property instances doesn't place us in a worse position with regards to omissions than other accounts of the causal relata. Therefore, this isn't a strong objection against the claim that property instances are the (fundamental) causal relata.

statements involving absences are always made true by positive causal relations in this world. Instead, they claim that these causal statements either report relations of 'causation*' (Dowe), or provide us with some other information about the causal history of an event (Beebe). In light of the pressure put on Noordhof's position by examples such as Lewis's void (see Lewis 2003b), someone who opted for this second approach might find themselves pressured into a position like Dowe's and Beebe's. But the jury is still out on this.

³¹ See especially Dowe (2000). He argues that causation by omission or causation* is "a hybrid fact usually involving certain actual real causation together with certain counterfactual truths about real causation" (2000, p.124). So causal statements about omissions, while not straightforwardly reporting relations of causation, do involve facts about the actual causation that occurred. Moreover, they tell us something about the causation in nearby possible worlds, where what is absent in this world occurs.

5. Singularism

In this chapter, the focus will turn away from the question of what sort of entities causes and effects are, to the issue of what it is for one entity to cause another. I shall argue that two assumptions should be incorporated into an account of the causal relation. The first of these concerns the local, intrinsic nature of the causal relation, whilst the second draws our attention to the fact that causal relations are part of more general patterns. I shall argue that CTP, if developed in the right way, gives us a plausible rapprochement of these two theses. This, I hope to show, will further two aims: first, it will illustrate that CTP has positive consequences for an account of the causal relation. Second, it will further support the thesis that property instances are the fundamental causal relata. Although, as we'll see, the account will place more demands upon the nature of these property instances.

5.1 The Intrinsicity Assumption

Causation is a *local* feature of a cause-effect pair. What makes one thing a cause of another is entirely a matter of the nature of the cause, of the effect, and of the transaction between them (Bigelow and Pargetter 1990, p.271).

Here, Bigelow and Pargetter express what I shall call the intrinsicity assumption. This states that two entities are causally related in virtue of local facts about that relation. This idea, whilst somewhat overlooked due to the influence of Hume, has enjoyed support from the singularist tradition.¹ Ducasse, for instance, argues that “the correct definition of the causal relation is to be framed in terms of one single sequence” (1926, p.124). This idea is also echoed in Anscombe’s analysis, as she claims that “causality consists of the derivativeness of an effect from its causes...If A comes from B, this does not imply that every A-like thing comes from some B-like thing” (1971, p.91-2). Claims like these have lead Menzies’s to identify singularism with the following thesis:

¹ As will become evident, there are different types of singularism, but I shall not try to draw a common essence from these accounts. While there is a family resemblance between the views, I shall just outline different theories which have been classified as singularist, when they become relevant to the discussion.

when a singular causal claim is true, it is made true by the holding of a certain kind of intrinsic relation...the uncontested point for all of them [i.e. the singularists] is that the truthmaker – the ontological ground – of a singular causal statement is a local, intrinsic tie between events (1999, p.315 brackets added).

I think that this identification is contentious. We'll see that many singularist accounts focus on charting the relationship between singular causal relations and causal laws. Nevertheless, I do agree that this claim is or should be central to singularist theories of causation, since I think that it constitutes the underlying appeal of the approach. Before explaining why this is, however, I shall first try to make the content of the intrinsicity assumption a little more perspicuous.

The intrinsicity assumption can be broken down into three claims: (i) Causation is a relation. (ii) The causal relation is an intrinsic relation. (iii) The obtaining of the intrinsic relation between a particular cause and effect is what makes true the relevant causal statement that C causes E. The first of these claims is relatively straightforward. It just captures the thought, emphasised by both Davidson (1967b) and Strawson (1985), that causation is a natural, objective relation which holds between things in nature. Due to the problem of the missing relatum, this once uncontroversial thesis has been subject to increasing scrutiny. But, as I mentioned in §4.4, here I shall assume that this thesis is true.

The second claim is a little more difficult to grasp. What does it mean to say that causation is an intrinsic relation? We've seen that this notion of intrinsicity is a tricky one.² But an intuitive gloss, similar to that given for intrinsic properties, conveys the basic idea. An intrinsic relation is one which can exist unaccompanied by all entities distinct from the relation. The relation of having the same mother, for instance, is not an intrinsic one. For whether or not two people stand in this relation depends upon more than the existence of the two people in the relation, it also depends upon the existence of a mother.

As it stands, however, this suggestion is unclear, as it could imply at least one of two things. First, we could take a relation to be intrinsic if it depended solely upon the nature of its relata, taken independently. So by fixing all the monadic properties of two particulars, A and B, all those relations that are thereby established are intrinsic to

² See §3.1.

A and B. For instance, if A is six foot and B is five foot, the relation *A is taller than B* will be intrinsic, because the monadic properties of A and B guarantee that this relation will hold between them. Lewis calls this kind of relation “intrinsic to its relata” (1983b, p.26). But it should be clear that this sense of an intrinsic relation is not the one relevant to the intrinsicity assumption. For causation depends upon other factors external to the monadic properties of its relata, such as the spatiotemporal relations between the cause and effect.

More appropriate here, is what Lewis calls a relation which is “intrinsic to its pairs” (1983b, p.26). A relation is intrinsic to its pairs if, roughly, it depends solely upon the relata and those relations which are generated by, and only require, the existence of those two relata. So the idea is that as well as those properties instantiated by A and B’s relata, a relation which is intrinsic to its pairs also includes all those relations which could hold between A and B, when A and B are unaccompanied by all other entities. Lewis tries to spell out this idea further by employing his duplication analysis. He writes,

Call the relation *intrinsic to its pairs* iff, whenever the pairs $\langle a, b \rangle$ and $\langle a', b' \rangle$ themselves are duplicates, then both or neither of them stand in the relation (1983b, p.26).

This elucidation is not very helpful, however. For earlier we saw that duplicates of a and b are just those entities with exactly the same perfectly natural properties as a and b. But then duplicates a' and b' are not guaranteed to have the same relations holding between them as a and b. For the fact that they have all the same perfectly natural properties doesn’t mean that the same spatiotemporal relations, for instance, will hold between them. So this doesn’t seem to be any advance on the intuitive characterisation given above.

Another reason for rejecting this duplication analysis is that, as in the case of the grounding intuition, it latches onto the wrong sort of intrinsicity. The intuitive force driving the intrinsicity assumption is this idea that causation is an entirely local affair. When one thing causes another, the thought is, we don’t need to go beyond that spatiotemporal region in order to find what is responsible for this fact. So what’s crucial to the intrinsicity thesis, isn’t the class of entities which can be instantiated

by duplicate pairs of a and b. Rather, the sense of intrinsicity we're after is that of the aforementioned interior entity.³ For it is this which captures the idea of a relation which is determined solely by the entities in and between the relata. A relation is intrinsic, therefore, in the sense relevant here, iff it does not depend upon anything extraneous to A and B and their juxtaposition.

The third thesis is the most contentious of the three. It claims that the obtaining of the intrinsic relation between a particular cause and effect is what makes true the causal statement that 'C causes E'. Another way of putting this is to say that the intrinsic causal relation between the cause and the effect is the sole truthmaker of the causal statement 'C causes E'. This formulation needs to be treated with caution, as talk of truthmakers is often associated with theories of facts.⁴ But in the present context, the term 'truthmaker' is only intended to convey the weaker idea that these intrinsic relations are responsible for the truth of causal statements. This idea can be expanded upon a little, however, if we are prepared to endorse Lewis and Bigelow's conception of truthmaking outlined earlier.⁵ If we think that "truth is supervenient upon being",⁶ then we can say that the truthmakers of a singular causal statement consist in the intrinsic property instances of the relata and relations that exist between them. So there could be no change in the truth of the statement 'C causes E', without some change in the property instances (relational and non-relational) instantiated by C and E.

According to the intrinsicity assumption, then, the causal relation is an intrinsic relation, and it is facts about this intrinsic relation which determine that one entity is the cause of another. In order to accord with this assumption, therefore, the analysis of what it is for one entity, C, to cause another entity, E, must appeal exclusively to local facts about the properties of the relata, C and E, and to facts about the intrinsic relations that obtain between C and E. Why should we accept this thesis? I think that it primarily recommends itself because it is strikingly intuitive. Whether or

³ See §3.1.

⁴ Armstrong, for instance, argues that facts or states of affairs are the truthmakers of sentences. These truthmakers have to guarantee (in all possible worlds) the truth of the corresponding sentence (see Armstrong 1997, p.115). Many have taken issue with this claim (see, for instance, Davidson 1967b, 1977, Lewis 2001 and Dodd 2002).

⁵ See §3.1.

⁶ See §3.1. The quote is from Lewis 1992, p.207. Also see Bigelow 1988 p.133 for the same conception.

not two entities are causally related looks as if it is a purely local matter, which depends solely upon what happens between those two things. The enjoyment caused by the chocolate Alice is now eating, for instance, doesn't seem to depend upon anything extraneous to the relation between Alice's eating chocolate and her enjoyment. In particular, it does not appear to rely upon past or future events of a similar sort, nor upon events occurring at different places in the universe. It seems only to concern what is going on at the time and place of Alice's eating the chocolate.

It is not difficult to see why this thesis appears so credible. If the truth of the causal statement 'C causes E' supervenes upon things extraneous to that causal relation, such as other C-types being constantly conjoined with other E-types, then it looks like C's ability to cause E isn't wholly within C's control. For it is plausible to claim that C cannot affect entities miles away from it, in the distant past or future. Consequently, if these sorts of extraneous facts are required for the truth of the claim that 'C causes E', then C and local facts about the relation between C and E cannot be fully responsible for the truth of that causal statement.⁷ When we think of mental causation statements, the implications of denying the intrinsicity assumption seem quite disturbing. If my deciding to sit down and write this chapter is the cause of my present typing in virtue of some facts which are extraneous to my decision and my sitting here typing, then it looks as if my sitting here typing is not something which I am in control of, as I am not in any position to bring about those extraneous facts necessary for my sitting here typing. There is, then, an explanation of why the intrinsicity assumption seems so appealing, since such a thesis seems to cohere better with our image of ourselves as the initiators of our own actions.⁸ But just

⁷ Obviously, there may be some things about the circumstances required to bring about E that are not within C's control but, in such cases, we tend to think that C is only part of the cause, or we pick it out as 'the cause' because it is the entity which is most pertinent to the occurrence of the effect.

⁸ The comments here are clearly very sketchy and are only meant to allude to why we might think that this thesis is attractive. Much more needs to be said in general about the ramifications of various positions within this area for an account of mental causation and free will. Recently, Beebe and Mele (2002) have contributed to this debate by arguing that Humean accounts of laws are an aid to free will. For what we choose to do determines what happens at a world, and hence what (Humean) laws hold at that world. Their interesting paper requires proper investigation, but I don't think that it affects this worry against the Humean approach to *causation*. For even if we grant that our actions can determine what laws there are, it still needs to be shown that, given the Humean approach, we are responsible for the things we think we bring about. It is this which seems threatened if we accept the standard Humean view that in a causal relation, entities in the far past or future are the truthmakers of the fact that C causes E.

because we can explain from whence the intuitiveness of the thesis arises, doesn't mean that it is false. On the contrary, I would suggest that unless there are good reasons to reject it, this should count in its favour.

If we look elsewhere in philosophy, moreover, we see that a permutation of this intuition is accepted by a great many philosophers. The thesis I have in mind is what I have been calling the grounding intuition. According to this, the causal powers of a particular are grounded in or determined by the intrinsic property instances of that particular.⁹ The intrinsicity assumption, on the other hand, states that the truthmakers of statements about causal relations are to be found among the intrinsic property instances of the relata and between the intrinsic relations that exist between them. What is the connection between singular causal relations and causal powers? The two are clearly different, because we think of causal powers as entities which an object can possess continuously and they tend to involve only one object. Causal relations, by contrast, implicate more than one object and they do not normally persist throughout an object's lifetime. Nevertheless, there is still an extremely close relationship between the two, because when a causal power is manifested, there will be a causal relation between the particular that instantiates this causal power and another entity.¹⁰ Consider, for instance, the causal power I have to move my hand away when it is being damaged. Suppose that, due to my careless cooking techniques, the right circumstances for a display of this causal power arises. Now we have a causal relation between event C, the event of my hand being burnt, and event E, me moving my hand away from the stove. The truthmakers for this causal relation, according to the intrinsicity assumption, are to be found in the intrinsic nature of the events of *my hand getting burnt* and *my hand moving away from the stove*. Primarily, then, we are looking at the intrinsic properties of the stove and me at certain times. If the

⁹ For examples of philosophers who accept this claim see, for instance, Mackie (1973), Harré and Madden (1975), Prior (1985), Armstrong (1996), Martin (1996), Jackson (1998), Mumford (1998), Cartwright (1999) and Campbell (2002a).

¹⁰ There may be a few counterexamples to this. Cartwright, for example, endorses the possibility of immanent causation (causation which remains within the entity but does not proceed via interactions between its parts), offering spontaneous radioactive decay as a possible instance of immanent causation (1989, p.109). If this is such a case, then displays of causal powers may fail to result in a causal relation between different entities. These cases are controversial, however, and although I do not want to rule them out, it seems clear that displays of causal powers do usually result in causal relations between distinct particulars.

intrinsicity assumption is correct, these entities, plus certain intrinsic relations between them, are going to have to account for the fact that there exists a causal relation between them.

This is exactly what we should anticipate, if the grounding intuition is true. According to the grounding intuition, the causal power that I have to move my hand away when it is being damaged, is something which is grounded in or determined by my intrinsic property instances. Consequently, when this causal power is manifested, due to the presence of the right conditions (in this case, the brush with the hot stove), what are, in large part, going to account for the fact that there is a causal relation between me and this hot stove, are facts about my intrinsic property instances which have determined that I will act in this way in certain circumstances. More generally, we can say that when a causal relation is the result of a display of an object's causal powers, the grounding intuition claims that what determines the behaviour of those objects are its intrinsic properties. Thus, apart from those trigger factors necessary for the display of the causal power, the truthmaker of the fact that this is a case of causation will be the relata's intrinsic properties.¹¹

The intrinsicity assumption, therefore, can be seen as stemming from the desire to respect the grounding intuition in an analysis of the causal relation. Those who are moved by the claim that the behaviour of a particular is determined by what lies within the confines of that particular, should also be moved by the claim that the truthmakers of causal statements are local facts about the intrinsic relation between the cause and effect. For an instance of causation involving two objects will involve displays of those object's causal powers, liabilities, tendencies etc. So if the truthmakers of statements about these powers are entities which are intrinsic to the object, these entities will also have to play a key role in determining what it is for C to cause E. The close connection between these two theses attests to the plausibility of the intrinsicity assumption. Most philosophers have wanted to hold on to the intuition that there must be something about the object which accounts for the fact that

¹¹ Earlier, in §3.1, I outlined two versions of the grounding relation, a weak and a strong. If the weaker relation is embraced, then as well as saying that the intrinsic properties of objects and their relations determine the truth of causal power ascriptions, we can also appeal to the laws of nature. But, as I said then, this appeal to the laws of nature, at least given most accounts, undermines the intuitive force of the grounding intuition. So I'll put this weaker version aside here.

its dispositional ascriptions hold true of it. The matter is very different, however, when we come to the debate about the causal relation. Due to the strong influence that Hume exerts in this area, this intuition has often been overlooked. But those who are attracted to the grounding intuition shouldn't disregard the intrinsicity assumption, since it is just a variation of it.

Are there any reasons to think that it would be beneficial to impose the intrinsicity assumption onto an analysis of the causal relation? One is suggested by the recalcitrant pairing problems that afflict some analyses of causation. The pairing problem appears in different forms, but the basic problem is always the same: the proposed analysis isn't able to pair the right cause with the right effect. Take, for instance, the simple counterfactual analysis. This states that C causes E iff, if C hadn't have occurred, E wouldn't have occurred either.¹² This account is vulnerable to the problem of preempting or back-up causes, as it will only be true that E won't occur in the absence of C, if there isn't some other suitably placed entity, C*, capable of bringing about the effect if C fails. In this sort of scenario, the simple counterfactual theory is unable to pair up the right cause with the effect, since C will no longer count as a cause according to this criteria.

Although more and more sophisticated counterfactual analyses have been developed which try to avoid this problem, it has been possible to formulate more and more sophisticated thought experiments in order to combat them.¹³ The mass of attempts to salvage the counterfactual analysis make it reasonable to surmise that this is a waning research programme, a fresh approach to analysing causation is required. It is instructive, however, to consider where the counterfactual analysis goes wrong, so we can then hopefully avoid these problems in future analyses. Part of the account's inadequacy, I think, stems from its failure to respect the intrinsicity assumption. On the counterfactual analysis, extraneous factors concerning what other entities are around at the time of the causal relation affect the truth of the statement 'C causes E'. This creates problems for the account, since it means that factors such as back-up causes which are extrinsic to the causal relation make an impact on whether C is the cause of E. By failing to respect the intrinsicity assumption, therefore, the

¹² For the classic statement of this view, see Lewis 1973.

¹³ See, for instance, Menzies 1996, Lewis, 2003a and Schaffer 2001a.

counterfactual analysis is left with the difficulty of trying to explain how causes with backups can be causes.

If we exclude all those factors which are extraneous to the causal relation in the analysis of that relation, the existence of a backup cause cannot affect the fact that C causes E. For the truth of the statement 'C causes E' will not depend upon anything outside the relation between C and E. So the intrinsicity assumption seems a good one to impose upon an analysis of causation, since it avoids the difficulty that counterfactual theorists find themselves in.¹⁴ The general lesson that can be taken away from preemption thought experiments is this: what matters in cases of causation is not what occurs in the environment around the cause and effect, but rather what happens between the cause and effect. The intrinsicity assumption incorporates this claim, by stating that it is the intrinsic relation between the cause and the effect which makes an entity the particular cause of its effect.

As well as committing us to the view that causation is a relation, what other implications does the intrinsicity assumption have for an analysis of causation? It is certainly not a trivial thesis, since it significantly restricts the range of causal analyses open to us. Nomological accounts of causation, for instance, which make laws the primary truthmakers of singular causal statements, are excluded. For on this view (regardless of whether the laws are Humean or non-Humean¹⁵) singular causal relations depend upon more general, nomological relationships that exist between the universals or event-types in question. So entities extraneous to the causal relation play a key part in determining that C causes E, contrary to the intrinsicity assumption. Similarly, the popular counterfactual or probability-raising views of causation are excluded. For, on these analyses, it isn't the intrinsic relation between C and E which is responsible for the truth of the causal statement 'C causes E'. Instead, the truth of 'C

¹⁴ Preempting or back-up causes also pose a serious problem for nomological and probability-raising accounts of causation.

¹⁵ For examples of Humean accounts of law, see Braithwaite 1927, Ayer 1953, Mackie 1974 and Lewis 1983b. For examples of non-Humean accounts of law, see Dretske 1977, Armstrong 1983 and Tooley 1987. Humean accounts claim that the metaphysical reality that underlies laws is nothing more than regularities in nature. The difference between laws and true generalisations lies at the level of linguistic reality. So law statements and statements about true generalisations are distinguished by the fact that they function differently in our theorising about the world. Non-Humeans, on the other hand, claim that there is some metaphysical reality which distinguishes laws from true generalisations. For more on this, see §7.2.

causes E' depends solely upon whether E is counterfactually dependent upon C, or upon whether C raises the probability of E. This, as we've seen, means that back-up causes, things external to the relation between C and E, can affect whether C causes E, contrary to the intrinsicity assumption.

Which analyses of causation are in line with the intrinsicity assumption? Certainly process theories of causation, which focus on the connecting line from cause to effect.¹⁶ For these take causation to be a local affair, determined by what goes on between the cause and effect. Aronson's (1971) and Fair's (1979) analyses of causation in terms of the transference of energy, for instance, or more recent, sophisticated developments of this view proposed by Salmon (1997) and Dowe (2000), would all accord with the intrinsicity assumption. The same is true of other process views such as Menzies's Camberra plan (1996), Ehring's persistence theory of causation (1997) and Rieber's property acquisition analysis (2002). Ehring takes the truthmakers to be persistent and partially persisting tropes,¹⁷ Rieber argues that (positive) causation should be understood in terms of the transference and acquisition of properties, while Menzies's leaves it open. But common to all these accounts is this claim that the truthmakers of singular causal statements are local and intrinsic to the process connecting cause to effect.

There is no need to commit to any one of these analyses. The intrinsicity assumption remains silent on what, for instance, is transferred, or whether any transference is required at all. Similarly, we need not suppose that causation is reducible to anything like energy transference or property persistence/acquisition. Instead, we could maintain some form of anti-reductionism about causation. Following Harré and Madden (1975), for instance, and developed by scientific

¹⁶ The modern, process views of causation have their origins in what O'Neill calls the "Scholastic influx model" (1993, p.37). This states that in causation, something is 'communicated' from the cause to the effect. This view is also echoed in the writing of the early moderns. Radner, for instance, argues that "Causation, for Descartes, is a matter of communication or impartment, and a thing cannot communicate or impart to another what it does not possess in itself" (1978, II) hence, Descartes causal principle that "There is nothing in the effect which did not previously exist in the cause" (1985, vol. II p.97). Whether or not this is the right way of reading Descartes, however, is a debated point (see, for instance, Clatterbaugh 1999).

¹⁷ Ehring does allow that laws may be involved too, so we would need to know more about the details of this aspect of his account (which he does not discuss in his book), in order to be positive that it would cohere with the intrinsicity assumption. However, even if it didn't, I suspect we could take what Ehring says as a starting point, and then develop a position which does.

essentialists such as Ellis and Lierse (1994, 2001), we could argue that particulars have fundamental causal powers. While these causal powers are constituted by the property instances of objects, the powers that property instances give rise to cannot be analysed in terms of anything non-causal. Causation could then be understood in terms of the manifestations of these local causal powers of particulars. So we could say something like this: C is the cause of E iff the property instances/causal powers of C and E give rise to a fundamental, intrinsic relation of causal necessitation between C and E.¹⁸

There are, then, a number of different accounts of causation which could be adopted, given the intrinsicity assumption outlined here. For this thesis doesn't offer us an analysis of causation, just a constraint on what that analysis should look like. This constraint, we've seen, places significant restrictions upon us. Any theory according to which the analysis of 'C causes E' depends upon widespread facts about the world, will be inconsistent with this assumption. However, I hope that this section has gone some way towards showing why we should place this restriction upon an account of causation. The claim that causation is an entirely local affair is not only motivated by the failures of the alternative approaches to causation, but also because it seems such a compelling, intuitive position to hold. It is thus worth hanging onto, if at all possible.

5.2 The Generality Assumption

The second of the two assumptions, the generality assumption, is based on the Humean observation that singular causal relations tend to be part of more general patterns. If C causes E, this seems to imply something about similar situations in which C-type and E-type entities are present. If, for instance, a balloon landed on something sharp and burst, we'd expect similar balloons to come to the same sticky end, if they landed on something equally sharp. This observation is often articulated by the principle of the nomological character of causation. This states that whenever there is a causal relation between two entities, C and E, there is a law that subsumes them (at least under some description of C and E).¹⁹ The idea is that because the

¹⁸ As we'll see (§5.4), Armstrong says something very similar to this.

¹⁹ See, for example, Davidson 1970 p.215.

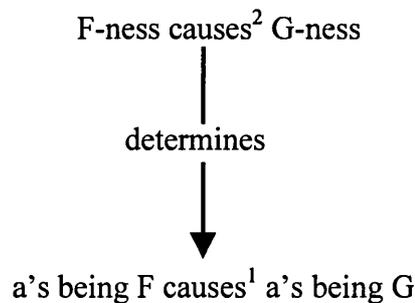
relation between C and E will be part of a more general causal pattern, it will follow from a more general causal truth or law.

The principle of the nomological character of causation has come under increasing pressure. I think that it is plausible, granted that we are flexible about the kind of laws causal relations can be subsumed under. If, for instance, we allow that causal relations can be subsumed under probabilistic and *ceteris paribus* laws, then it seems likely that the principle will be true. I shall not offer a defence of it here, however, as I don't think that such a strong principle is required to capture the intuition behind the generality assumption. But what I do want to defend is the related, slightly weaker, claim that singular causal relations give rise to more general causal patterns or general causal truths.²⁰ For this seems to be an observable fact about our world. If I drop acid on litmus paper and it turns red, other acidic solutions will turn litmus paper red if they come into contact with it. Similarly, if I break my toe by dropping a 10 stone boulder on it, boulders of a comparable weight are likely to break my toes if they are dropped on them. The converse is also true. If it's a general causal truth that smoking causes cancer, this indicates something about my chances of getting cancer if I smoke. Similarly, if there's a general causal truth stating that water evaporates at 100°C, then this implies something about the water in this kettle if it is heated to 100°C.

This close relationship that exists between singular and general causal truths requires, I think, some explanation. We need to say what the connection is between, for instance, 'those bad prawns caused Ellie's stomach ache' and 'bad prawns cause stomach ache in humans'. It is clear that they bear some relationship to each other. The fact that bad prawns cause stomach ache is not wholly irrelevant to the fact that Ellie gets stomach ache from those bad prawns. The generality assumption can be viewed as a demand for an explanation of this relationship. We find that singular causal relations are part of more general causal patterns and that they are thus subsumable under general causal truths which describe these patterns. So an adequate account of the causal relation must offer some account of this fact.

²⁰ This is only meant to apply to physical causation. Unfortunately, I don't have space to discuss the special issues that arise for mental causation here.

Singularists often seek to provide an account of the relationship between singular and general causal truths.²¹ They see themselves in opposition to the more traditional Humean or neo-Humean accounts of that relationship. It will be useful to the forthcoming discussion to distinguish between three different way of viewing the relationship between singular and general causal truths. The first of these is that of the generalist, which broadly states the position of Humeans and neo-Humeans.²² On this view, that there is a general causal relation between, for instance, bad prawns and stomach aches in humans, accounts for the fact that there is a singular causal relation between those bad prawns Ellie ate and her stomach ache. The generalist's model, then, looks something like this:



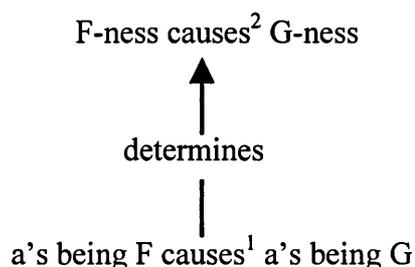
The fact that there is a singular causal relation (cause¹) between a's being F and a's being G requires that there is a general causal relation (cause²) between F-ness and G-ness.²³ For there is a relation of determination from the general causal relations to the singular causal relations. What does this mean? This notion of determination has proved difficult to analyse, but to illustrate the idea, consider the nomological account of causation. This states that singular causal relations, such as a's being F causes¹ a's being G, have to be subsumable under laws. For there only exists a causal relation between a's being F and a's being G *because* there is this nomic relationship between F-ness and G-ness. So laws are the primary truthmakers of singular causal statements.

²¹ This concern is central to many who are classified as singularists. Ducasse (1926), Anscombe (1971), Foster (1985) and Tooley (1987) all place great emphasis on this issue.

²² This terminology is taken from Ehring (1997).

²³ I shall suggest that there aren't two fundamentally different types of causal relation here, cause¹ and cause², whatever causation is going on is at the local level. But it has been said by some that cause¹ and cause² are different kinds of causal relations (see, for example, Eells 1991), so we shouldn't just assume that they are the same.

In the opposing, singularist camp, two different analyses of the relationship between singular and general causal truths are evident. I shall call the first of these AT-singularism, as it is suggested by both Anscombe and Tooley.²⁴ AT-singularism does not offer a positive characterisation of the relation between singular and general causal truths. Rather it can be identified with the negative claim that singular causal truths are *not* determined by general causal truths. The second singularist thesis, however, does offer a positive conception of this relationship, as it states that general causal truths are determined by singular ones. I shall call this view D-singularism, as I think it was first suggested by Ducasse.²⁵ The D-singularist's model, then, looks something like this:



The general causation between F-ness and G-ness is determined by the singular causal relations holding between instances of a's being F and a's being G. How does this work? The idea is that the truthmakers of the general causal truths are the many instances of singular causal relations holding between something's being F and something's being G. Therefore, the singular causal facts are prior to and determine the general causal facts, because the latter are built up out of the former.

What is the relationship between the intrinsicality assumption and AT or D-singularism? The connection between the two is weaker than we might expect. If we endorse the AT-singularist's claim that singular causal facts cannot be determined by general causal facts, nothing follows from this about the nature of the relation between the cause and the effect. It could be that entities extraneous to the causal relation affect whether this is an instance of C causing E. In which case, the relation would not be an

²⁴ See, for example, Anscombe 1971 p.104 and Tooley 1990 p.458.

²⁵ See, for example, Ducasse 1926 p.129. Bigelow and Pargetter (1990) also suggest this view. They write, "The transactions don't count as causal because they are subsumed under laws. The connection works the other way round: the causal laws hold because of the presence of local causal connections" (p.272).

intrinsic one, contrary to the intrinsicity assumption. For the same reason, we could say that singular causal facts fix all the general causal facts, without thereby being committed to the claim that the intrinsic relation between C and E is the sole truthmaker of the statement 'C causes E'. Therefore, neither AT nor D-singularism entail the intrinsicity assumption.

What of the converse claim? Does the intrinsicity assumption imply either AT-singularism or D-singularism? This is the question of the next section. There, I shall outline Anscombe's claim that the intrinsicity assumption is incompatible with the generalist's conception of the relationship between singular and general causal truths. If this is correct, then the plausibility of the intrinsicity assumption strengthens the case for the singularist's conception of this relationship. All is not well, however. For I've argued that both the generality and intrinsicity assumptions are highly plausible, so we need an analysis of the causal relation which can respect each of these theses. But in the next section, we'll see that these two theses are difficult to reconcile.

5.3 Anscombe's Challenge

Anscombe, in her famous article "Causality and Determination", gives us reason to doubt that the intrinsicity assumption can be justifiably combined with the claim that singular causal relations imply more general ones. She believes that singular causal facts cannot be determined by facts about general causal truths or laws, because general causal truths (which state, for example, what always happens) are unable to capture what is essential to causation. She writes,

Effects derive or come out of their causes...Now analysis in terms of necessity or universality does not tell us of this derivedness of the effect; rather it forgets about that. For the necessity will be that of laws of nature; through it *we* shall be able to derive knowledge of the effect from its cause, or vice versa, but that does not show us the cause as source of the effect (1971, p.92).

Anscombe thus endorses the intrinsicity assumption, because she thinks that an analysis of the causal relation must be given in terms of particular causes and effects. She goes on to suggest, however, that accepting this thesis undercuts our right to think that singular causal relations will be part of more general patterns.

Why is Anscombe sceptical of this claim? Her main argument for this, I think, is that the correct account of causation provides us with no reason to think that it will. According to Anscombe's analysis, A causes B if and only if B 'derives from', 'arises out of' or 'comes from' A.²⁶ But from this fact alone, we cannot deduce anything about other A-types or B-types. She writes,

If A comes from B, this does not imply that every A-like thing comes from some B-like thing or set-up or that every B-like thing or set-up has an A-like thing coming from it; or that given B, A had to come from it, or that given A, there had to be B for it to come from. Any of these may be true, but if it is, that will be an additional fact, not comprised in A's coming from B (1971, p.92).

Consequently, we have no reason to believe or expect that there will be general causal truths involving A and B, from the mere fact that A causes B.

What significance does this have for us? If we endorse an account of causation which incorporates the intrinsicity assumption, as Anscombe does, then it looks like we'll be in the same boat as her. For we will be committed to the claim that the truthmakers of causal statements are found solely among the intrinsic features of the causal relations reported. This makes it difficult to see what reason we could have for supposing that singular causal relations will be part of more general patterns, since these general patterns play no part at all in determining the causal relation. So general causal truths reporting these causal patterns will not be required in the account of what makes a singular causal statement true. It looks like we'll have to conclude, with Anscombe, that no singular causal statements, such as 'a's being F causes¹ a's being G', will give rise to more general causal statements of the form 'F-ness cause² G-ness'. The latter must be an "additional fact, not comprised" of a's being F causing¹ a's being G.

Trying to combine the intrinsicity assumption with the generality assumption, therefore, appears problematic. If we claim that the truthmakers of singular causal statements are entities which are local and intrinsic to those causal relations, then it looks like we have to deny that general causal facts determine the causal relations, as they will be extraneous to these relations. This leads to scepticism about the thesis that causal relations will be part of more general patterns. For why

²⁶ See Anscombe 1971, p.92.

should C causing E imply anything about other C-types causing E-types if, as the intrinsicity assumption claims, C is the cause of E solely in virtue of local, intrinsic facts of the relation between C and E? Anscombe's paper thus poses a challenge: why should the generality assumption hold, if we maintain, in accordance with the intrinsicity assumption, that general causal truths are not part of what makes singular causal statements true?

5.4 Armstrong's Reply

Armstrong's account of causation may be thought to offer a generalist's response to Anscombe's challenge.²⁷ He argues that all causal relations are subsumable under a law of this form: the universal F-ness nomically necessitates the universal G-ness.²⁸ Furthermore, he thinks that a relation's being subsumable under a law is what makes it a causal, rather than non-causal, relation.²⁹ Take, for instance, the causal relation, *Fa causes Gb*. The relata a and b, we are assuming, instantiate the universals F-ness and G-ness respectively. Similarly, as the irreducible relation of nomic necessitation is a universal, we can suppose that this too is instantiated by the relation. Indeed, according to Armstrong, what makes this relation a causal one, is that it instantiates this irreducible nomic necessitation universal. So his account respects the generalist's thesis that causal laws determine causal relations.

It is clear, then, that Armstrong's account of causation complies with the generality assumption, since being an instance of a law is what makes a relation causal. Thus, any relation which is causal will be part of a more general causal pattern which is subsumable under some law. It is far from clear, however, how the theory coheres with the intrinsicity assumption. How could the truthmakers of causal

²⁷ Comments made in his 1999b paper illustrate that Armstrong thinks that his account combines the nomic character of causation, with the claim that "the causal structure of a process is determined solely by the intrinsic character of that process" (p.184).

²⁸ Although he maintains that this is an a posteriori, not a priori, truth (see 1997 §14.7).

²⁹ I believe that this is a fair interpretation of Armstrong's view. In his 1999b paper, it looks initially as if he is not committed to this claim. He writes, "The solution that I recommend to the problems posed by the neuron diagrams is very simple. Where there is an arrow in a diagram showing that one neuron brings it about that another neuron fires, take it that there is a genuine two-term relation of singular causation holding between cause and effect...This is the open door" (p.176). Later on, however, it becomes clear that Armstrong does intend to offer a deeper analysis than this suggests, as he writes, "my claim is that instantiation of a law gives us the essence...of singular causation...singular causation is instantiation of a certain sort of law" (p.184).

relations be entities which are intrinsic to those particular relations, if general causal truths or laws are what determine that these relations are causal? The answer to this question seems to lie in Armstrong's characterisation of universals. His account can be summarised by the following three claims. First, universals are said to be spatiotemporal entities which are spatially 'in' their instances. Armstrong wants to endorse this claim because he believes that "the world, the totality of entities, is nothing more than the spacetime system" (1997, p.5). Second, universals are thought of as inherently general entities, because they can exist at many different places at the same time. He thus has to deny that spatiotemporal entities have one spatiotemporal location at a time. Finally, universals are said to be entities which are "present completely in each instantiation" (1999b p.184).

Given this characterisation of universals, it looks like Armstrong's theory is consistent with the intrinsicity assumption. Take, again, the causal relation, *Fa causes Gb*. The relation's instantiation of the irreducible universal of nomic necessitation is what makes it a causal relation. In other words, this is the primary truthmaker of the singular causal statement 'Fa causes Gb'. If this account is going to cohere with the intrinsicity assumption, therefore, this universal of nomic necessitation will have to be an intrinsic feature of the causal relation. This is what Armstrong's conception of universals seems to ensure. On his account, an instantiation of the universal of nomic necessitation is not only a spatiotemporal entity located within the boundaries of its instantiator, it is also something which is wholly present in it. So an instantiation of the nomic necessitation universal is something which can be intrinsic to causal relations, as the intrinsicity assumption demands.

Armstrong's proposed reconciliation of the generality and intrinsicity assumptions, therefore, challenges Anscombe's claim that respecting the intrinsicity thesis requires denying that the truthmakers of causal relations are general causal truths. Armstrong's account of universals makes it look possible to satisfy this requirement, while still having laws as the truthmakers of singular causal statements. But how convincing is his reconciliation? I shall argue that it falters, because of the theory of universals it relies upon.

Armstrong states that universals are concrete (i.e. spatiotemporal) entities, but this sits uneasily with the claim that universals can be wholly present in many

different places at the same time. Why does this tension arise? It does not stem from the claim that a concrete entity can be at different places at the same time. Although this is controversial, it seems reasonable to say that a scattered pack of cards, for instance, is in different places simultaneously.³⁰ The difficulty rather arises because Armstrong says that a concrete entity can be *wholly present* at different places at the same time. This seems implausible. How could something be completely at one location at the same time as being completely somewhere else? Perhaps this could be true of an abstract object (at least given some non-spatial correlate of ‘present’), since these are not part of the spatiotemporal system and so are not subject to the same spatiotemporal boundaries that concrete entities are. But it seems almost a truism to say that a spatiotemporal entity, scattered in space, is only partly present at any one of its locations.

No doubt Armstrong will object that I am utilising intuitions tailor made for concrete particulars, not concrete universals. But if I am right to think that, at least in common usage, being ‘wholly present’ just means something like ‘all the parts of the entity are present in that particular spatiotemporal region’, there is a serious problem for the proposed reconciliation here. It is clear that Armstrong has to mean something else by the phrase ‘wholly present’, since he has to claim that there are spatiotemporal entities which are wholly present at different places simultaneously. So he owes us another account of what he means by this phrase. But even if Armstrong offered us such an analysis, which he doesn’t, it is unlikely that his theory will be able to reconcile the intrinsicality and generality assumptions. For the intrinsicality assumption demands the usual meaning of ‘wholly present’, as when we say that the truthmakers of causal statements are intrinsic and local to those relations, we are depending upon this idea that the entity’s parts are all in that spatiotemporal region. Therefore, not only does an extremely counter-intuitive thesis lie at the heart of Armstrong’s account, it is doubtful whether this thesis gets us the reconciliation we want any way.

If we avoid Armstrong’s problematic account, and opt instead for the view that universals are abstract entities, the intrinsicality assumption still eludes us. For an instantiation of the universal of nomic necessitation will depend upon an entity

³⁰ Goodman (1951) and Quine (1960), for instance, talk about scattered objects.

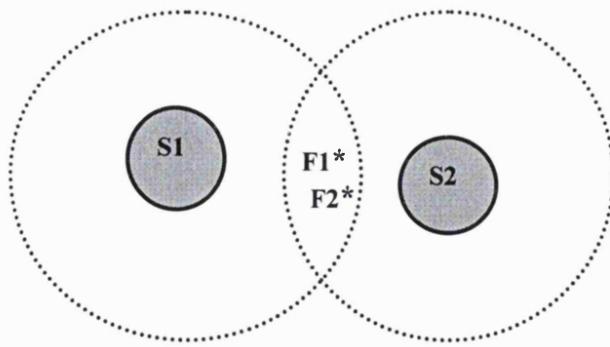
extraneous to the spatiotemporal location of the relata and the relation. If the non-spatiotemporal realm was eradicated, this relation would no longer instantiate nomic necessitation, so this feature is not intrinsic to the relation. The kind of generalist account Armstrong offers, therefore, fails to reconcile the intrinsicity and generality assumptions, whether or not an abstract or concrete theory of universals is endorsed.

5.5 An AT-Singularist's Response

AT-singularists do not offer a positive conception of the relationship between singular and general causal truths. Nevertheless, if we follow the lead of Foster and Tooley, their view might be thought to offer a response to Anscombe's challenge. Foster and Tooley are AT-singularists because they believe that not all singular causal facts are reducible to or supervenient upon facts about causal laws.³¹ Foster and Tooley endorse this thesis because they believe that there are possible situations in which the indeterministic causal laws and non-causal facts fail to determine which cause is paired up with which effect. Consider, for instance, a case in which two heated spheres are the cause of two flashes.³² Foster asks us to suppose that there is a law stating that when any spherical lump of a certain kind of metal, K, reaches temperature t, a flash will appear half a second later at some unspecifiable point near the K-sphere. Imagine that two K-spheres are put sufficiently close together so that the surrounding flash areas overlap. Then both lumps are heated and reach the critical temperature at the same time. Half a second later, we see two simultaneous flashes occur within the region of overlap.

³¹ See Foster 1985 part III, §7 and Tooley 1987, Ch. 6 and 1990. Foster formulates the thesis in terms of reduction and Tooley in terms of supervenience, but I have grouped the two views together because they are both trying to convey the same idea. It is not entirely clear, however, whether Foster ends up endorsing this position, because he concludes that the generalist can respond to the arguments in favour of this thesis (see p.263). Nevertheless, he does seem very sympathetic towards the view, so for the purposes of this exposition, I shall just presume that he does hold it.

³² This is Foster's version of the argument (see 1985, p.256). Tooley offers a very similar thought experiment, which works in the same way (1990, p.459-460). I think that this is the most convincing of the examples offered by Foster and Tooley, as it involves the least controversial assumptions.



Possible Options:

- a) S₁ causes F₁ and S₂ causes F₂
- b) S₁ causes F₂ and S₂ causes F₁

In this case, Foster argues it is intuitively plausible to suppose that each flash is caused by one of the spheres at temperature t .³³ The law and the non-causal description of all that happens, however, does not suffice to determine which sphere caused which flash, as either flash could belong to either event. In other words, we cannot distinguish between possibilities a) and b), because the non-causal description and the law are neutral between these alternatives.

Both Foster and Tooley argue that in these kinds of cases, there is a fact about the matter concerning which sphere caused which flash. Since these pairings cannot be determined by the laws and the non-causal facts, we have to postulate an irreducible singular causal fact to do this. So the causal laws (plus non-causal facts) cannot determine all the singular causal facts. This conclusion, however, is compatible with the thesis that causal relations are subsumable under causal laws. For we might think that, despite this failure of supervenience, all relations have to be subsumed under some law in order to be causal. This would be the case, for instance, if we thought that laws, while not the sole truthmakers of singular causal statements, were nevertheless among them. Indeed, this sort of position was once endorsed by Tooley.³⁴ Thus, their view does not force us to reject the generality assumption.

We have, then, an account which can combine the generality assumption with the AT-singularist's thesis that singular causal facts are not determined by general causal facts. Furthermore, AT-singularism is compatible with the intrinsicity assumption, as Anscombe's theory clearly incorporates both these claims. This alone,

³³ This is another illustration of the pairing problem. It is intuitive to claim that F₁, for instance, is either caused by S₁ or S₂ because we think that causation involves an intrinsic relation between the cause and the effect. The fact that the generalist's reply to this kind of case (namely, that there are no causal pairings beyond that determined by the laws and the non-causal facts) seems unsatisfactory, again attests to the intuitive character of the intrinsicity assumption.

³⁴ See Tooley 1987 §6.6 and Ch. 8. In his 1990 paper, however, he seems to have changed his mind.

however, does not suffice to show that the account meets the intrinsicity assumption, since earlier we saw that the assumption doesn't follow from AT-singularism. To illustrate this, consider the case just given. There, laws were said to be part of the truthmakers of causal relations. To make the example more concrete, suppose that laws are analysed via Lewis's 'best-system theory'. If these are among the truthmakers of causal statements, then the intrinsicity assumption will not hold. For Lewis's laws are just those global patterns of reoccurrence which appear in the best scientific systems. So the causal relations will be partly determined by entities which are extraneous to them.

Is there any reason to suspect that those AT-singularist accounts which incorporate the generality assumption will be unable to respect the intrinsicity assumption? There are grounds for scepticism. AT-singularists, such as Tooley, were only able to integrate their view with the generality assumption by undermining the intrinsicity assumption. In other words, rather than saying that facts about the causal relation determine all the singular causal truths, those local facts were only said to *partly* determine the singular causal truths. This makes room for the claim that causal relations are subsumable under laws, as now we are able to say that these laws are among the truthmakers of singular causal statements. So every singular causal statement will still entail the existence of a law. But this explanation of why causal relations satisfy the generality assumption was just the one that Anscombe wanted to avoid. For it manages to reconcile AT-singularism with the generality assumption by utilising the generalist's justification for this thesis. Given this proposal, part of what it is to say that 'A causes B' is comprised by the general causal truth 'A-types cause B-types', as this general causal statement is incorporated into the analysis of the singular causal statement. But Anscombe's challenge was to explain why we should accept the generality assumption, once we have rejected this generalist claim.

Now Anscombe did have good reason to insist upon this, since it is difficult to see how the intrinsicity assumption could hold even if laws are just one of the truthmakers of causal statements. Laws concern facts about what happens in all aeons. Therefore, if they are responsible for the truth of causal statements, it looks as if causal relations cannot be determined solely by entities which are local and intrinsic to those relations. This unresolved difficulty justifies the limited conclusion that Foster and

Tooley's account gives us no indication of how we might go about reconciling the generality and intrinsicity assumptions.

5.6 CTP Solution

The coveted reconciliation, then, still seems a long way off. But, fortunately, Ducasse's version of singularism points a way forward. He writes,

a causal connection explains the regularity of the succession, but it is not constituted by such regularity, which is but a corollary of the causal connection (1926, p.130).

The thought here is that although regularities of E-types following C-types are not part of what makes the singular causal statement 'C causes E' true, these regularities are nevertheless implied by the causal relation. But why is this? Ducasse suggests that it is because causal laws are constituted out of "a class of resembling facts", each of which "already happened to be a causal fact individually and in its own right" (1926, p.129).

The form of singularism presented by Ducasse, therefore, offers a radical break from the generalist's tradition. Rather than just saying that singular causal facts are not determined by general causal facts, as AT-singularism says, D-singularism goes further and claims that it is singular causal facts that determine the general causal facts. Unfortunately, however, it is still not clear why the generality assumption should hold. It does not follow from Ducasse's claim that laws are constituted out of "a class of resembling facts". For the assertion that laws are generalisations over singular causal facts does nothing to guarantee that these singular causal facts will always give rise to more general ones. Consequently, there may be causal relations which aren't part of more general causal patterns, and so which aren't subsumable under any law. But even if every singular causal relation, as it happened, was part of a general causal pattern, this fact would be coincidental, it would not follow from anything Ducasse says. So we still have no explanation of why singular causal relations are subsumable under general causal truths. What we require, then, is an analysis which incorporates the D-singularist's thesis with an account of why the existence of a causal relation guarantees the existence of a more general causal pattern. For this would enable us to explain why singular causal relations are subsumable under general causal truths, without thereby endangering the intrinsicity constraint.

Is it possible to formulate such an account? I think it is, if facts about the tropes of singular causal relations are taken to be the truthmakers of singular causal ascriptions. The proposal being put forward then, appeals to the aforementioned trope analysis of property instances.³⁵ Although there are different characterisations of these entities, we've seen that trope theorists are united in taking tropes to be *sui generis* property instances. In other words, they are property instances which cannot be analysed in terms of anything more basic, like instantiations of universals or as members of sets of possible particulars. In addition to trope theory, the account will also appeal to CTP, as this analysis of the nature of properties can provide us with what Ducasse's suggestion lacked, namely, a reason for thinking that causal relations will give rise to more general causal patterns.

How does CTP, in combination with the claim that facts about the tropes of singular causal relations are the truthmakers of singular causal ascriptions, provide us with an account which satisfies both the generality and intrinsicity assumptions? Let's begin by considering the generality assumption. Tropes are unlike properties in that they only have one spatiotemporal location. Nevertheless, because they are *property* instances, many trope theorists claim that universals or properties can be reduced to equivalence classes of these tropes.³⁶ How are these tropes grouped together into classes? Usually this is done by appealing to the relation of exact resemblance. But this is rather imprecise and it is not clear that we want to say that all tropes in a class have to exactly resemble each other. If the tropes were property instances of the determinable red, for instance, then this criterion would be too stringent. This standard account can be improved upon, however, if we utilise CTP, as this offers a more precise account of how tropes (which stand for genuine properties) are grouped into classes. According to CTP, tropes belong to the same class, and are thus instances of the same property, iff they meet certain causal requirements laid down by the functional definition for that property. Trope c^1 and trope c^2 , for example,

³⁵ See §3.3. For defenders of trope theory see, for instance, Husserl (2001), Stout (1921), Williams (1953), Mulligan, Simon and Smith (1984) Campbell (1990), Bacon (1995), Ehring (1997), Schaffer (2001b). See §6.4 for a brief discussion of the nature of tropes.

³⁶ See §1.5 and §3.3. A trope theorist does not have to say this, although it is the majority view (see, for instance, Stout 1921, Williams 1953, Campbell 1990 and Bacon 1995). We could think of the universal as something distinct from the class of tropes, but this raises issues that I shall be addressing in the next chapter.

are members of the set which stands for the universal of charge, iff both c^1 and c^2 realise the causal role described by the functional definition for the property of charge.³⁷

All I am employing here, therefore, is the functional reading of CTP outlined earlier. If we adopt the functional role analysis proposed, then this already includes an ontology of tropes. If, on the other hand, we just want to endorse this functional two-level criterion of identity, then this, combined with tropes, gets us our generality assumption. In order to see this, consider the causal relation, *Jack's fall caused his crown to break*. The events *Jack's falling* and *his crown breaking* possess certain tropes, the most crucial ones being Jack's trope of falling and his crown's trope of breaking. These tropes belong to sets, the members of which all realise the same nexus of causal relations. Thus, if Jack's trope of falling and his crown's trope of breaking are causally related (given certain circumstances), then Jill's trope of falling and her crown's trope of breaking will also be (given the same circumstances), since Jill's trope of falling and her crown's trope of breaking enter into the same causal relations as Jack's trope of falling and his crown's trope of breaking.

The existence of the causal relation between Jack's trope of falling and his crown's trope of breaking, therefore, guarantees the existence of a more general causal pattern. This general pattern will have a law-like description of the form: in circumstances C, the property of falling nomically necessitates the property of breaking. So we can say that the singular causal statement, 'Jack's fall caused his crown to break' is subsumed under a law, because it is an instance of this more comprehensive causal truth. The same will also hold for all other causal relations, but this won't be because causal laws are truthmakers of singular causal statements, rather it is because singular causal relations automatically give rise to laws.

On this view, therefore, the generality assumption will be true of a world, if there are properties at that world. For, given CTP, tropes are guaranteed to give rise to more general causal truths about properties. So far in this analysis, however, no appeal to tropes has been made. For property instances, regardless of their analysis, would do the job just as well. Why, then, do we need an ontology of tropes? The reason for this

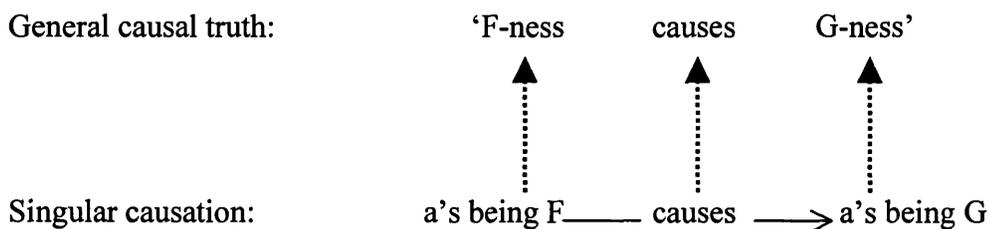
³⁷ More on how tropes should be constructed into classes and what, more generally, the relation is between properties and tropes will be discussed in the next chapter.

is because, without them, we would be unable to satisfy the intrinsicity assumption. We want this version of D-singularism to not only cohere with the generality assumption, but also with the claim that the truthmakers of singular causal statements are entities which are intrinsic (in the sense outlined earlier) to that relation. How do tropes enable us to do this? Unlike instantiations of universals, tropes do not partly consist of an entity which is abstract (i.e. outside space and time). Nor are they constituted out of Armstrong's universals, so we don't have to endorse the dubious claim that scattered spatiotemporal entities can be wholly present in each instance. Instead, a trope is a simple entity with a single spatiotemporal location. So tropes can be thought of as wholly present in the entities which possess them.³⁸

In the next chapter, I shall outline two different ways of thinking about tropes, both of which render them intrinsic to their objects. But from the bare-boned conception given here, I think it should be clear that tropes are the kind of entities which can be wholly present and intrinsic to the entity that possesses it. For they are not analysable or dependent upon anything other than that entity. Consequently, a trope will not entail the existence of any other entity outside that which has it. The upshot of this is that if tropes are the truthmakers of singular causal statements, then those truthmakers will be intrinsic to the causal relations. The account will thus be able to satisfy the intrinsicity and generality assumptions.

So functional role theory, or some other version of CTP combined with an ontology of tropes, provides a cogent rapprochement of the intrinsicity and generality assumptions. This, I have argued, is no mean feat to pull off. But it is a worthwhile objective, because of the plausibility of both of these theses. The resulting theory is a form of D-singularism, since it claims that facts about singular causal relations are the truthmakers of general causal truths. Diagrammatically, then, it looks something like this:

³⁸ Some trope theorists even think that their spatiotemporal location is what individuates them. See, for instance, Schaffer 2001b.



The normal arrow stands for the relation of singular causation, whereas the dashed arrows stand for the relation of truthmaking. The diagram addresses the issue of what makes true general causal claims, such as ‘smoking causes cancer’, and how these truths relate to singular causal relations, such as Jim’s smoking caused his cancer. On this CTP singularist view, whatever causation is going on is at the local level. All general causal truths are determined by these singular causal relations. For laws are relations between properties, and these properties are comprised out of tropes – the truthmakers of singular causal statements.

5.7 Causal Relata Again

Before leaving this singularist model of causation behind, I first want to take a step back, and consider what implications the discussion here has for an analysis of the causal relata. Part of thinking about what the relation is between singular and general causal truths, involves considering this question: what is the relationship between the singular causal relata and the general causal relata? Two accounts of this relationship suggest themselves. The first is to say that the singular causal relata are ontologically dependent upon the general causal relata. In other words, the singular causal relata could not exist without the general causal relata. The second is to say that the general causal relata are ontologically dependent upon the singular causal relata. For while the general causal relata do not directly depend upon any particular singular causal relata, they are built up out of an array of them.

In what follows, I shall defend the second of these views, as I shall argue that the considerations of this chapter support the claim that tropes are the singular causal relata. In chapter four we saw that fine-grained entities make good causal relata. But the argument there didn’t lead exclusively to tropes, or even property instances more broadly construed, since other accounts of the causal relata, such as fact theories or

Lewis's analysis of events, also postulate fine-grained entities.³⁹ Here, however, I hope to show that at least some of these alternatives face problems respecting one of the following: the generality assumption, the intrinsicity assumption or a plausible extension of the intrinsicity assumption. I shall begin by outlining this extension of the intrinsicity assumption. Once this is in place, I shall then explain why I think a causal ontology of tropes has the edge over its fine-grained competitors.

We've seen that the intrinsicity assumption is motivated by the desire to respect the claim that causation is an entirely local affair. But what if the cause and effect are not themselves entirely present in the spatiotemporal region where the causing takes place? So, for instance, if we thought that events were spread out over actual and possible spatiotemporal regions, then part of the cause and the effect would not be present in the spatiotemporal region where this ball hits and breaks this window. Such a position would not, strictly speaking, contravene the intrinsicity assumption, because the relation between the cause and effect could still be an intrinsic one. But it is not difficult to see why this claim goes against the spirit of the intrinsicity assumption, since part of the causal relata is not present where the causation is taking place. The motivation behind the intrinsicity assumption, therefore, supports this further claim: in an instance of causation, the cause and effect is wholly present in the spatiotemporal region where the causation is taking place. Let's call this the local relata principle. We'll see that some analyses cannot accommodate this plausible principle.

Consider first the view that causes and effects are instantiations of universals.⁴⁰ This analysis has difficulties meeting the local relata principle. For if we analyse property instances as instantiations of universals, then property instances must either

³⁹ See Lewis 1986b. There he argues that events are classes of actual and possible spatiotemporal regions. The event of Socrates's death, for instance, is the class of actual and possible spatiotemporal regions in which the subject, Socrates, dies. Because Lewis thinks that events are classes of spatiotemporal regions, they have a mereology. Every spatiotemporal region of a world which is a member of an event class can be treated as a part of that class. Lewis uses these parts in his account of the essences of events. We can capture the claim that an event is essentially a death, for instance, if in each of its actual and possible spatiotemporal regions or parts of the class, a death occurs. This means that Lewis's events can be very fine-grained, and he is perfectly willing to allow that the demands of our causal theory may lead us to postulate two different (though not entirely distinct) events when, intuitively, we would only have thought there to be one.

⁴⁰ The following will also hold of those accounts which take Russellian facts, i.e. complexes of particulars and universals, to be the relata of causation. Fregean facts would face similar difficulties too.

be partly constituted out of an abstract entity which is outside space-time, or a concrete, general entity which exists at many different places at the same time. The first conception fails to meet the local relata principle because part of the cause exists in an abstract realm, it isn't present where the causation is taking place. The second conception doesn't seem to face this problem, if we combine it with the claim that universals are wholly present in each place they are. But, as I argued in §5.4, this analysis of universals drains the expression 'wholly present' of its standard meaning, and it is this standard meaning which the local relata principle appeals to.

This isn't all, however. In order to explain why singular causal relata give rise to more general causal truths about types of entities, we need to postulate some kind of dependence relation between the two. Now if we embrace a causal ontology of instantiations of universals, the key constituents of the causal relata are going to be these universals or general entities. So we have to appeal to them in order to explain why singular causes and effects behave as they do. As a result, the singular causal relata will be ontologically dependent upon the general causal relata. For the ability of particular causes to function as they do will depend, at least in part, upon what general causal relations its universals can stand in. This leads to troubles with the intrinsicity assumption. The generalist's model of the relation between singular and general causal truths is forced upon us, as singular causal truths will be partly fixed by what general causal or nomic relations these universals occupy. So part of what makes singular causal statements true will be entities which are extraneous to their relations, contrary to the intrinsicity assumption.

I am not claiming that all those who think that the causal relata are instantiations of universals have asserted this. My point is only that once we turn our attention a little further from the relata of causation, and consider what the truthmakers of singular and general causal statements are, then what we say about the ontology of the causal relation will have important ramifications. Hence, these consequences should be taken into consideration. This point has, to a large extent, been overlooked. Consider, for instance, Menzies's example of reductive singularism.⁴¹ He writes,

⁴¹ By singularism, Menzies means any view which endorses the intrinsicity assumption, see §5.1. But, as we'll see, he has a different sense of 'intrinsicity' in mind.

Let us assume that there are no causal relations among the perfectly natural relations instantiated in the actual world, over and above the spatiotemporal relations; and that causal relations are intrinsic in the specific sense that they supervene on the perfectly natural properties of spacetime points and point-sized bits of matter that make up their relata and on the spatiotemporal relations holding among them. Then it follows that these causal relations conform to Humean supervenience (1999, p.325).

Menzies states that this view conforms with the intrinsicity assumption. But I think that unless we say more about the nature of these properties and relations, the issue is not clear cut. For these properties may be abstract entities, so the truthmakers of singular causal statements will not be local and intrinsic to the causal relation. This disagreement can be explained by the fact that Menzies endorses Lewis's analysis of intrinsicity in terms of perfectly natural properties and duplicates. But earlier I argued that this understanding of intrinsicity doesn't capture the intuition motivating singularism (and the grounding intuition).⁴² So in order to get a reductive account which respects the intuitive force of singularism, which Menzies wants, we need to say more about the ontology of causation than is said here.

Granted this interior notion of intrinsicity, therefore, a causal ontology of instantiations of universals leads to difficulties with the local relata principle and the intrinsicity assumption. But what of the other possible fine-grained causal relata, like Lewis's events? He argues that the relata of causation are events (with essences), so we are free to say that the general causal relata are ontologically dependent upon the singular causal relata, because the event-types which appear in general causal truths are built up out of particular instances of events which appear in the singular causal truths. Adopting this view, however, does not ensure that the local relata principle will hold. For Lewis's events are classes of actual and possible spatiotemporal regions. So the event of Don's gripping lightly, for instance, isn't wholly present at the time and place of Don's gripping, as the event is composed of many other actual and possible regions of space-time.⁴³

If we leave this aside, however, there is still a question about what explanation an account like Lewis's can offer of the fact that singular causal relata usually imply

⁴² See §3.1 and §5.1.

⁴³ I think that Yablo's account of events (1992a), which also appeals to essences, avoids this problem while still being fine-grained. But I think that it is vulnerable to the generality assumption, in the same way that Lewis's is.

the existence of more general causal truths about types of those relata. Lewis can say that we group events into event-types on the basis of shared similarities. So we'd get classes of spatiotemporal regions being grouped with further classes of spatiotemporal regions, on the basis of shared properties which are themselves just classes of possible particulars! But even if we are prepared to accept this, there is still no reason to expect that types of events, grouped together on the basis of shared similarities, will behave in the same way as instances of events.

If we embrace CTP trope theory, we have an explanation of why singular causal relata give rise to more general causal truths about properties. Tropes belong to the same class and are thus instances of the same property iff they meet certain causal requirements laid down by the functional definition for that property. So singular causal relata give rise to more general causal truths about the behaviour of properties, because properties are constructed solely out of tropes, all of which have to realise the same functional role. Tropes are also in a unique position to respect the local character of the causal relata because, unlike universals and sets of spatiotemporal regions, tropes can be wholly present where the causation is taking place. A causal ontology of tropes, therefore, preserves the local character of causation, without thereby losing sight of the fact that more general causal truths about types of entities are implied. This, I think, is a significant advantage of the trope account over other analyses. Although we've not looked at every theory, we've seen that universal instantiations and Lewis's events are not in a position to do this. So the considerations of this chapter further support the claim that property instances are the fundamental causal relata. Although now, of course, more demands have been placed upon the nature of these property instances.

6. Properties and Their Instances

The question which I shall focus on in this chapter is this: given CTP and the other considerations put forward here, how should properties and their instances be viewed? What we need, is a causal theory of properties which not only avoids the problems which were raised earlier against CTP,¹ but which also provides a causal ontology which accords with the fine-grainedness of causes and the intrinsicity and generality assumptions. Within the general outline, we'll see that divergent theses could be adopted. These alternatives will result in different causal analyses of properties and their instances, but my aim is not to decide between them. Instead, I shall limit myself to indicating what the various options and their consequences are.

6.1 CTP Expounded

The heart of CTP can be expressed by this two-level criterion of identity:

Two property instances/tropes are instances of the property F-ness iff they both realise the causal role stated in the RCL definition for property F-ness.

This criterion, regardless of whether it is read reductively or non-reductively, commits us to a certain view about what it is for an object to instantiate a property:

An object instantiates the property F-ness iff the object satisfies causal requirements R, stated in the RCL definition for property F-ness.

One of the things CTP offers, then, is an account of property instantiation, as it analyses what it is for an object to instantiate a particular property in terms of the causal powers that object has.² If an object instantiates the property F-ness, that object must meet certain causal requirements laid down by the RCL definition for F-ness. All

¹ See §2.4, 2.5 and 3.1.

² We can extend this analysis to events. We shouldn't talk of properties bestowing causal powers onto events, as we only ascribe causal powers to those entities which we think of as persisting. But we can say that properties bestow causal efficacy onto events. So a property is instantiated by an event, if that event is able to occupy a certain causal role because of that property.

the objects which meet these requirements can thus be said to be alike in a certain respect, because they all instantiate the property F-ness. Another way of putting this is to say that all the objects have a property instance or trope of that property, since these are just particular instances of properties. According to CTP, it is these property instances or tropes which enable an object to meet the causal requirements specified by the RCL definitions for universals. So property instances or tropes are responsible for an object's causal powers.

This account of property instantiation introduces two entities onto the scene. First, there are the properties. These are general entities which can be instantiated in many different places at the same time. Second, there are the particular instances of properties. These are the particular ways that objects (and events) are, which enable them to behave as they do. According to CTP, a property is instantiated by an object if that object possesses a trope or property instance which enables that object to meet the causal requirements of that property. But this account of instantiation leaves a number of questions unanswered. Most crucially, what is the relationship between properties and their instances? And, what is the relationship between a property instance or trope and its object? This section will deal with the first of these issues. In §6.4. something will be said about the second question.

On the issue of the relationship between properties and their instances, there seem to be two options which would cohere with CTP and the causal considerations given already. The first of these is trope nominalism. This is the view of traditional trope theorists like Stout (1921), Williams (1953) and Campbell (1990). They argue that *sui generis* properties or universals need not be postulated, as classes of tropes can do the work that universals are supposed to. Not just any class of tropes could sensibly count as a property, however, since then there would be innumerable properties. The class of {red trope, wise trope, charm trope}, for instance, would count as a property if this were the case, because these tropes are members of a class. This can't be right. If two objects instantiate the same property, then they must be similar in some respect, as properties are meant to ground genuine resemblances between objects or 'carve nature

at the joints'.³ Trope theorists must thus impose some condition of entry onto a class which stands for a property, so that these classes are not just disparate collections of tropes. Typically, this is done by appealing to the relation of resemblance. Williams and Bacon, for instance, say that tropes are grouped together into classes on the basis of their "similarity", while Campbell talks of tropes being "like" each other. So on this view, tropes are grouped together into classes because of their primitive resemblances to each other.⁴

This proposal, however, is vulnerable to the complaint that we don't have a good grip on what condition a trope has to meet in order to count as a member of a certain class. Trope theorists sometimes limit themselves to the relation of, what Williams calls, "precise similarity" (1953 p.117). If trope a and trope b are part of the same class then, according to Williams, they must resemble each other exactly. But this is still not as clear as we might hope. For even if we grant that a trope is a member of the F-ness class iff it is exactly similar to all the other tropes in this class and not to any tropes in other classes, we still have no account of what makes them exactly resembling tropes of F-ness, rather than exactly resembling tropes of G-ness, or H-ness. Furthermore, in some cases, it might not be clear what it is for two tropes to exactly resemble each other. Two tropes of charge will, presumably, not be identical in every respect because, according to trope theorists, these entities are particular. Consequently, we have to say that their qualitative (non-particularised) aspects are precisely similar. But what does this amount to for charge tropes? We can't just imagine what two precisely similar tropes of charge look like, as we can with tropes of red. Therefore, more elucidation as to what exactly this resemblance consists in would, I think, be beneficial.

³ This famous saying is inspired by Plato, who in *Phaedrus* writes, "The second principle is that of division into species according to natural formation, where the joint is, not breaking any part as a bad carver might" (1970, 265d-266a).

⁴ See Williams 1953, p.116, Bacon 1995, p.13 and Campbell 1990, p.31. Stout (1921) offers another account of how tropes should be grouped into classes (this is the only other view I have come across). He argues that some tropes form primitive natural classes. These primitive natural classes are distinguished by their distributive unity, an ultimate and unanalysed feature of some classes. This account has not found much favour in recent times, however, probably because we have an intuitive idea of what the relation of resemblance involves, unlike that of distributive unity. So, in what follows, the traditional trope view shall be identified with the claim that tropes are grouped together into classes because they resemble each other.

This is what CTP trope nominalism can offer. CTP provides a way of grouping tropes together into classes which stand for natural properties. It says that any object which satisfies causal requirements R, laid down by the RCL definition for property F-ness, has a trope of F-ness. So the class of tropes which is the property F-ness, according to CTP trope nominalism, is the class of tropes which realise the causal requirements definitive of property F-ness. How does this compare to the resemblance account? Favourably, I think, because it offers some elucidation of what this relation of resemblance amounts to. Take, for instance, a trope of charge and let us suppose that tropes of charge are ‘precisely similar’ to each other.⁵ What does this relation of precise similarity amount to, given that the entities in question are not literally identical? CTP provides an answer to this: for charge trope₁ to be precisely similar to charge trope₂ is for charge trope₁ to realise exactly the same functional role as charge trope₂. Moreover, CTP provides an explanation of why trope₁ is an instance of the property of charge rather than an instance of the property of mass. For we can say that trope₁ realises the RCL definition for the predicate ‘is charge’ rather than the RCL definition for the predicate ‘is mass’. So CTP offers a more precise specification of what membership into a class of tropes, which stands for a natural property, involves.

If we adopt CTP trope nominalism, therefore, the class of tropes which stands for a property is the class of tropes which satisfy the condition specified by CTP. For instance, the class of tropes which is the property of being knife-shaped is the class of tropes which realise the causal relations stated in the RCL definition for the predicate ‘is knife-shaped’. This account needs expanding upon somewhat, in order to account for the differences between determinates and determinables, but the basic claim is clear enough. This theory, then, offers a very definite conception of the relationship between tropes and properties. Properties are just collections of tropes which meet certain conditions, so a trope is literally part of a property, while a property is literally made-up of a number of tropes.

This account of the relationship between tropes and properties coheres well with the causal claims made here and can be viewed as a form of strong CTP. It provides a way of understanding the claim that properties are nothing over and above their causal profiles. For if properties are just sets of tropes which realise such-and-

⁵ This makes them instances of determinate universals, see §6.2.

such causal roles, then properties are exhaustively characterised by the causal profiles of other entities. As a result, the account is committed to the claim that properties are individuated by the causal contributions of their tropes. The view also provides the right backdrop for the singularist theory of causation outlined in the last chapter. If tropes are the truthmakers of causal statements and properties are literally constructed out of tropes in the way suggested here, then local facts about the tropes of a causal relation will give rise to more general causal truths involving properties.

Despite these advantages and the parsimony of this analysis, however, some may think that there are other, independent considerations which tell in favour of properties, understood as *sui generis* entities or universals. Therefore, although I am keen on this trope nominalist reading, because of its wider metaphysical implications for issues further afield, it is worth seeing what CTP would look like, if we endorsed an ontology of universals. As this view denies that properties are classes of tropes, it leaves us with some form of weak CTP. Understood as a non-reductive criterion of identity, weak CTP states that although properties are something over and above sets of property instances, nevertheless there is a necessary correlation between a property and a set of property instances under the relation, *sameness of causal role*.

Given this reading, how could CTP theorists conceive of the relationship between universals and tropes? We could adopt some version of trope universalism. This view accepts the existence of both *sui generis* tropes and *sui generis* properties or universals. What would CTP trope universalism look like? I suspect that this conception would only make sense if universals are viewed as abstract entities.⁶ These universals, in accordance with weak CTP, would not be exhaustively characterised by the causal roles which their instances realise. Furthermore, universals would not be constituents of property instances, as these instances are *sui generis* entities too. Nevertheless, we can still think of tropes as instances of these abstract universals, because we can suppose that there is a necessary correlation between the identity of a universal and the causal relations its instances realise. Tropes can be thought of as instances of a particular abstract universal, because they reflect its (causal) nature by bestowing a certain set of conditional causal powers onto their objects.

⁶ An explanation of why this is shall soon be offered.

If this account is adopted, then CTP's classes of tropes would mark the extension of universals, they would not themselves be universals. Unfortunately, the resulting conception of the relationship between tropes and universals is less perspicuous than the trope nominalist's. But, still, it does offer an account of what it is to be an instance of one universal rather than another. An object only instantiates the universal of F-ness if it has a trope which displays the causal nature of that universal. So the analysis coheres with weak CTP, as universals can be individuated through the causal contributions of their instances. The account also accords with the further constraints brought to bear upon CTP by the discussion. It meets the fine-grainedness constraint because the local, fine-grained tropes are still doing the causing and effecting. Similarly, if tropes are taken to be the truthmakers of causal statements, the intrinsicity assumption is satisfied. For as tropes are not partly composed of an abstract universal, they can be thought of as intrinsic to their objects/relations. Finally, there is still a strong connection between the behaviour of tropes and wider causal truths regarding universals, as tropes display the causal nature of their universals. So relations between tropes will give rise to more general causal truths about universals, in accordance with the generality assumption.

Would any concrete form of trope universalism be viable? I doubt it. If universals are concrete *sui generis* entities, then they are spatiotemporal entities which can be 'wholly present' in different places at the same time. Given this conception, an account of property instances in terms of instantiations of universals seems unavoidable. For if concrete universals are characterised by the fact that they can be in many different objects simultaneously, then an object which instantiates a universal must partly consist of this spatiotemporal entity. So property instances would have to be understood as complexes of universals, objects and times, as this is what it is to be an instance of a concrete universal. If we take universals to be concrete entities, therefore, we will have to adopt a form of no-trope CTP.

CTP and the no-trope analysis appear an incongruous pair. For if we think that property instances are instantiations of universals, this naturally leads to Moreland's view that "b [the universal] is the nature of a [the property instance]; that is, b answers the question 'what is a an instance of?'" (1985, p.169, brackets added). This analysis is incompatible with CTP, as CTP claims that properties are individuated by the causal

relations that their property instances stand in. So we cannot say that property instances are individuated by their component universals, on pain of circularity. In order to avoid this difficulty, CTP no-trope theorists must say that the nature of a universal is such that two (completely determinate) instantiations of the same universal have to make the same causal contributions to the objects that instantiate them. So there is a necessary correlation between a universal and an equivalence set of property instances (understood as instantiations of universals) under the relation, *sameness of functional role*. No-trope CTP, therefore, is a form of weak CTP.

While CTP theorists needn't endorse an ontology of tropes, however, we've seen that the resulting view cannot respect the further causal considerations brought to bear here. Although it can deal with the fine-grainedness of causes and the generality assumption, it faces difficulty with the intrinsicity assumption and the related local relata principle. In what follows, therefore, I shall set this no-trope view aside, and just concentrate upon the other conceptions of the relationship between properties and their instances.

6.2 Determinables and Determinates

How can the basic proposal outlined above, account for the relationship between determinables and their determinates? The distinctive issues raised by the determinable/determinate relationship, are frequently overlooked by theories of properties. A satisfactory analysis of properties, however, must consider this issue. For even if we ultimately decide to take a sceptical attitude towards the existence of determinable properties, we still need to offer some explanation of why the entities referred to by determinable predicates stand in a special relationship to those picked out by certain determinate predicates. What are these special features of the determinable/determinate relationship? We can distinguish four key characteristics of this relationship, which a theory of determinables and determinates needs to account for:⁷

1. Determinates are more specific than their determinables.

More precisely we can say,

⁷ This list draws upon features cited by Johnson (1921), Prior (1949), Searle (1959) and Armstrong (1997).

- (a) Having a determinable entails having some determinate of that determinable, but no particular determinate property is entailed.
 - (b) Having a determinate entails having the determinable property which the determinate falls under.
2. The determinable/determinate relation differs from the genus/species relation because we cannot define a determinate by means of a determinable plus some independent differentia. Consider, for instance, the determinable colour and its determinates red and blue, it seems clear that “we cannot give verbal expression to the differentiating element except by using the species name itself, red or blue”.⁸
 3. If a particular instantiates a determinate property, it cannot instantiate another determinate property, which is on the same level and falls under the same determinable, at the same time and place. Johnson (1921), whose classic discussion first introduced the terminology of determinates and determinables, writes, “if any determinate adjective characterises a given substantive, then it is impossible that any other determinate under the same determinable should characterise the same substantive” (p.181).
 4. There are “ultimate differences” between the highest of the determinables which render them incomparable.⁹ But the determinates which fall under one determinable are all comparable with each other. Moreover, they resemble each other to different degrees. So we can say that yellow is more like orange than blue is, and so on.

In this section, I shall consider how the proposal outlined above can account for the special features of the determinable/determinate relationship.

Trope nominalists often deal with the phenomenon of determinables by appealing to less than exact relations of resemblance. Williams, for instance, writes, “the tropes approximately similar to the given one provide a less definite universal” (1953, p.81). Campbell concurs arguing, “The closeness of resemblance between the tropes in a set can vary. These variations correspond to the different degrees to which properties are specific” (1981, p.134). This sort of approach is unsatisfactory,

⁸ Prior 1949, p.5.

⁹ See Johnson, 1921 p.175. Colour, shape and pitch are said to be amongst the highest determinables (i.e. these properties are not determinates of anything else). They are distinguished from a family of determinables and its determinates by their complete “otherness” (p.176).

however, largely due to its vagueness. The appeal to approximate degrees of similarity results in an imprecise conception of what the class of red tropes, for instance, consists in. Another cause for concern is that it looks like trope theorists will have to substantially increase the number of primitive relations postulated by their theory. For now we not only have the relation of precise similarity, but also the relation of similarity to degree n_{th} . Bacon argues that “The cost is high”,¹⁰ and so he is only prepared to countenance the existence of a single relation of exact similarity. But how else can a trope nominalist account for determinate and determinable properties? Can CTP trope nominalism provide us with an account of this relationship?

The completely determinate properties (i.e. those properties which are not the determinables of any other properties) are easy to deal with. These are the classes of tropes that are precisely similar to each other. Given CTP, this amounts to each trope in a completely determinate class bestowing upon its object the exact same cluster of conditional causal powers as the other members of its class. Thus, if trope₁ and trope₂ are part of a class which stands for a completely determinate universal, then every causal power that trope₁ bestows upon its object, trope₂ also bestows upon its object, and vice versa. Matters are more complicated, however, when we get to the determinables. These, as Johnson writes, characterise their particulars “less determinatively”.¹¹ So objects that instantiate a determinable property permit of more variation than their completely determinate counterparts. Two tropes of redness, for instance, might not bestow exactly the same cluster of causal powers onto their objects. For one may be an instance of crimson and so have the power to appear crimson to us, while the other may be an instance of scarlet and so have the power to appear scarlet to us.

How can we reflect this in the account? If we are trope nominalists, then the most natural way is to treat determinables as constituted out of classes of their determinates. This proposal was originally suggested by Stout. He writes,

‘redness’ is a subclass of the more general class ‘colour’ as ‘red things’ is a subclass of ‘coloured things’ (1921, p.398).

¹⁰ See Bacon, 1995, p.18.

¹¹ Johnson 1921, p.174,

On this view, then, the class of tropes which is the determinable of colour, is the class of tropes which includes all the tropes of its determinates. So the class of yellow tropes, blue tropes, red tropes, etc. are all subclasses of the class of colour tropes. Such a class of tropes forms a much looser cluster of causal powers than the subclasses of yellow, blue, red etc. Thus, two tropes of colour would not necessarily bestow exactly the same cluster of causal powers onto their objects.

This suggestion is promising as it satisfies the first criteria of adequacy for an account of determinables/determinates. Any trope of colour has to be a determinate shade of colour, since the class of colour tropes consists entirely in tropes of more determinate shades. But because a colour trope can be one of a number of different determinates, a trope of colour does not have to be any particular determinate shade. Similarly, every trope of red is also a trope of colour, because by being a member of the class of red tropes, it is thereby included in the class of colour tropes. This account, moreover, is not limited to trope nominalists, other CTP theorists could also utilise it. CTP trope universalists, for instance, could say that determinables are universals that are individuated by the causal relations that instances of the determinable's determinates realise. Or, they could deny the existence of determinable universals and say that when we speak of redness, what we are referring to are classes of instances of more determinate universals.

Despite the flexibility of this proposal, however, as it stands it is inadequate. For although it seems plausible to say that the class of red tropes includes all the tropes of its determinates, as yet we have no idea of what this red class consists, because we don't know what it is to be a determinate of red. Why are tropes of scarlet, crimson, magenta etc. members of the class of red tropes, and not those of scarlet, lime green and turquoise? We still require some conception of what being a member of the class of red tropes involves.

CTP's framework suggests a way of responding to this. We can keep hold of the suggestion that the class of red tropes contains all of the tropes of its determinates, and then supplement it with an account of how determinables are formed from classes of determinates. How might the account go? Shoemaker has recently offered a suggestion which draws upon his CTP account. He writes,

sometimes the conditional powers bestowed by one property will be a proper subclass of those bestowed by another. This will be true where the one property is a determinable of which the other is a determinate. The class of conditional powers bestowed by redness will be a proper subclass of the conditional powers bestowed by scarlet, for example. The different determinates of redness will each confer its distinctive class of conditional powers – but these will have in common the class of conditional powers conferred by redness (1998b, p.78).

So CTP theorists could say that determinates are sorted into more or less similar groups by the number of causal powers they share. Property F-ness and G-ness, for instance, are determinates of the same determinable iff they share a subclass of their causal powers. The more causal powers two or more determinates share, the more unified the resulting determinable will be. The determinates of red, for instance, will have more of their causal powers in common than the determinates of colour. Consequently, red is a more unified (and so more natural) property than colour is.

Tropes of determinates, then, are members of classes which stand for determinables because of their shared causal powers. What does the resulting account look like? Imagine that scarlet has causal powers 1-18, crimson has causal powers 3-15, magenta has causal powers 5-25, and so on. According to the present suggestion, because we can extract from these determinates a common core of causal powers, namely, causal powers 5-15, this class of determinates forms a determinable. Why say that this determinable is red, rather than blue or green? Because according to the RCL definition for the predicate ‘is red’, the word red is associated with an entity which has causal powers 5-15, whereas the predicates ‘is blue’ and ‘is green’ are associated with different clusters of causal powers. Thus, as these determinates with overlapping causal powers most resemble what we name red, they form the determinable of redness.

This account offers a clear conception of what membership of a class which stands for a determinable involves. But, unfortunately, it is not without its difficulties. One worry is that it seems to lead to a proliferation of determinables, as all those determinates with a shared subclass of causal powers can be said to be a determinable. Imagine, for instance, that the class of {crimson, scarlet} shares a slightly larger subclass of causal powers than {crimson, scarlet, magenta}. In this case, given what has been said, we can say that the class of {crimson, scarlet} forms one determinable and {crimson, scarlet, magenta} forms another. I suspect that this possibility just has

to be accepted, but this doesn't mean that sensible distinctions between more or less natural determinables can't be made. Naturalness can be viewed as a trade-off between the number of shared causal powers and the number of determinates included in the determinable class. Why might the latter be important? Because properties are meant to signify important resemblances between objects. But if one object instantiates the determinable {crimson, scarlet} and another the determinable {crimson, scarlet, magenta}, which has extremely similar causal powers to that of {crimson, scarlet}, this will not be a substantial difference between those objects. So if we want properties to carve nature at its joints, we will look for determinables which mark greater differences in the objects that instantiate them.

A potentially more serious source of disquiet is caused by the assumption that there will always be a proper subclass of causal powers which determinables share with their determinates. As I set it up, determinates are grouped into possible determinables on the basis of shared causal powers. But perhaps there are some determinates of determinables which do not have any causal powers in common. I accept that this may represent a possibility and thus the suggestion is open to empirical refutation. Nevertheless, I think that the account outlined is very credible. If the world is carved up into completely determinate properties, all of whose tropes have exactly the same causal powers, the fine-grainedness of these entities makes it seem highly likely that there will be shared subclasses of causal powers between different determinates. From these first-level determinables (i.e. determinables of completely determinate properties) further shared subclasses of causal powers will probably emerge, which in turn will create yet more determinables, to which the first-level determinables will stand as determinates.¹²

The real test for the account, however, is whether it can satisfy the criteria of adequacy proposed for an analysis of the determinable/determinate relation. Unfortunately, the suggestion seems to fall at the first hurdle, (1a). The account does make determinates more specific than their determinables, because a determinate has to meet narrower, more exacting conditions in order to be instantiated by an object.

¹² In this respect, I think the account fares better than Armstrong's (1978, 1997). His theory commits us to the contentious claim that all determinates are complex properties. As a class of determinates which fall under a determinable is unified by relations of partial identity, and two properties can only stand in the relation of partial identity if they share some common constituent, i.e. a universal.

But as things stand, there is no guarantee that an object has to instantiate a determinable by instantiating a particular determinate of that determinable. Why not? The shared causal powers of red's determinates are claimed to be essential to the determinable of redness. A trope is a trope of redness, on this view, because it can realise a common nexus of causal relations. Nothing has been said, however, to rule out the possibility of tropes *just* realising this common core. So there could be tropes of red per se, i.e. tropes of red which aren't also tropes of a particular shade of red.

As this possibility should be excluded, the suggestion needs augmenting further. The class of causal powers which an object has to instantiate in order to be red (in the example given above, causal powers 5-15) should be regarded as necessarily incomplete.¹³ In other words, nothing can instantiate causal powers 5-15 without thereby instantiating some further causal powers, because causal powers 5-15 nomologically entail some further causal powers. This then allows us to say that when an object is red (in virtue of instantiating causal powers 5-15), it will thereby be a particular shade of red. Because in order to instantiate causal powers 5-15, the object will have to instantiate some other causal powers definitive of the determinate shades of red. So an object could not instantiate the property of redness per se, as (1a) claims.¹⁴

This proposal satisfies (1b), because if we think of determinables as classes which include determinates as subclasses, a trope in a determinate class will automatically be part of any determinable which has that determinate as a subclass. We can also explain why the determinable/determinate relation differs from the genus/species relation. Before it would have been possible to view a determinate as a combination of the determinable plus some differentia. Scarlet, for instance, could have been analysed as those tropes which ground causal powers 5-15 (i.e. the red tropes) and also 1-4, 16-18 (i.e. the differentia). But given the extended analysis, we cannot separate the causal powers of the determinates into those which are shared with

¹³ The necessity in question can be regarded as either nomological or metaphysical. However, if we accept that CTP's criterion of identity is a transworld identity thesis, as most CTP theorists do, then nomological necessity will just be a form of metaphysical necessity (see chapter seven).

¹⁴ Although if we accept Yablo's account, which I shall outline shortly, this claim will need qualifying slightly. For although red tropes will always be coincident with determinate tropes of red, they will nevertheless be distinguishable from them (see §6.3). The account outlined here will still be required, however, in order to explain why red tropes have to be coincident with their determinates.

the determinable and the rest. So the determinable/determinate relationship cannot be treated in the same way as the genus/species relationship.

Why do determinates exclude each other? Although the account supposes that the causal powers of determinates overlap to some extent, they cannot completely overlap as then, according to CTP, they would be the same determinate property. So a trope of scarlet, for instance, cannot also be a trope of crimson, as they have to bestow different causal powers onto their objects. But why couldn't an object possess both a trope of scarlet and a trope of crimson at the same time and place? The answer to this must lie in the fact that some of scarlet's causal powers exclude some of crimson's causal powers. One of the causal powers that a trope of scarlet must bestow, for instance, is that of making an object look scarlet to us. This causal power cannot be instantiated alongside crimson's class of causal powers, as its causal powers include that of making objects appear crimson to us.

We can explain why determinates exclude each other, therefore, by saying that the different classes of causal powers associated with different determinates cannot be instantiated together. Some may be concerned that this proposal doesn't guarantee that determinates will exclude each other, as it may be possible to have determinates whose classes of causal powers can be instantiated together. I accept that this is still a possibility, but this maybe because objects could instantiate determinates of the same level. Although they standardly don't, Armstrong asks, "is it impossible for the very same thing to be simultaneously sweet and sour? Sounds appear to be so ordered. But are they incompatible with each other in the way that shapes and colours are?" (1978, p.113). Armstrong argues that these questions need to be decided empirically not a priori, and I think that this is right.¹⁵ So I doubt that it is a failing of the account that it allows for the possibility of determinates not excluding each other.

Finally, the account does justice to Johnson's observation that the determinates of the same determinables are comparable, unlike the completely determinable properties. The fact that some properties are not comparable with others can be

¹⁵ Although we can weaken this slightly and say that, sometimes, it needs to be decided empirically, other times it may be deducible a priori. If, for instance, it was part of our concept of red and green that their causal powers excluded them from being instantiated at the same time and place, say because one was defined via certain appearances and the other by certain contrasting appearances, then it would be an a priori matter that objects couldn't be red or green all over.

explained by the fact that they share no overlapping causal powers whatsoever. They thus appear completely ‘other’. Determinates of the same determinables, on the other hand, can be compared as they all share some subclass of causal powers. Furthermore, they can be more or less similar to each other, depending on how many causal powers they have in common. Those determinates which are like each other will have a larger class of overlapping causal powers than those determinates which aren’t.

If we are prepared to accept the augmented proposal, therefore, CTP offers a plausible account of relationship between determinates and their determinables which satisfies the criteria of adequacy. Before moving on from this topic, however, I first want to consider how this account bears upon a related debate, that has arisen between MacDonald and Yablo.¹⁶

6.3 MacDonald versus Yablo

MacDonald argues for the thesis that “different properties may share a single instance” (1989, p.161). In order to defend her claim, she asks us to consider the relationship between the determinate property of being red and its determinable of being coloured. She writes,

No one would suppose that in order for an object to possess both properties, it must first instance the former property, and then, in addition, instance the second. An object’s instancing of the former property just is its instancing of the latter: nothing further is required...But if this is so, then any case in which an instance of the property of being red is causally efficacious is one in which an instance of the property of being coloured is also causally efficacious (1995, p.65).

Here, then, MacDonald is claiming that it is plausible to suppose that a single property instance can be both an instantiation of the property of colour and an instantiation of the property of redness.

This seems an attractive position to adopt, at least in cases of determinates and determinables, as it respects the thought that an object instantiates redness by instantiating some specific shade of red. If we say that one particular trope is an instance of both the property of scarlet and the property of redness, we can hold onto

¹⁶ See MacDonald (1989 and 1995) and Yablo (1987, 1992a and 1992b). Robb (1997) and Ehring (1999) also endorse MacDonald’s thesis.

the claim that, in this instance, instantiating red just is instantiating scarlet. Moreover, we can do this while still maintaining that there are two distinct properties here, since the property of redness and scarlet have different extensions. Is MacDonald's thesis compatible with the account of determinables and determinates outlined above? Yes, one trope can be both an instance of a determinate and an instance of a determinable, because tropes can be members of more than one class. My boots, for instance, possess a trope which is both a member of the class of red tropes and a member of the class of scarlet tropes. When this occurs, we can say that an object possessing this trope, instantiates all of the properties (i.e. classes of tropes) this trope is a member of, in virtue of having this single trope.

Yablo, however, objects to MacDonald's thesis, arguing that properties are not co-instantiable in a single instance. Yablo notes, I think correctly, that there are more problematic cases which count against this thesis.¹⁷ Consider, for instance, an example which is put forward by MacDonald herself. She asks us to suppose that there is a piece of putty resting on a metal mesh. Over time, there is a change in the putty's shape, due to a change in its microphysical parts, which causes the putty to fall through the mesh. During that same time, the microphysical parts are also responsible for the expansion in the volume of the putty. MacDonald comments,

On the co-instantiation model, it looks as though the change in shape and the expansion in volume will be co-instanced, so that if one is causally efficacious, then so is the other. But if this is so, then we are forced to the conclusion that the expansion of the putty must be held to be causally responsible for the putty's falling through the mesh (1995, p.67).

This seems an unwelcome conclusion. MacDonald tries to deal with it by invoking Davidson's distinction between citing the cause and providing a good causal explanation.¹⁸ But she is still forced to deny the intuitive claim that it was the putty's change in shape, not its expansion in volume, which was the cause of the putty's falling through the mesh. So it is important to ask whether CTP commits us to MacDonald's conclusion.

¹⁷ See Yablo 1992b, p.259. Shoemaker (1998b, p.80-1) also objects to this view.

¹⁸ See §4.2.

I don't think that it does. Indeed, it seems to commit us to rejecting this conclusion. Why? The causal profile which characterises the property of volume X (X being the volume of the putty at time t) and the property of having shape S (S being the shape of the putty at time t) are, we can suppose, different. For we want to say that the latter but not the former is responsible for the putty's falling through the mesh. So two tropes of these properties have to be different too, as they bestow these different causal profiles onto their objects. We thus have the resources to say that it was the changing shape of the putty, not its increased volume, which was responsible for the object's falling through the mesh. Why does this case differ from ones involving determinates and determinables? After all, in these instances, the conditional powers which characterise the determinable also differ from the ones that characterise its determinates. There is an important difference between these two cases, however, because in the case of red and scarlet, the causal powers of red are a subclass of the causal powers of scarlet. In other words, the causal powers of red completely overlap with the causal powers of scarlet. In the putty scenario, by contrast, although there is an area of overlap between the causal powers of the shape of the putty and the causal powers of the increased volume of the putty, they do not completely coincide with each other. This means that we cannot say that the trope of the putty's shape is a member of both the class which stands for the property of being shape S and the class which stands for the property of being volume X, because two tropes are required to realise the different causal profiles definitive of each property.

CTP theorists should thus reject MacDonald's description of the putty case. For on their view, co-instantiation will only take place when the causal profiles of two properties overlap completely. This is a welcome conclusion, however, as it allows CTP theorists to say that two properties can be co-instantiated if those properties are determinables and their determinates, but it doesn't allow us to say this in cases such as the putty one, where the thesis seems far less plausible. If we adopt CTP, therefore, we can embrace this intuitive analysis of property instantiation in cases of determinables/determinates, without thereby being committed to this thesis in cases where the analysis seems inappropriate. So at least in this respect, CTP's analysis of the relationship between properties and their instances surpasses that offered by MacDonald.

Yablo takes issue with MacDonald's claim that properties can be co-instantiated. He argues that two properties cannot be instantiated in a single instance, even in cases involving determinables and their determinates. Why does he want to say this? Because he thinks that, sometimes, we need to distinguish between cases in which determinables are the causes and cases in which their determinates are the causes. Consequently, if we identify instances of red with instances of scarlet, we will be unable to make these fine-grained distinctions. Consider this example given by Yablo:

Imagine a pigeon, Sophie, conditioned to peck at red to the exclusion of other colours; a red triangle is presented, and Sophie pecks. Most people would say that the redness was causally relevant to her pecking, even that it was a paradigm case of causal relevance. But wait! I forgot to mention that the triangle is a specific shade of red: scarlet (1992b, p.257).

This example is utilised by Yablo to show that the causal exclusion principle is overdrawn. For if it were true that "if an event is causally sufficient for an event *y*, then no event distinct from *y* is causally relevant to *y*",¹⁹ we would have to conclude that the triangle's being red is causally irrelevant to Sophie's pecking, as the fact that the triangle was scarlet was causally sufficient for the effect. What we require, then, is a restatement of the causal exclusion principle, which respects the claim that both red and scarlet are causally relevant in this case. So far so good. MacDonald also wants to claim that determinables are causally efficacious when their determinates are. So we can all agree that, in this case, the triangle's being red is the cause of Sophie's pecking. Yablo goes further than this, however, arguing that the triangle's being red, but *not* the triangle's being scarlet, should be taken to be the cause of Sophie's pecking.

This is where MacDonald and Yablo part ways. She wants to claim that both the triangle's being red and its being scarlet are the cause of Sophie's pecking; since they are the very same thing they must both be the cause. But Yablo believes that it is mistaken to think of the triangle's being scarlet as the cause, because this is not proportionate to its effect. Sophie would have pecked if the triangle had been red but not scarlet (if it has been crimson, for instance). It was the redness of the triangle that

¹⁹ Yablo 1992b, p.247.

'made the difference' between the effect's occurring and it not occurring, so this was the cause of Sophie's pecking. Unfortunately, this claim seems to stir up problems for the proposed account of determinables and determinates. For if Sophie's pecking was caused by the triangle's being red and not also by the triangle's being scarlet, it looks like red has a causal power that scarlet doesn't have, namely, the causal power to make Sophie peck. This causes difficulties for the account because it states that the condition of entry into the red class is that the tropes realise a *subclass* of the determinates causal powers. So if there are causal powers had by red but not by red's determinates, this condition of entry is false.

This leaves us with two options: either we can reject Yablo's claim, or we can try to accommodate it within the account proposed. There is plenty of scope for the first method of attack. Yablo's argument rests heavily upon his proportionality principle, which states that,

Nothing causes an effect that leaves out too many relevant factors, or brings in too many irrelevant ones. True causes are...commensurate with their effects (1992a, p.404).

As a general rule of thumb, this seems to be a good one. But I shall argue that we can accept the general soundness of the proportionality principle, while still having reason to deny its applicability in Sophie's case.

There seem to be important differences in cases which could be classified as having 'too much detail'. Consider, for instance, an example given by Ducasse. He counter-intuitively claims that if, at the instant a brick strikes a window, the air waves of a canary song reach the window, the cause of the window's smashing will be the canary's song as well as the brick.²⁰ In this case, Yablo's claim that the cause includes too much detail seems highly plausible – we should say that the brick's striking is alone the cause of the window's breaking, as this is all that is required to bring about its destruction. This case, however, seems quite different from Sophie's. First, in the canary case, if we say that the cause is just the brick's striking, we thereby fix on something which is more proportional to the effect. But it is not clear that we can do this in Sophie's case. For we would simply be exchanging one determinate of red for

²⁰ See Ducasse, 1926, p.133.

another, which isn't anymore proportional to its effect. Second, we can make the cause more proportional to its effect in the canary case, without changing the thing that seems to be doing the causing, because in both we still have the same brick striking the window. But in Sophie's case, it looks like we will change the cause, because the cause will go from being scarlet, to being crimson, magenta, or some other determinate of red.

Yablo, however, has a response to this. He argues that it is an instance of the property of redness, not any of red's determinates, which is the cause in this case. This instance of redness is more proportional to its effect than any of red's determinates. For it wouldn't have mattered if the triangle was scarlet, crimson etc, so long as it was red the effect would still have occurred. Some may feel that this reply misses the point. According to the criteria of adequacy for determinates and determinables, every instance of the determinable red must also be some determinate shade of red. So as we want to say that it is the *instantiated* red of this triangle which is doing the causing, the singular cause of the pecking will have to be red *and* either scarlet, or crimson, or magenta, etc. This, however, begs the question against Yablo, as it presupposes an ontology which doesn't distinguish between an instantiation of red from that of scarlet. Yablo allows that whenever there is an instance of scarlet, crimson, etc. there is also a distinct instance of redness. So we are able to say that it is this instantiation of redness, not that of scarlet, which is more proportionate to its effect and hence the cause in this case.

It looks as if we have a stand off, then. If we have an ontology which differentiates between an instantiation of red and scarlet, an instantiation of redness will be more proportionate to the effect. But if we have an ontology like MacDonald's, then we won't be able to make such a distinction. This conclusion, however, is enough to seriously damage Yablo's argument. Why? Yablo implicitly assumes his preferred ontology when utilising the proportionality principle. For in order to be able to say that, in Sophie's case, redness is more proportional to the effect than scarlet, we have to suppose that instantiations of scarlet are distinct from instantiations of redness. But this, of course, is the point at issue. If we endorse MacDonald's ontology, we will deny that there is this distinction between an instantiation of red and that of scarlet. So

we will have good reason for claiming that the proportionality principle cannot be applied to determinable/determinate cases in the way that Yablo suggests.

Yablo may object: the proportionality principle is an argument for this ontology. Since, with it, we get causes which are more proportional to their effects. This deployment of the proportionality principle is inadmissible, however. While this principle is a good heuristic device for seeking the cause among a number of possible alternatives, this is only the case when these possible alternatives are conceived of as *distinct existences*. If we do not think that the entities in question are distinct existences, as in the case of co-instantiated property instances, the principle cannot be used to show that one and the same thing isn't the cause. For, ex hypothesi, both are the same, so they are going to be equally proportional to their effect.

It may be objected that I'm employing double standards here. Davidson doesn't allow that an event like *Don's gripping* and a property instance of an event like *Don's gripping lightly* are distinct existences. But didn't I employ something very like the proportionality principle against him in chapter four? No, I never used the proportionality principle as an argument for my preferred ontology. I only said that if you are *not* a nominalist then, using the proportionality principle, there is reason to think that property instances can be causes and effects. Yablo, however, does use the proportionality principle to support his ontological analysis, and it is this which I'm claiming is inadmissible. Yablo can, of course, try to establish his ontological analysis by other means. Perhaps, for instance, he could argue, as I did against Davidson, that his ontology makes for a better account of causal explanation. But he does not appeal to any such considerations, and it is doubtful whether such a case would be compelling in the case of determinables/determinates.²¹

Yablo, therefore, fails to undermine MacDonald's claim that determinables and determinates are co-instantiated. However, if CTP theorists do want to side with Yablo, there is a way of accommodating his views within the proposed account. How? First, we need to show that the account can distinguish between tropes of red and tropes of red's determinates. This can be done by appealing to Yablo's relation of coincidence. We can say that the red trope that the triangle displays is coincident with,

²¹ Although, as I argued against MacDonald, I do think that these considerations will prove persuasive if we use the co-instantiation model more widely than this.

but not identical to, the scarlet trope that the triangle displays. The triangle is red because it displays causal powers 5-15, but it is also scarlet because it displays causal powers 5-15 in virtue of displaying causal powers 1-18. Now we could say that these two properties are both instantiated by the same trope of scarlet in this instance, like MacDonald does. But we could also say that these two tropes are distinct because they have different modal features. The trope of red has causal powers 5-15 essentially but 1-4 and 16-18 only accidentally, whereas the trope of scarlet has all of the causal powers 1-18 essentially. So we can say that the two tropes are coincident, because all of their causal powers are the same in this instance. However, they are not identical, because they possess their causal powers differently.²²

How does this help? One of Yablo's other central theses is that the manner in which something is possessed, i.e. whether an entity has the item essentially or accidentally, affects what it can cause.²³ This claim, combined with the coincident version of the determinables account, gives us room to manoeuvre. For we can use the difference in the manner in which an object instantiates a certain class of causal powers, to ground the difference between what objects with red, as opposed to scarlet, tropes can cause. So in Sophie's case, we can suppose that it is the trope which possesses causal powers 5-15 essentially and 1-4 and 16-18 accidentally which causes Sophie to peck. For in nearby possible worlds, the triangle has causal powers 5-15, but not 1-4 or 16-18.

If we take this line, how would the proposal now go? We can still claim that the class of red tropes is {crimson tropes_{1-n}, scarlet tropes_{1-n}, magenta tropes_{1-n} etc.}.²⁴ And we can still say that in order to be a trope of red, a trope must realise a subclass of the causal powers of red's determinates. Let's suppose that this subclass is causal powers 5-15, so every red trope has causal powers 5-15 essentially. A red trope is thus

²² Yablo writes, "x is coincident with y iff they have their categorical properties in common" (1992a, p.408). I cannot talk about properties in the above context, but something similar is meant since which causal powers are essential or accidental depends on how things stand in possible worlds other than the actual one. In the actual world, we are supposing that there is nothing to distinguish the causal powers of red and scarlet tropes, thus, it seems appropriate to call them coincident.

²³ See, in particular, Yablo 1992a.

²⁴ The story is basically the same if we adopt trope nominalism or trope universalism. As I noted earlier, if we accept trope nominalism, then the set of red tropes is the universal. This would also be the case if we were trope universalists who didn't want to endorse the existence of determinable universals. However, if we did believe in determinable universals, then this set of tropes, rather than being the universal, would instead only reflect its nature by giving its identity conditions.

distinct from a scarlet, crimson etc. trope, in having some of its causal powers accidentally and some essentially. This can make a difference to what red objects can cause, as we can suppose that since causes must be proportional to their effects, sometimes, a trope which only has causal powers 5-15 essentially and all the others accidentally will be more proportional to the effect than those which have all 1-18 essentially. Does this mean that the condition that tropes have to meet in order to belong to the red class is no longer a subclass of its determinate's causal powers, contrary to what is claimed by CTP? I don't think that this follows. We can still say that tropes belong to the red class if they share a certain subclass of causal powers. Furthermore, this subclass of causal powers exhausts what it is to be the property of redness. All we are saying now is that because an object possesses some of its causal powers in a different manner by having a trope of red as well as a coincident trope of scarlet or crimson etc, an object's being red is more proportionate to some effects than an object's being scarlet. This is not to say that the triangle's being scarlet doesn't suffice for the effect - clearly it does - Yablo does not dispute this. Hence, there is a sense in which the triangle's being scarlet has the causal power to bring about the effect. It is just denied the title of cause because of the extra constraints on this notion.

Thus, I think that we can remain faithful to the essence of the proposal outlined above, even if we accept Yablo's claim that the triangles being red, not its being scarlet, is the cause of Sophie's pecking. This is an agreeable conclusion, as it means that the account of determinables and determinates proposed here can remain neutral in the debate between MacDonald and Yablo. We can say that there are some properties which can be co-instantiated in one property instance - although there are more restrictions on which properties can be co-instantiated than MacDonald allows. But, equally, we can deny this claim, as Yablo does, even in cases involving determinables and determinates.

6.4 The Nature of Tropes

If the proposal outlined here is to be believed, tropes play a key role in CTP and the resulting account of causation. What, then, are these things called tropes? The myriad of titles they go by has somewhat obscured their prevalence in philosophy. Tropes are also referred to as 'modes', 'individual accidents', 'moments' 'abstract particulars'

and even ‘events’.²⁵ Although there are differences in the precise nature of the entities postulated, most of these can be put aside here. We can, for instance, remain neutral on the question of whether tropes are persisting entities, as Ehring argues, or more momentary, event-like entities, like others have claimed.²⁶ Similarly, we can remain open upon whether tropes have to be qualitatively simple, or whether their qualitative natures can be extremely complex.²⁷ What does matter for the purposes of this discussion, however, is that tropes are entities which have these characteristics:

- (1) They are instances of more general attributes found in the world.
- (2) They are *sui generis* entities, not analysable in terms of particulars and universals or as members of classes of possible particulars.
- (3) They have a single spatiotemporal location.
- (4) They are causally efficacious entities.
- (5) Non-relational tropes are intrinsic to the entities that possess them.

Any entity, whatever its name, which meets this specification is what I call a trope, since that entity can satisfy the role that has been sketched for tropes here.

Tropes have figured so crucially in the account of properties and causation outlined, because of the desire to respect the grounding intuition and the related intrinsicality assumption. Nothing which has been said, however, commits us to a view about the position of tropes in the wider metaphysical picture. We could take them to be the most fundamental of the ontological categories, out of which objects,

²⁵ For talk of ‘modes’ see Lowe (1998). He attributes this word to the scholastics. Also Locke’s use of ‘mode’ seems to refer to something very like a trope (see, for example, 1975, §III, vi. 42). Leibniz talks about ‘individual accidents’ (1981, p.38). Husserl and his followers call them ‘moments’. Campbell (1981) names them ‘abstract particulars’, while Quinton (1979) and Hausman (1998) refer to tropes as events.

²⁶ See Ehring, 1997 Ch. 4. Others have argued that tropes have a more event-like nature. Quinton, for instance, writes, “It occurs or happens, rather than exists” (1979, p.211). This leads us to view tropes as momentary entities.

²⁷ Quinton takes the first view (1979, p.211), while Bennett the latter (1988, p.92-3). Bennett also attributes his view to Leibniz, because included in Leibniz’s list of individual accidents is the birth of Jesus Christ (1981, p.328), which Bennett takes to be a complex trope. In order to respect the fine-grainedness of the causal relata, tropes have to be qualitatively simple at least sometimes. But I doubt any trope theorist would deny this, all that is at issue is how complex an entity a trope could be.

universals and events are built.²⁸ But we could also deny all of these claims. Tropes are not intended to analyse the nature of events or objects, nor are they said to offer a satisfactory answer to the problem of resemblance. Instead, I have argued that these entities should be part of our ontology because, with them, we can respect the intuition that the locus of an object's causal powers stems from the intrinsic nature of an object. So causal relations involving objects are causally related in virtue of local, intrinsic features of that relation.

Given this objective, it is clear that in the list of trope characteristics, the last is extremely important. In order to satisfy the role set for them here, non-relational tropes must be intrinsic to the particulars that possess them. If they're not, then I've been barking up the wrong tree. Earlier I suggested that tropes did fulfil this criterion,²⁹ but the case given was largely negative, as I argued that nothing impeded them from being regarded in this way. We can say more than this, however. For if we look at the two most popular and well-worked out accounts of the nature of tropes, we see that they render tropes intrinsic to the particulars that possess them.

The first of these approaches is put forward by Stout (1921), Williams (1953), Campbell (1990) and Bacon (1995). They believe that tropes are something akin to mini-objects, as they literally make-up or constitute objects. Campbell, for instance, writes,

On the view that tropes are the basic particulars, concrete particulars, the whole man and the whole piece of cloth, count as dependent realities. They are collections of co-located tropes, depending on these tropes as a fleet does upon its component ships (1981, p.128).

On our standard, commonsensical conception, objects are taken to be self-standing entities which properties are dependent upon. But this account states that the converse is the case. Concrete particulars are dependent entities, as they cannot exist apart from the tropes that form them. Whereas tropes are the self-standing entities, which are not reliant upon anything else. Tropes are thus viewed as the basic building blocks of the universe. As a matter of fact they come together in clusters and form more complex,

²⁸ Campbell calls it "a one-category ontology" (1990, p.1). For philosophers who try to make out this claim, see, for instance, Williams (1953), Campbell (1990), and Bacon (1995), who also adds possible worlds to the list.

²⁹ See §5.6.

ordinary particulars such as trees and apples. But they are independent parts of the objects they compose.

The second approach is pretty much the polar opposite of this. Instead of taking tropes to be self-standing entities, separable from the objects of which they're a part, this approach views tropes as wholly dependent upon the objects that possess them. Husserl's discussion of parts and wholes presents this view of tropes (or "moments") as entities which are essentially dependent upon their wholes.³⁰ He distinguishes between two different ways an entity can be "part" of a whole. An entity can be part of a whole by being, what Husserl calls a "piece" of that whole. A page, for instance, is a piece of a book, or a table leg is a piece of a table. What characterises them is that they are detachable (at least in principle) from the whole of which they are part. They can exist separately from that whole, and thus can be thought of as wholes themselves. The second way an entity can be a part of something is by being a dependent part of that whole. A whole's colour or extension, for instance, cannot exist independently of that whole. We cannot just take away an object's colour and hold it in our hand, like we can the leg of a chair. Similarly, we cannot simply wrench apart the colour of an object from, say, its shape. This is why Husserl refers to moments as 'interpenetrating' (vol. II p.4). They are not only necessarily tied to their wholes, they also depend (either reciprocally or unilaterally) on many other moments in the whole. Thus, the moments inter-link and permeate each other.

We have, then, two very different conceptions of what tropes are like, but both render tropes intrinsic to their objects, events or relations.³¹ Earlier, an intrinsic entity was characterised as one which is internal to its particular, so that particular can possess it regardless of what is going on outside of it. The Stout *et al.* account fits this characterisation, because tropes stand in the traditional part/whole relationship with their particulars. These tropes are what Husserl would call 'pieces', because they are self-standing entities which can exist away from the bundle of which they're part. But tropes are still internal to, or in, the particulars they comprise, like the pip of an orange

³⁰ See, for example, 2001 vol. II, p.4. Some followers of Husserl have also adopted this conception of a moment or trope. See, for instance, Sokolowski (1968), Smith (1982), Mulligan, Smith and Simons (1984) and Simons (1987).

³¹ For simplicity's sake, I shall talk of particulars rather than objects, events or relations. But all of these things are implied.

is internal to, or in, its orange. This means that tropes do not require the existence of any entity outside the boundaries of its particular. Just like the pip of an orange can exist independently of entities external to its orange, tropes too are not reliant upon entities extraneous to their particulars.

Husserl's characterisation also renders tropes intrinsic to or immanent in their particulars. However, instead of this being guaranteed by the ontological priority of tropes over particulars, almost the converse is the case – tropes are intrinsic to the particulars that possess them, because they are wholly dependent upon them. A trope is possessed by its particular (or in Husserl's language 'founded' in its particular) because without it, that trope would not exist. Thus, a trope is not internal to a particular like a pip of an orange is, rather a trope is internal to its particular because it depends upon that particular (including some of its other tropes), but on nothing outside of it. Both of these analyses, therefore, render tropes intrinsic to the particulars that possess them. For tropes are wholly present within the confines of their particulars, so the existence of anything extraneous is not entailed.

As both Stout's tropes and Husserl's moments satisfy the criterion outlined, these analyses underline the point that this criterion leaves the nature of property instances underdetermined. While no doubt significant restrictions have been placed upon the nature of these entities, whether or not we think of them as events (*a la* Quinton), or as self-standing simple entities (*a la* Campbell), or as dependent moments (*a la* Husserl) doesn't matter. All that is required is that the entities satisfy the five features outlined above, since then they will be able to fulfil the role sketched for tropes in this account of properties and causation.

This result, in a sense, is rather unsatisfactory, as it leaves many important issues unresolved. Consider, for instance, the question mentioned earlier: what is the relationship between a trope and its object? It is clear that the proposed CTP does not prescribe an answer to this question, as what we say about this turns upon the details of our account of tropes. If we opt for Stout's view, for instance, then a trope will stand in the part/whole relationship to its object, as it will just be part of the collection of tropes which constitute the object. But this still leaves questions unanswered, as we have to explain what 'glues' these tropes together into objects. On Husserl's view, on the other hand, we have the task of trying to account for this 'founding' relationship

between a trope and its object. Another issue which is left outstanding is what the principle of individuation is for tropes. Should we say, with Schaffer, that they are individuated by their spatiotemporal location? Or, do we have to accept Campbell's conclusion that tropes are primitively quantitatively distinct?³²

These difficult issues will have to be left pending here. The aim of this section cannot be to offer a complete analysis of the nature of tropes, as this is too large a task. Instead, I just wanted to make clear what we do need to say about these entities, if we adopt CTP and the further causal considerations given. This, as we've seen, does not result in a comprehensive account of the nature of tropes. However, looking on the brighter side, it has shown that the proposed CTP is committed to less controversial claims regarding the nature of tropes, than it might at first appear.

³² See Schaffer (2001b) and Campbell (1990, p.69).

7. The Laws of Nature

An adequate CTP, I have argued, should offer an informative analysis of properties, which avoids the grounding objection and the others difficulties mentioned earlier. Although we've seen that different CTPs can satisfy these requirements, a generalised form of functional role theory, which appeals to an ontology of tropes, is the most ambitious of these theories. By endorsing the functional role reading of CTP, we get a very definite conception of the nature of properties. For it offers a reductive two-level criterion of identity, a theory of property instantiation and a new and improved version of trope nominalism. Furthermore, this analysis has positive ramifications for a theory of causation. The fine-grainedness of tropes makes them suitable causal relata, and the resulting CTP can capture the motivation behind both singularism and generalism. These advantages for an analysis of causation, however, can also be accrued by combining a weaker version of CTP with an ontology of tropes. So long as CTP is stated using tropes, objections can be rebutted and plausible causal commitments ensue.

So far, however, the discussion has remained firmly world-bound. For all the argued benefits of CTP hold independently of any theses about properties in other possible worlds. In particular, there has been no appeal to an oft-cited argument for CTP, namely, that the form of nomic necessity that results from this theory is a virtue.¹ Now, then, it's time to look at the modal implications of CTP, as formulated here.

7.1 Modal Implications

Many philosophers reject CTP because of its perceived modal implications. Standardly, CTP theorists claim both that the laws of nature are relations between properties, and that the causal profile of a property is essential to it. This renders the laws metaphysically necessary, because if a property's causal profile is essential to it, then its causal relations with other properties cannot alter. Consequently, laws reporting relations between the *very same properties*, cannot vary either. Matters do

¹ See, for instance, Swoyer (1982), Fales (1990), Shoemaker (1998a) and Ellis (2001).

not change given the version of CTP expounded above. If we take a property to be a class of tropes, the condition for entry being that each trope realises the RCL definition of that property, then the causal relations which this property will enter into (through its instances) will not vary. So once we've combined this thesis with the claim that laws state relations between properties, as I argued we should in order to meet the generality assumption, we end up in the same boat – the laws of nature are metaphysically necessary.

Since the days of Hume, this claim that the laws are metaphysically necessary has struck many as obviously false. Perhaps because of this, some CTP theorists have tried to avoid this commitment.² Mellor, for instance, writes,

no mere definition of a property can entail any actual property satisfies it. There may well be worlds where none of the laws in which mass occurs holds, and in those worlds mass will not exist; just as I would not exist in a world that contained no one with any of my actual properties (1995, p.172).

Mellor's point here is that because there are possible worlds where the properties in this world are not instantiated, our laws will not govern these worlds. Consequently, the laws of nature are not metaphysically necessary, because these possible worlds are governed by different laws.

Although I agree whole-heartedly with Mellor's claim that there are possible worlds governed by different laws, I think that he draws the wrong conclusion from this. Mellor is right to say that CTP doesn't exclude the possibility of worlds where mass isn't instantiated, and hence worlds where the laws concerning mass don't apply. But this doesn't mean that the laws of nature aren't metaphysically necessary, just as the fact that water doesn't exist in every possible world doesn't damage the truth of Kripke's claim that water is necessarily H₂O. In order to do this, we would need to establish that *if* water/mass existed in one of these worlds then, in at least some of them, water=H₂O or the law F=MA wouldn't obtain. So the fact that there may well be

² Although most adherents of this view do accept this commitment (see, for example, Swoyer 1982, Fales 1990 and Shoemaker 1998a), Mellor and Oliver (1997) do not. They offer a slightly different form of CTP, because instead of defining properties via their causal contributions, they do so by their place in the laws of nature. Mellor, for instance, writes, "There is nothing to mass but the laws of mechanics, nothing to temperature but the laws of thermodynamics" (1995, p.195). The similarities between these views are so striking, however, that it is sensible to think of their account as a permutation of CTP.

many worlds where mass and the laws involving it do not obtain, does not establish Mellor's conclusion that the laws are contingent.

In order to avoid confusion, therefore, we need to distinguish between a law being true at a world and a law governing or obtaining at a world. A law governs a world or obtains in it, when the properties in that law are instantiated. Whereas a law is true at a world if, had its properties been instantiated there, that law would have governed the world. Unfortunately, however, once this distinction has been recognised, a difficulty for the proposed account emerges. CTP does not rule out the possibility of alien properties, i.e. properties which are not instantiated in this world, but which make an appearance in other possible worlds. But now consider the class of alien properties which, in other possible worlds, interact with properties found in our world. Granted this possibility, we have to say that if those alien properties were instantiated in our world, then they would have interacted with our properties in such-and-such a way.

Now we've seen that laws can be true at worlds, even if they do not obtain in them, because they state relations involving uninstantiated properties. So, in the interests of consistency, a CTP theorist must say that in the actual world, there may be true laws involving our properties and these alien properties. Suppose, for instance, that there is a law which states the relation between force in our world and an alien property, schmass. Because there is a definite fact of the matter concerning how schmass would have interacted with force, this is part of force's causal profile. For the RCL definition for force must state how that property interacts with all properties, whether or not they are instantiated.³ The upshot of this is disagreeable. Since if part of a property's causal profile states how it relates to other, uninstantiated properties, we will never be able to discover the entirety of a property's causal profile at this world. The analysis thus places an unwelcome epistemological burden upon us, as part of a property's nature may remain forever opaque to us.

³ I don't think that we should be tempted to deny this claim, in order to avoid the current problem. Consider, for instance, a very rare property F-ness, which does X on exposure to G-ness. Now suppose that all the substances that possessed this property were eradicated, at least for a certain period of time. In such a case, we would still want to say that there is a truth about how the property G-ness interacts with this extinct property F-ness. So this must be part of the comprehensive RCL definition for G-ness.

This commitment stands in sharp contrast to Shoemaker's claim that we should accept CTP because of its epistemological advantages.⁴ Contrary to what Shoemaker suggests, this analysis allows that two property instances in the actual world could be instances of different (determinate) properties, even though they bestow exactly the same set of causal powers onto their particulars. For these instances may have behaved differently, had an alien property been instantiated. Whatever we make of Shoemaker's arguments, therefore, it is clear that the analysis has its epistemological downside.

How could CTP theorists respond to this difficulty? One option would be to uphold the unified RCL definitions for properties. Although part of these RCL definitions maybe unknowable, this should not deter us. For we still apprehend enough of a property's nature to suffice for all our practical needs. Moreover, the epistemological burden that this analysis places upon us is outweighed by the virtues of this theory. The second option is to abandon the claim that properties are transworld entities. As we'll see in §7.3, this would mean rejecting unified RCL definitions for properties. But such a move would allow us to say that no part of a property's RCL profile remains forever unknown to us, because no part of that profile will state interactions with alien properties.

These two defences are also important to the wider debate. For they reflect two ways a CTP theorist could respond to the challenge that it is implausible to claim that the laws of nature are metaphysically necessary. The first way tries to make this thesis acceptable, by arguing that the advantages of this analysis outweigh its shortcomings. Whereas the second concedes that the opponent is right, and so tries to reconcile CTP with the claim that the laws are not metaphysically necessary (henceforth, I shall refer to this claim as the contingency thesis). In what follows, I shall pursue both of these responses to the objections raised here. This, I hope, will go some way towards showing that the modal implications of CTP are not an insuperable obstacle to the theory.

⁴ See Shoemaker 1980a and 1980b. Most people have rightly found Shoemaker's epistemological arguments for CTP unconvincing. (For criticisms see Swinburne 1980 and Owens 1992). Indeed, now even Shoemaker seems to have abandoned them (see 1998a, p.47).

7.2 An Argument For Metaphysically Necessary Laws

A strong case for the thesis that the laws of nature are metaphysically necessary can be made. There is much to be said for the claim that CTP offers the most plausible transtemporal and transworld individuation conditions for properties. Not only are there few alternatives that cohere with the realist's view of properties, the causal features of natural properties seem central to the role they play in our theorising.⁵ If reflections on the nature of properties thus point to the claim that the laws of nature are necessary then, in the absence of compelling arguments to the contrary, this is what we should conclude. However, there are reasons to think that the metaphysical necessity of laws should be endorsed regardless of any considerations concerning the correct individuation conditions for properties. This is all the better for CTP theorists, since if they can show that CTP is committed to an independently plausible thesis, no objection, indeed support, is offered to their theory.

One of the main arguments for the thesis that the laws are metaphysically necessary is based upon the perceived necessity of laws.⁶ It has long been observed that laws purport to tell us not merely what has or will happen, but also what *must* happen. This necessity seems absolutely central to our conception of laws, since it accounts for many other features which we take to be characteristic of them.⁷ For instance, if we think that G must follow F, this explains why we are happy to project this state of affairs into unobserved or counterfactual situations, thus making it a useful tool in prediction. Because of this, a number of philosophers have argued that a proper understanding of nomic necessity is a precondition for an adequate account of laws.⁸ But many analyses of laws arguably fail this test. In what follows, I shall argue that the analysis of laws which ensues from CTP, respects their necessity better than other non-Humean theories, such as those offered by Dretske (1977), Armstrong

⁵ See, for instance, Shoemaker (1980a and b), Swoyer (1982) and Blackburn (1991) for a defence of this latter claim.

⁶ Swoyer (1982), Fales (1990) and Ellis (2001) all employ this kind of argument. Due to the limits of space, I shall restrict the discussion to this argument, but for one of a different sort, see Bird (2002).

⁷ Dretske (1977 p.262-3) lists six features which we take to be characteristic of laws but not of true universal generalisations. These features, he argues, are manifestations of "ontological ascent" (p.263), which mark the shift from talking about individual events, to talking about laws.

⁸ See, for instance, Van Fraassen (1989, §2.4 and 2.5) and Ellis (2001, §6.4).

(1983) and Tooley (1987).⁹ I am taking these as my opponents because Humean accounts have a very different starting point, which cannot be accepted by those sympathetic to CTP and the resulting singularist model outlined here. So the account to concentrate on is the most popular, alternative non-Humean theory.

Armstrong argues that laws are second-order states of affairs, such as $N(F,G)$, in which two first-order universals (in this case, F and G) are related by a certain dyadic second-order universal, N , the relation of necessitation. According to Armstrong, it is a contingent matter which universals are related by the relation of necessitation. But if $N(F,G)$ obtains, then the first order regularity $\forall x(Fx \rightarrow Gx)$ has to obtain in that world. Consequently, if $N(F,G)$ and Fa obtain in a world then, granted that there are no other states of affairs that could act as defeaters, Ga will necessarily follow. Armstrong's claims that laws are contingent and yet involve relations of necessitation seem unlikely bedfellows. But they are not inconsistent. To say that a relation of necessitation holds between F and G does not imply that the relation holds necessarily. However, Armstrong owes us an account of how this relation of necessitation is supposed to capture the thought that $N(F,G)$ must (in some sense) hold, given his commitment to the claim that the law, $N(F,G)$, may not have obtained.

In order to make the issue more concrete, consider these, we'll suppose, true universal generalisations: 'sugar dissolves in water' and 'every time I go to Blackpool it rains'. According to Armstrong, the former nomic generalisation obtains out of a certain necessity, and this is reflected by the fact that it holds not only in this world, but also in nearby possible worlds where the laws are the same. This alone, however, isn't enough to capture the thought that the laws tell us what *must* happen, not merely what has or will. After all, we can equally say that 'every time I go to Blackpool it rains' is true in the actual world and in all those possible worlds where this generalisation obtains. But this, we all agree, is not necessary in any way. So the claim that nomic generalisations obtain in worlds with the same laws as ours fails to latch onto an interesting sense of necessity. For, by parity of argument, we can establish that

⁹ For a basic characterisation of Humean versus non-Humean accounts of laws, see §5.1. In this discussion, I shall utilise the account offered by Armstrong, as this is a typical and influential instance of this kind of non-Humean approach.

true universal generalisations are necessary in this sense, since they too obtain in all those worlds where the same (non-nomic) generalisation holds.

Fortunately, however, this does not exhaust Armstrong's analysis of nomic necessity. He argues that the claim 'sugar dissolves in water' is distinguishable from 'every time I go to Blackpool it rains', because only the former is subsumable under a law of the form $N(F,G)$, and so reports a relation of necessitation taking place. Once this is granted, Armstrong can offer some account of the differing modal status of nomic and true universal generalisations. The class of possible worlds where the laws hold is significantly different from the class of worlds where the same true universal generalisations hold, as in the former worlds, the very same states of affairs are related by this special relation of nomic necessitation. So Armstrong can say that a generalisation *must* obtain iff it obtains in all possible worlds where our laws hold.

With this relation of nomic necessitation then, Armstrong can claim that there is an important difference between the generalisations which obtain in all worlds with the same laws, and those which do not. But, as it stands, this analysis does nothing to explain why this class of possible worlds where the laws hold is so special. What is it about being true in all these worlds that legitimises the claim that these states of affairs *must*, in some sense, happen? The relation of necessitation is supposed to provide an answer to this question. But when asked what it is about this relation that accounts for the unique modal dimension of laws (and so for the fact that they support counterfactuals etc.) the account falls silent. Armstrong writes, "The inexplicability of necessitation just has to be accepted. Necessitation...is a primitive or near primitive that we are forced to postulate" (1983, p.92.)

Now this, in a way, is fair enough. Armstrong has come clean and just said, 'look there isn't much we can say about this relation which explains, amongst other things, the necessity of laws'. But the dearth of such an analysis has serious repercussions for the theory. First, it leaves non-Humeans wide open to the Humean objection that this talk of necessary connections is unintelligible. Not only are such connections obscure, they also fail to do the work which they were introduced to do. For we're given no non-circular analysis of the necessity of laws. Although we can say that a state of affairs is nomically necessary iff it is true in all those worlds where the laws hold, this hardly counts as a satisfying analysis of the phenomena. Furthermore,

this shortcoming casts doubt on the theory's ability to deal with other nomic characteristics. For it is the claimed necessity of laws which legitimises our practice of projecting them into unobserved and counterfactual scenarios.

Second, without a well-defined notion of necessity, the inference, $N(F,G)$, Fa , therefore, Ga , becomes suspect. This move looks justified because if it is necessary that all F s are G s, every actual and possible instance of F will be G . But Armstrong's 'necessitation' isn't the standard notion of necessity, captured by the idiom 'it couldn't have been otherwise'. So why does this relation ensure that if Fa , then Ga ? As Armstrong doesn't have an alternative, positive conception of this relation of necessitation, we just have to accept the inference from $N(F,G)$ and Fa , to Ga . But this, as Lewis elegantly remarks, is less than satisfying:

Whatever N may be, I cannot see how it could be absolutely impossible to have $N(F,G)$ and Fa without Ga ...The mystery is somewhat hidden by Armstrong's terminology. He uses 'necessitates' as a name for the lawmaking universal N ; and who would be surprised to hear that if F 'necessitates' G and a has F , then a must have G ? But I say that N deserves the name of 'necessitation' only if, somehow, it really can enter into the requisite necessary connections. It can't enter into them just by bearing a name, anymore than one can have mighty biceps just by being called 'Armstrong' (1983b, p.40).

At least as it stands, then, Armstrong's account is seriously deficient.

A much more robust and perspicuous analysis of nomic necessity is offered by CTP. According to CTP, the necessity of the laws stems from the nature of the properties involved in those laws. The essence of a property consists in its causal profile. So if it is part of the causal profile of F -ness that all its instances are co-instantiated with G -ness, there couldn't be an instance of F -ness which wasn't co-instantiated with G -ness. As a result, nomic necessity is just a species of the better understood metaphysical necessity. While there could be possible worlds governed by laws different from those in this world, the same properties couldn't be instantiated in a world and yet different laws obtain. For if a possible world has the same properties as the actual world, then that world must be governed by the laws that hold here. So the account captures the idea that our laws of nature couldn't have been otherwise, for they govern every possible world where the relevant properties exist.

By endorsing this stronger form of necessity, a number of problems which beset Armstrong's view are overcome. First, it justifies our making modal claims about F-ness. All instances of F-ness must be co-instantiated with G-ness since, otherwise, they wouldn't be instances of F-ness. As a result, our custom of holding laws fixed in counterfactual situations is vindicated because, very simply, they are true in all these possible situations.

Second, if nomic necessity is just a form of metaphysical necessity, the inference, $N(F,G), Fa, \text{ therefore, } Ga$ is justified, since straightforward metaphysical necessity does imply that if something is necessarily true, then it is true in all possible worlds. So we can legitimately infer Ga from $N(F,G)$ and Fa .

Third, the analysis offers a clear way of distinguishing between laws and true universal generalisations. A CTP theorist can say that there is a significant difference between the two. The nomic generalisation 'sugar dissolves in water' obtains in every possible world where the properties are the same as they are here. But the true generalisation 'every time I go to Blackpool it rains' won't obtain in all possible worlds with properties just like ours. For even if two worlds have the same deterministic laws, they needn't have the same particular history as they could have had different initial conditions. So this generalisation can't be projected into all possible worlds with the same properties. This provides a better analysis of the modal disparity between laws and true universal generalisations. It avoids the circularity inherent in Armstrong's account, as CTP theorists can delineate the possible worlds where the nomic generalisations hold, without appealing to those very laws.

CTP, therefore, promises to offer a substantial and meaningful analysis of the necessity of laws, since it captures the idea that things couldn't be just like or qualitatively similar to the ways things are here, and yet have turned out differently.¹⁰ Some may object, however, that the sort of necessity that ensues from CTP is trivial or misplaced.¹¹ For nomic generalisations are only necessary because there is a certain

¹⁰ Unless, of course, the laws that govern the world are probabilistic.

¹¹ This objection was inspired by Fine (2002, §3). Fine considers a slightly different account of nomic necessity, as he argues against the thesis that nomological necessity can be defined in terms of metaphysical necessity plus certain facts about what properties or kinds are instantiated here. This thesis isn't an instance of CTP, as nomic necessity isn't said to be subsumable under metaphysical necessity. However, it is difficult to see why anyone would endorse this kind of position, unless they shared CTP's approach to questions about how properties should be individuated. For alternatives to CTP (i.e. those accounts offered by the Humean approach and Armstrong's non-Humean approach) allow that

class of possible worlds where all the properties of this world are instantiated. But, we may object, surely this mislocates the source of nomic necessity? For the mere fact that there are these possible worlds with the same properties as ours does not constitute an adequate account of what this force, or form of nomic necessity might be. We need to say more than this in order explain why properties *have* to behave in the way that they do.

I think that this objection is forceful against one possible way of understanding the proposal being put forward. But I hope that further clarification of what, at least, I have in mind, will undermine it. I am not suggesting that laws are necessary *because* they obtain in all possible worlds with the same properties as ours. I do not want to claim that the class of possible worlds with the same properties as ours is responsible for the truth of statements asserting nomic necessities. The possible world talk just provides a useful way of spelling out the kind of necessity nomological necessity is. According to CTP theorists, if it is true that mass acts according to the inverse square law then this is true in all possible worlds, hence nomological necessity is just a form of metaphysical necessity. But this does not commit CTP theorists to the counter-intuitive claim that the laws are nomologically necessary because there is a certain class of possible worlds where only the properties of this world are instantiated. They are free to offer another, more plausible analysis of the source of nomological necessity, as they are only committed to the claim that the laws are metaphysically necessary.

Now there are, no doubt, numerous explanations which could be given of why our laws obtain in all the worlds with the same properties as ours. Sidelle, for instance, in his criticism of this view, tries to foist one particular explanation upon its adherents. He argues that if the laws are metaphysically necessary, this is only in an uninteresting sense. For the metaphysical necessity is created by our semantic conventions governing what we would call what. If it metaphysically necessary that $F=MA$, for instance, this is only because our linguistic practices are such that we wouldn't refer to a property as force if it stood in different relations to the properties of mass and

our properties can be governed by different laws. So the mere fact that a possible world instantiated all our properties would not be enough to ensure that it was a nomological possibility. Nevertheless, it is worth outlining a variant of the objection posed by Fine, as it may be thought to tell against CTP.

acceleration. According to Sidelle then, the necessity of laws stems from conceptual truths about our linguistic practices, it does not reflect any deep metaphysical claims about the nature of properties.¹²

I do not wish to dispute the claim that this is a possible explanation of why the laws are necessary, but the tradition that I have been pursuing here rejects this analysis of the necessity of laws. For it locates the source of nomic necessity in the nature of properties, and claims that there is something about those very natures which accounts for their necessary connections with one another. Although our semantic conventions may latch onto these deep metaphysical truths about the nature of properties, the necessity in question is metaphysical rather than conceptual. Now, admittedly, a complete analysis of how this nomic necessity issues from the natures of properties has not been given and so, because of this, more needs to be said in order to avoid the charge of obscurity made against Armstrong's account. In this thesis, I have suggested that the tropes of particulars are the source of causal necessity and nomic necessity should be seen as being built up from instances of the necessity evident when C causes E. This account clearly needs augmenting further, with an analysis of what it is for C to causally necessitate E.¹³ But, going back to the objection made earlier, it is far from clear that the account does mislocate the source of necessity in the world. For if we accept the plausible singularist's thesis and the allied grounding intuition, the necessity is located in exactly the right place, namely, in the intrinsic features of particulars.

What I think the Sidelle and Fine-inspired criticisms draw attention to then, is the fact that there is still much work for CTP theorists to do. However, they do not undermine the coherence of the kind of account on offer, nor do they seriously damage the claim that the CTP analysis has significant advantages over its non-Humean rival. Although both require further elucidation, CTP still offers a non-circular, if only partial, analysis of nomic necessity. This allows CTP theorists to distinguish between laws and true generalisations, without appealing to the concept of laws which is

¹² See Sidelle 2002 p.321.

¹³ As I noted in §5.1, due to the enormity of the undertaking, I am unable to offer a comprehensive account of the causal relation here. But I did suggest that any such analysis of what it is for C to cause or causally necessitate E should respect the intrinsicity assumption. And while this does not commit us to any particular account, it does significantly restrict the range of analyses open to us.

implicit in the notion of the relation of nomic necessitation. Moreover, CTP's commitment to the claim that the laws are metaphysically necessary furnishes us with a better analysis of the necessity involved in laws, since it is just the notion of necessity we employ elsewhere. So CTP theorists needn't shy away from the claim that the laws are metaphysically necessary, as there are good reasons, independent of CTP, for endorsing this thesis.

7.3 CTP and the Contingency Thesis

Despite the attractions of the thesis that the causal profiles of properties are essential to them, some may still find this view difficult to accept. The epistemological burden it results in, not to mention the allure of the contingency thesis, may be thought too great. Therefore, it is worth looking at whether some form of CTP could be adopted, without any of these modal commitments.

One way of avoiding these implications would be by adopting this restricted criterion of identity for properties: in worlds with the same laws, two tropes/property instances are instances of the same property iff they realise the same causal profile. This version of CTP is clearly compatible with the contingency thesis, because it rejects the claim that the causal profiles of properties are essential to them. Furthermore, as the criterion is restricted to worlds with the same laws, questions about how these properties interact with alien properties absent from our world doesn't arise. Unfortunately, however, this thesis, by itself, is too insubstantial to be called a CTP. It diminishes CTP to the claim that properties have the causal features they do in worlds with the same laws, hence we can identify instances of properties via them. But this tells us very little about the nature of properties. In particular, it does not state that a property is exhaustively characterised by its causal profile (strong CTP), or even that the most crucial characterisation of a property is given by its causal profile (weak CTP). Indeed, a vocal opponent of CTP, namely Armstrong, could endorse this thesis, since he thinks that the causal features of properties are stable in worlds with the same laws. Therefore, it is at best misleading to call just this restricted criterion of identity a CTP.

Another way of trying to avoid these modal commitments is by arguing that properties are world-bound. In other words, they are entities which can only exist in

one possible world. This way, we can still maintain that the property of F-ness is identical with the set of tropes which realise the RCL condition for F-ness. It's just that, now, this property of F-ness does not exist in any worlds other than our own. This allows us to endorse the contingency thesis, as we cannot say that the causal profiles of properties are essential to them. Similarly, since properties cannot exist in other possible worlds, questions about how they interact with properties not instantiated here are bypassed. Avoiding the modal implications of CTP in this way, however, appears rather desperate. For it looks like we are forced to deny simple modal truths, such as 'this room would have been hot, even if it had been 35°C, rather than 36°C'. Why? If properties are world-bound, then the property of being hot or of having a temperature of 35°C or 36°C cannot exist in worlds other than our own. So this room couldn't have had a temperature different from 36°C.

CTP theorists, however, can avoid this counterintuitive consequence by making the same move as Lewis. In the debate about the transworld identity of particulars, Lewis argues that *de re* modal truths about particulars can be saved by employing counterpart theory. This states that the object relevant to the truth of a sentence, such as 'Sally could have had blonde hair', isn't Sally herself, but rather a counterpart of Sally. So this sentence is true iff a counterpart of Sally has blonde hair. But which particulars are these counterparts of Sally? Although Lewis refuses to give necessary and sufficient conditions for the counterparthood relation, he does say this,

Your counterparts resemble you closely in content and context in important respects. They resemble you more closely than do the other things in their worlds. But they are not really you. For each of them is in his own world, and only you are here in the actual world (1986a, p.112).

Although the counterpart relation is not precise, therefore, Sally's counterparts at other possible worlds are those individuals which resemble her the most in the relevant respects.

If we apply counterpart theory to properties, then CTP theorists can say that the entities relevant to modal truths involving properties aren't those very properties, but rather counterparts of those properties. Following Lewis, a counterpart of a property can be said to be that set of tropes which realises the causal profile most

similar to that property's causal profile in the actual world.¹⁴ So, in nearby possible worlds, a counterpart of F-ness will be a set of tropes whose members all realise the functional definition of F-ness in this world. However, in more distant worlds, the counterpart of F-ness will be a set of tropes whose members all realise the causal profile most similar to the causal profile of F-ness in this world.¹⁵ Thus, a modal claim such as, 'the property of being 100°C could have had a slightly different causal profile', is true iff the property of being 100°C has a counterpart whose tropes all realise a slightly different causal profile from the tropes of 100°C in this world.¹⁶

By endorsing this counterpart theory for properties, therefore, we are free to accept the contingency thesis. But is it still sensible to call the resulting account a *CTP*? I think it is, for as well as the restrictive two-level criterion of identity mentioned earlier, we also have an analysis which respects the claim that the identity and transworld identity of a property is determined by the causal features of its tropes.¹⁷ This is important because it means that there is still this crucial connection between the nature of a property and its causal profile. If the identity of a property is determined by the causal roles that tropes realise in all possible worlds, we can maintain that the most significant characterisation of a property is given by its causal profile, in accordance with weak *CTP*.

If we embrace trope universalism, this is all we will be able to say, as the abstract universal is something distinct from the set of tropes. However, if we opt for trope nominalism, we get more than this, as we can preserve strong *CTP*'s claim that properties are exhausted by the causal profiles of their tropes. While there isn't one lot

¹⁴ If we endorse trope universalism instead of trope nominalism, then the set of tropes won't be the property but rather reflect its identity conditions.

¹⁵ This analysis raises the difficult issue of how we can distinguish a possible world where there is a counterpart of F-ness from one where no such property is instantiated. In reply, I suspect that we could plausibly argue, with Lewis, that there are "no determinate right answers to questions about representation de re" (1986a, p.251). When determinacy in virtue of the causal character of tropes gives out, determinacy itself gives out.

¹⁶ If we wanted to preserve talk of properties being transworld entities then, following Lewis, the property of F-ness could be identified with the class of tropes that realise the causal profile of redness in this and nearby possible worlds, plus those counterparts of tropes which realise slightly different causal profiles in further away possible worlds. This formulation requires that we accept unrestricted composition, i.e. the view that any mereological sum can be said to compose an entity (see Lewis 1986a, p.213 for details). But apart from this, the two views aren't substantially different.

¹⁷ This mirrors Lewis's definition of anti-haecceitism, i.e. the view that two worlds cannot differ in representation de re concerning some individual, without differing qualitatively in some way (1986a, p.221).

of causal behaviour which serves to individuate a property, properties are still nothing over and above the behaviour of their instances. For they are just sets of tropes which realise certain causal profiles. On the present analysis, these sets of tropes are unable to capture the modal features of these properties. But by appealing to counterparts of properties, which are themselves just further sets of tropes meeting certain causal conditions, we get the full picture. So we can adopt this proposal while remaining faithful at least to the spirit of CTP.

As counterpart theory is controversial, this proposed uniting of CTP with the contingency thesis will be too. But as I think this is the only way the two can be brought together, CTP theorists face a choice: either they can deny the contingency thesis or adopt counterpart theory for properties. We are thus left weighing the pros and cons of each position. All those problems that philosophers have claimed befall counterpart theory will be set aside here. But before moving on, I should first comment on a tension between the proposal here and what was said previously. Earlier I argued that two objects genuinely resemble each other with respect to F-ness iff they possess a trope which realises the RCL condition for F-ness. According to the proposal here, however, there could be worlds where objects are alike with respect to F-ness, but which do not have tropes realising this RCL condition. So this raises the question: why suppose that, in this world, genuine resemblance between objects with respect to F-ness requires that they possess a trope which realises a particular RCL condition, while in other possible worlds, a near match to this RCL condition will suffice?

In reply, CTP theorists can say that the nomological character of our world gives us reason to think that genuine resemblance with respect to F-ness is analysed via tropes realising a particular RCL condition. Since instantiating the property of F-ness gives rise to similar causal powers and patterns of causal relations in this world. If we are sympathetic to the contingency thesis, however, then genuine resemblance between objects with respect to F-ness may not be so constrained in other possible worlds. Why? On the view proposed here, if there are tropes at a world, then this

world is similar to ours in that it too must have a nomological character.¹⁸ But, given the contingency thesis, there's nothing to stop this nomological character being manifested in different ways in different possible worlds. Consequently, as there are still the same relations between tropes, properties and causation in these worlds, genuine resemblance with respect to F-ness will have to be analysed via the causal capacities objects have in virtue of their tropes there.

If we are prepared to grant counterpart theory for properties, therefore, we can have a CTP which endorses the contingency thesis. Such a theory works better if we embrace trope nominalism, as then we have a more informative theory. But a very weak variant of CTP can be formulated given trope universalism.¹⁹ I argued earlier, however, that there were advantages to be had by denying the contingency thesis. So before finishing, it's worth looking at whether CTP theorists should feel under any pressure to accept this counterpart theory, due to forceful arguments for the contingency thesis.

7.4 The Conceivability Argument

'Tis an establish'd maxim in metaphysics, That whatever the mind clearly conceives includes the idea of possible existence, or in other words, that nothing we can imagine is absolutely impossible (Hume 1978, p.32).

Philosophers often argue from premises about what situations are conceivable to conclusions about what is possible. At least some of our conceivings, it is claimed, provide us with knowledge about what is and isn't possible.²⁰ It would be imprudent to reject this thesis, for our access to modal truths seems to require that the conceivable, at least often, is possible. But if we accept that conceivability is a good

¹⁸ See §5.6. Tropes, by their nature, are instances of properties. So, granted the aforementioned claim that tropes are the truthmakers of causal statements, any world with tropes will give rise to more general nomological relations between properties.

¹⁹ The same is not true, however, given a version of weak CTP which claims that property instances are instantiations of Armstrong-type universals. For, then, there is some constituent of the property instance which is literally identical in all its instances (see §6.1), and this will suffice to settle questions about the transworld identity of universals. If we opt for this no-trope view, therefore, either we have to say that there is a necessary correlation between a universal and the behaviour its instances manifest, so we are committed to the view that the causal profiles of universals are essential to them. Or, we can endorse Armstrong's view and so deny CTP. For it is not the causal features of its instances which determine a universal's identity across possible worlds, rather it is the universal's nature or "quiddity" (see Armstrong 1999a, p.28).

²⁰ See Gendler and Hawthorne (2002) for a discussion of this thesis.

guide to possibility, then this argument can be formulated against the necessity of laws:

- 1) Conceivability is a good guide to possibility.
- 2) We can conceive of a situation in which the laws in our world do not obtain.
- 3) Therefore, there is a possible world where the laws in our world do not obtain.

As it stands, however, this is a very weak argument against CTP, as CTP theorists are not committed to the claim that *our* laws obtain in every possible world. Hence, the thesis is not damaged by the mere fact that we can conceive of possible worlds where nomologically impossible events (by the standards of our world) happen. In these cases, CTP theorists can simply say that although the properties that exist in these worlds may look superficially similar to our own, as they have different causal profiles, these are distinct properties. So CTP theorists can easily explain away the fact that we are able to imagine nomologically impossible events happening in different possible worlds, as they can say that what we are imagining in these cases, are possible worlds where different laws obtain, because different properties are instantiated in these worlds.

For this conceivability argument to have bite against CTP, therefore, we need to substitute 2), for this stronger premise:

- 2*) We can conceive of the very same properties in this world being governed by different laws.

I have to confess, I am suspicious of the claim that our conceivings furnish us with so much theoretical content. But even if we grant that there is a *sense* in which 2*) is true, can we be at all sure that this sense of conceivability really is a good guide to possibility?

In order to get a better grasp on what is being asked, we need to consider which conceivings or mental states are supposed to be good guides to possibility. Some say that it is our non-sensory conceivings which are our guide, the idea being that if we can suppose that a certain scenario obtains without any incoherence, then

that scenario is possible.²¹ Alternatively, we could say that it is our sensory imaginings which are our guide. In sensory imagination, we do not imagine in the sense of “suppose” or “entertain a thought”, rather we “imagine from the inside being in some conscious state”.²² Some argue that it is this experiential form of imagination - the capacity to imagine a certain experience, say of seeing a tiger - which provides us with defeasible evidence for a situation’s possibility. For experience is our primary indicator of what is actually the case, so imagining experiencing a certain situation provides prima facie evidence for the possibility of that situation.²³ These two options are obviously not mutually exclusive, but to begin with, I shall concentrate on the claim that CTP is undermined because our non-sensory conceivings are a good guide to possibility.

Since Kripke (1972), it has become widely recognised that not all (non-sensory) conceivings are good guides to possibility. While there is a sense in which we can conceive of water being something other than H₂O, for instance, or Hesperus being something other than Phosphorus, these sorts of conceivings are thought to be misleading guides to possibility. In order to make the conceivability argument against CTP convincing, therefore, opponents need to show why the sense of conceiving which is a good guide to possibility excludes CTP’s claim that ‘necessarily, F=MA’, but doesn’t rule out generally accepted cases of metaphysical necessity, such as ‘necessarily, water = H₂O’ and ‘necessarily, Hesperus = Phosphorus’. This, I think, will be a very difficult thing to show. Take, for instance, Kripke’s defence of the claim that ‘necessarily, water = H₂O’. He argues that although we seem able to conceive otherwise, this is just an illusion created in one of two ways. Either what we are

²¹ Descartes, for instance, argues that it is our intellectual conceiving, not our sensory imaginings, which are our guide to what is possible. We cannot, for instance, sensorily imagine a figure with one hundred sides, as we could not distinguish it from one with ninety-nine sides (see 1996, §72). However, we can non-sensorily conceive it or suppose it to be true without contradiction, so we know that such a figure is possible.

²² Peacocke 1985, p.21.

²³ See Peacocke 1985, p.31. This idea is also echoed in Campbell (2002b). He argues that “our conception of what is and is not possible for the demonstrated object, are grounded in our experience of the thing” (p.140), the idea being that our experience gives us knowledge of categorical facts about the way the object actually is, and this experience somehow illuminates the way an object could be. For a note of caution on this idea, see Gendler and Hawthorne (2002 §1.2). It is not clear, however, whether Peacocke and Campbell think that sensory imaginings are the only sorts of conceivings which are a guide to possibility. A more moderate line would be to say that both are, although perhaps our sensory imaginings provide us with a better guide than our non-sensory conceivings.

supposing is a situation in which the substance we refer to as water (and which we now know to be H₂O) could have been discovered to be something other than H₂O. This is a case of epistemic possibility. Our ignorance or imagined ignorance means that we find it possible to believe that the hypothesis could turn out either way. Or, we can seem to conceive of a situation where water isn't H₂O, by imagining that the mode of presentation which we have come to recognise water by, is correlated with a substance which has a different molecular structure.

In both sorts of cases, Kripke wants to claim that we are not conceiving of the scenario in a way relevant to proving its metaphysical possibility. But if we are prepared to grant that these are cases of mere seeming conceivability, then CTP theorists can make analogous claims for their theory. They can say that what we seem to be conceiving when we imagine that $F \neq MA$ only reflects the epistemic possibility that it could have been otherwise (in light of the available evidence) for all we knew. It does not establish that, granted they are true, they are only contingently so. Or, they could say that what we are really conceiving, are not cases where the very same properties are related by different laws, but cases where different properties, whose modes of presentation are very similar to our own, are related by different laws.

Yablo (1993) attempts to bolster the claim that conceivability, once it is properly understood, is a good guide to possibility. He characterises the relevant sense of conceivability as this: "I find p conceivable if I can imagine, not a situation in which I truly believe that p , but one of which I truly believe that p " (p.26). Then he takes us through apparent counter-examples to illustrate why they do not constitute a rebuttal of this principle. So, for instance, he argues that this sense of conceivability does not allow us to conceive of a situation in which Hesperus (H) \neq Phosphorus (P). For to imagine truly believing that $H \neq P$ is to imagine truly believing that Venus \neq Venus, and this clearly cannot be done. What explains our appearances to the contrary? Yablo argues that this can be put down to the fact that "I can imagine believing *something* true with my Hesperus \neq Phosphorus-thought" (p.24). Although my thought, in the actual context, expresses the proposition that 'necessarily, $H \neq P$ ',

in other hypothetical contexts it may have expressed a different proposition.²⁴ For instance, we can imagine that if Mars had been responsible for our sighting of Phosphorus, then our thought would have expressed a proposition with truth conditions for Venus \neq Mars. In which case, we can say that I can truly imagine believing the thought expressed in this hypothetical situation. But this isn't a case where I truly believe the thought expressed by my actual proposition.

Yablo's characterisation of the relevant sense of conceivability, however, not only lets Kripke's a posteriori identities through the net, it can also accommodate the claim that the laws are metaphysically necessary. We can say that what we are imagining when we entertain the thought that 'F \neq MA' isn't a case where I truly believe that this is so. For once we've fully comprehended what force, mass and acceleration are, we'd realise that, like 'water = H₂O', force couldn't be anything else. However, we can still imagine *something* true by the thought 'F \neq MA', as we can imagine a hypothetical situation in which a different proposition was expressed by this thought. For instance, we can imagine that if force had been associated with a different causal profile, then our thought 'F \neq MA' would have expressed another proposition with distinct truth conditions. But, again, this isn't a case where I truly believe the thought expressed by my actual proposition.

Given, then, Kripke's and Yablo's definitions of those non-sensory conceivings which are a good guide to possibility, it is not clear that we can non-sensorily conceive, in the relevant sense, of properties behaving differently. But what of our sensory imaginings? Do they give us reason to think that properties could behave differently? Peacocke's principle that if we can experientially imagine having an experience, then such an experience is possible, does seem to have some plausibility.²⁵ Furthermore, it looks like it can be utilised to support the claim that 'necessarily, H=P' and 'necessarily F=MA' are disanalogous. For we can say that in the case of H=P, we cannot imagine an experience of seeing H without also seeing P, because (given that H=P) we would have to imagine an experience of both seeing

²⁴ By 'thought' here, Yablo means "the internal state or act that determines, in context, which proposition I believe" (p.24). So a thought in one context can express a different proposition than the same thought in another context.

²⁵ See Peacocke 1985, p.31.

Venus and not seeing Venus. But when we consider the laws of nature, it seems quite easy to imagine having an experience of, for instance, a ball hitting another and yet of one not moving. Therefore, in this case, sensory imagination does provide us with defeasible evidence for the claim that the laws are contingent.

Does this succeed in demonstrating that the burden of proof is on those who wish to deny that the laws are contingent? I doubt it, as it is far from clear that we can sensorily imagine a case where the laws are not true. Suppose, for instance, that we sensorily imagine a ball hitting another ball and the second failing to move, does this constitute an experience of non-Newtonian behaviour? No, not unless we add in a commentary claiming, for instance, that there are no other Newtonian forces at work on the second ball which stop it from moving, that the mass of the second ball is not so much greater than that of the first that it doesn't move with the impact of the second, that the properties in this scenario are the same as those instantiated in our world, etc. The experience alone does not suffice to show that what we perceive is a case where Newton's laws do not hold. We need to add in the extra commentary for this conclusion to follow.

It may be objected that this imposes a false distinction between the content of an experience and the interpretation we put upon it. But this doesn't matter, because the force of Peacocke's principle rests on the fact that we can separate what we experience from what we merely non-sensorily conceive. Potential experience, on this view, is the source of modal knowledge. Mere coherent suppositions do not protect us from impossibility. Hence, if the objector is right to say that interpretations completely permeate our experience, then our sensory imaginings won't be a better guide to possibility than our non-sensory conceivings. Either way, we are left having to establish that non-sensory conceivings are a good guide to possibility, in order for the conceivability argument against CTP to go through.

It is not clear, therefore, that given the sense of conceivability required for the maxim 'conceivability is a good guide to possibility', we can conceive of the very same properties being related by different laws. This has two positive consequences. First, it means that CTP theorists who deny the contingency thesis can still maintain an epistemology for modality which takes, as its starting point, facts about what we

can and cannot conceive. Although a lot needs to be said on the topic of modal error,²⁶ if we accept CTP, we can at least explain why we cannot conceive, in Yablo's sense, of properties being governed by different laws. This should come as a welcome relief to CTP theorists, because they, no less than anyone else, have to appeal to what is conceivable in their modal reasoning. Second, although much more still needs to be said about the issues raised here,²⁷ the discussion has gone some way towards weakening the most forceful argument against the claim that the laws are metaphysically necessary. So CTP theorists need not be unduly worried by their commitment to this thesis.

²⁶ See Yablo 1993 §XII.

²⁷ In particular, more needs to be said about two-dimensional approaches to modal discourse. Using the machinery of two-dimensional modal logic, these accounts try to portray both how a sentence represents things to be, and how a sentence would have represented things to be had it been uttered in a different setting. This approach raises difficult and complex issues which cannot receive adequate attention here. But for more on this, see Chalmers 1996 and 2002.

Conclusion

What makes a property the property it is, what determines its identity, is its potential for contributing to the causal powers of the things that have it (Shoemaker 1980a, p.212).

How exactly this statement should be interpreted, as we've seen, takes some working out. But I think that it repays the effort, as this statement, and the account it expresses, has the potential to provide illuminating answers to many different metaphysical questions. I've tried to outline different versions of CTP which do justice to Shoemaker's original claim that "properties are causal powers"¹ and to his later, weaker formulations of the theory. Although problems were encountered, I think that the theses outlined offer plausible formulations of CTP. For they avoid both the pitfalls of the grounding objection and related criticisms, while resulting in a cogent two-level criterion of identity for properties.

On the issue of causation, two claims were made. First, it was argued that whatever causes and effects are, they need to be fine-grained. The theory of causal relata which ensues from CTP respects this, because it states that property instances are the fundamental causal relata. Second, it was argued that an analysis of the causal relation should respect both the intrinsicity and generality of that relation. Anscombe shows why this is no mean feat to pull off. But by combining CTP with trope theory, we have a sketch of an account of causation which can satisfy both the intrinsicity and generality assumptions. This combination leaves us with an intuitive model of causation, which is the polar opposite of the Humean approach. It is a singularist theory, as the determiners of causal relations and laws are said to be local entities, which are intrinsic to the causal relations.

The ramifications of CTP for an account of properties are just as positive. CTP provides an illuminating account of the nature of properties, which results in a two-level criterion of identity for these entities. There are different ways of developing the basic account, so some flexibility is available in this area. But CTP offers a definite conception of what it is for a property to be instantiated by an object. Furthermore, it

¹ Shoemaker 1980a, p.210.

can be used to group tropes together into natural classes, since CTP presents us with a conception of what resemblance between tropes involves. This aids the cause of trope nominalism, by elucidating which classes of tropes stand for which universals, in cases of determinates and determinables.

There are objections which have prevented many from endorsing CTP, but I have tried to show why these need not deter us. The oft-cited grounding objection is ineffective against the causal theory of properties expounded here. The criticism based on CTP's commitment to the laws being necessary is a little harder to dispense with. I tried to show why this consequence of CTP should be embraced rather than rejected. But there are ways of avoiding this commitment, if you remain unconvinced. Moreover, all the virtues of CTP which have been canvassed here, hold independently of these contentious issues surrounding the necessity of laws. So even if we want to preserve their contingency, the wider case for CTP remains unaffected.

CTP, therefore, has proved to be a fertile line of enquiry. It connects, in a way not done previously, the central notions of properties and causation, explaining why these notions have to be understood in tandem. The result is both a plausible causal ontology and a worked out analysis of the nature of properties.

References

- Achinstein, P. (1974) "The Identity of Properties" *American Philosophical Quarterly*, vol. 11.
- Anscombe, G. E. (1971) "Causality and Determination" reprinted in *Causation*, eds. E. Sosa, and M. Tooley (1993) Oxford: Oxford University Press.
- Armstrong, D. M. (1970) "The Causal Theory of the Mind" reprinted in Armstrong's *The Nature of Mind and Other Essays* (1980) Queensland: University of Queensland Press.
- Armstrong, D. M. (1978) *Universals and Scientific Realism: A Theory of Universals*, vol. 2, Cambridge: Cambridge University Press.
- Armstrong, D. M. (1983) *What is a Law of Nature?* Cambridge: Cambridge University Press.
- Armstrong, D. M. (1996) "Dispositions as Categorical States" in *Dispositions: A Debate*, ed. T. Crane (1996) London: Routledge.
- Armstrong, D. M. (1997) *A World of States of Affairs*, Cambridge: Cambridge University Press.
- Armstrong, D. M. (1999a) "The Causal Theory of Properties: Properties according to Shoemaker, Ellis and Others" *Philosophical Topics*, vol. 26.
- Armstrong, D. M. (1999b) "The Open Door: Counterfactual versus Singularist Theories of Causation" in *Causation and Laws of Nature*, ed. H. Sankey (1999) Dordrecht: Kluwer Academic Publishers.
- Aronson, J. (1971) "On the Grammar of Cause" *Synthese*, vol. 22.
- Austin, J. (1950) "Truth" *Proceedings of the Aristotelian Society*, suppl. vol. 24.
- Ayer, A. J. (1956) "What is a Law of Nature?" *Revue Internationale de Philosophie*, vol. 10.
- Bacon, J. (1995) *Universals and Property Instances: The Alphabet of Being*, Oxford: Blackwell.
- Beebe, H. and Mele, A. (2002) "Humean Compatibilism" *Mind*, vol. 111.
- Beebe, H. (2003) "Causing and Nothingness" forthcoming in *The Counterfactual Analysis of Causation*, eds. J. Collins, N. Hall and L. Paul, Massachusetts: The MIT Press.

- Bennett, J. (1988) *Events and their Names*, Oxford: Clarendon Press.
- Berlin, I. (1950) "Empirical Propositions and Hypothetical Statements" *Mind*, vol. 59.
- Bigelow, J. and Pargetter, R. (1990) *Science and Necessity*, Cambridge: Cambridge University Press.
- Bigelow, J. (1988) *The Reality Of Numbers*, Oxford: Oxford University Press.
- Bird, A. (2002) "On Whether Some Laws are Necessary" *Analysis*, vol. 62.
- Blackburn, S. (1991) "Filling in Space" *Analysis*, vol. 50-1.
- Blackburn, S. (1993) *Essays in Quasi-Realism*, Oxford: Oxford University Press.
- Block, N. (1980) "Introduction: What is Functionalism" in *Readings in Philosophy of Psychology*, vol. 1, ed. N. Block, London: Methuen.
- Block, N. (1990) "Can the Mind Change the World" in *Meaning and Method: Essays in Honour of Hilary Putnam*, ed. G. Boolos (1990) Cambridge: Cambridge University Press.
- Braithwaite, R. (1927) "The Idea of Necessary Connection" *Mind*, vol. 36.
- Burgess J. and Rosen G. (1996) *Subject with No Object*, Oxford: Clarendon Press.
- Campbell, J. (1994) *Past, Space and Self*, Massachusetts: MIT Press.
- Campbell, J. (1995) "Molyneaux's Question" *Philosophical Issues*, vol. 7.
- Campbell, J. (2002a) *Reference and Consciousness*, Oxford: Clarendon Press.
- Campbell, J. (2002b) "Berkeley's Puzzle" in *Conceivability and Possibility*, eds. T. Gendler and J. Hawthorne (2002) Oxford: Clarendon Press.
- Campbell, J. "Intelligible Causation" (manuscript copy).
- Campbell, K. (1981) "The Metaphysics of Abstract Particulars" reprinted in *Properties*, eds. D. Mellor and A. Oliver (1997) Oxford: Oxford University Press.
- Campbell, K. (1990) *Abstract Particulars*, Oxford: Blackwell.
- Cartwright, N. (1989) *Nature's Capacities and their Measurement*, Oxford: Clarendon Press.

- Cartwright, N. (1999) *The Dappled World: A Study of the Boundaries of Science*, Cambridge: Cambridge University Press.
- Chalmers, D. (1996) *The Conscious Mind*, Oxford: Oxford University Press.
- Chalmers, D. (2002) "Does Conceivability Entail Possibility" in *Conceivability and Possibility*, eds. T. Gendler and J. Hawthorne (2002) Oxford: Clarendon Press.
- Chisholm, R. (1957) *Perceiving*, Ithaca: Cornell University Press.
- Clatterbaugh, K. (1999) *The Causation Debate in Modern Philosophy, 1937-1730*, London: Routledge.
- Davidson, D. (1963) "Actions, Reasons and Causes" reprinted in his *Essays on Actions and Events* (1980) Clarendon Press: Oxford.
- Davidson, D. (1967a) "The Logical Form of Action Sentences" reprinted in his *Essays on Actions and Events* (1980) Oxford: Clarendon Press.
- Davidson, D. (1967b) "Causal Relations" reprinted in his *Essays on Actions and Events* (1980) Oxford: Clarendon Press.
- Davidson, D. (1969) "The Individuation of Events" reprinted in his *Essays on Actions and Events* (1980) Oxford: Clarendon Press.
- Davidson, D. (1970) "Mental Events" reprinted in his *Essays on Actions and Events* (1980) Oxford: Clarendon Press.
- Davidson, D. (1977) "The Method of Truth in Metaphysics" reprinted in his *Inquiries into Truth and Interpretation* (1984) Oxford: Clarendon Press.
- Davidson, D. (1985) "Reply to Quine on Events" in *Perspectives on the Philosophy of Donald Davidson*, eds. E. Le Pore and B. McLaughlin (1985) Oxford: Blackwell.
- Denkel, A. (1996) *Object and Property*, Cambridge: Cambridge University Press.
- Descartes, R. (1985) *The Philosophical Writings of Descartes*, vol. 2, trans. and eds. J. Cottingham, R. Stoothoff, D. Murdoch and A. Kenny, Cambridge: Cambridge University Press.
- Descartes, R. (1996) *Mediations on First Philosophy*, trans. J. Cottingham. Cambridge: Cambridge University Press.
- Dretske, F. (1977) "Laws of Nature" *Philosophy of Science*, vol. 44.
- Dodd, J. (2002) "Is Truth supervenient Upon Being?" *Proceedings of the Aristotelian Society*, vol. 102.

- Dowe, P. (2000) *Physical Causation*, Cambridge: Cambridge University Press.
- Ducasse, C. J. (1926) "On the Nature and Observability of the Causal Relation" reprinted in *Causation*, eds. E. Sosa and M. Tooley (1993) Oxford: Oxford University Press.
- Dummett, M. (1973) *Frege: Philosophy of Language*, London: Duckworth.
- Dummett, M. (1991) *Frege: Philosophy of Mathematics*, London: Duckworth.
- Dunn, M. (1990) "Relevant Predication 2: Intrinsic Properties and Internal Relations" *Philosophical Studies*, vol. 60.
- Durham, J. (2002) "The Permeability of the Causal Criteria of Identity for Universals" *Pacific Philosophical Quarterly*, vol. 83.
- Eells, E. (1991) *Probabilistic Causality*, Cambridge: Cambridge University Press.
- Ehring, D. (1997) *Causation and Persistence: A Theory of Causation*, Oxford: Oxford University Press.
- Ehring, D. (1999) "Tropeless in Seattle: The cure for insomnia" *Analysis*, vol. 59.
- Elder, C. (2001) "Contrariety and the Individuation of Properties" *American Philosophical Quarterly*, vol. 38.
- Ellis, B. (2001) *Scientific Essentialism*, Cambridge: Cambridge University Press.
- Ellis, B. and Lierse, C. (1994) "Dispositional Essentialism" *Australasian Journal of Philosophy*, vol. 72.
- Evans, G. (1980) "Things Without the Mind" in his *Collected Papers* (1985) Oxford: Clarendon Press.
- Evans, G. (1982) *The Varieties of Reference*, ed. J. McDowell, Oxford: Clarendon Press.
- Fair, D. (1979) "Causation and the Flow of Energy" *Erkenntnis*, vol. 59.
- Fales, E. (1990) *Causation and Universals*, London: Routledge.
- Fine, K. "The Varieties of Necessity" in *Conceivability and Possibility*, eds. T. Gendler and J. Hawthorne (2002) Oxford: Clarendon Press.
- Foster, J. (1985) *A. J. Ayer*, London: Routledge.
- Geach, P. (1957) *Mental Acts*, London: Routledge.

- Geach, P. (1975) "Names and Identity" in *Mind and Language*, ed. S. Guttenplan (1975) Oxford: Clarendon Press.
- Gendler T. and Hawthorne J. (2002) "Introduction: Conceivability and Possibility" in *Conceivability and Possibility*, eds. T. Gendler and J. Hawthorne (2002) Oxford: Clarendon Press.
- Goodman, N. (1951) *The Structure of Appearance*, Massachusetts: Harvard University Press.
- Hale, B. (1987) *Abstract Objects*, Oxford: Blackwell.
- Harrè, R. and Madden, E. H. (1975) *Causal Powers*, New Jersey: Rowman and Littlefield.
- Hausman, D. (1998) *Causal Asymmetries*, Cambridge: Cambridge University Press.
- Hawthorne, J. (2001) "Causal Structuralism" *Philosophical Perspectives: Metaphysics*, vol. 15.
- Hempel, C. (1965) *Aspects of Scientific Explanation*, London: Collier Macmillan
- Hobbes, T. (1839) *The English Works of "Thomas Hobbes of Malmesbury"*, vol. 1, ed. W. Molesworth, London: Bohn.
- Humberstone, I. (1996) "Intrinsic/Extrinsic" *Synthese*, vol. 108.
- Hume, D. (1978) *A Treatise of Human Nature*, second edition, ed. L. A. Selby-Bigge, rev. by P. H. Nidditch, Oxford: Clarendon Press.
- Hume, D. (1975) *An Enquiry Concerning Human Understanding*, third edition, ed. L. A. Selby-Bigge, rev. by P. H. Nidditch, Oxford: Clarendon Press.
- Husserl, E. (2001) *Logical Investigations*, trans. J. N. Findlay, rev. D. Moran, London: Routledge.
- Jackson, F., Pargetter, R. and Prior, E. (1982) "Three Theses about Dispositions" *American Philosophical Quarterly*, vol. 19.
- Jackson, F. (1998) *Mind, Methods and Conditionals*, London: Routledge.
- Johnson, W. E. (1921) *Logic: Part 1*, Cambridge: Cambridge University Press.
- Kim, J. (1972) "Phenomenal Properties, Psychological Laws and the Identity Theory" *Monist*, vol. 56.

- Kim, J. (1973) "Causation, nomic subsumption and the concept of event" *The Journal of Philosophy*, vol. 70.
- Kim, J. (1976) "Events as Property Exemplifications" reprinted in Kim's *Supervenience and Mind* (1993) Cambridge: Cambridge University Press.
- Kim, J. (1982) "Psychophysical Supervenience" *Philosophical Studies*, vol. 44.
- Kim, J. (1998) *Mind in the Physical World*, Massachusetts: The MIT Press.
- Kripke, S. (1972) *Naming and Necessity*, Oxford: Blackwell.
- Langton, R. and Lewis, D. (1998) "Defining Intrinsic" reprinted in Lewis's *Papers in Metaphysics and Epistemology* (1999) Cambridge: Cambridge University Press.
- Leibniz, G. W. (1981) *New Essays Concerning Human Understanding*, trans. and eds. by P. Remnant and J. Bennett, Cambridge: Cambridge University Press.
- Leibniz, G. W. (1988) *Philosophical Texts*, trans. and eds. by R. Woolhouse and R. Francks, Oxford: Oxford University Press.
- Lewis, D. (1966) "An Argument for the Identity Theory" reprinted in his *Philosophical Papers*, vol. I (1983) Oxford: Oxford University Press.
- Lewis, D. (1970) "How to Define Theoretical Terms" *The Journal of Philosophy*, vol. 67.
- Lewis, D. (1972) "Psychophysical and Theoretical Identifications" *Australasian Journal of Philosophy*, vol. 3.
- Lewis, D. (1973) "Causation" reprinted in Lewis's *Philosophical Papers*, vol. II (1986) Oxford: Oxford University Press.
- Lewis, D. (1980) "Mad Pain and Martian Pain" in his *Philosophical Papers*, vol. I (1983) Oxford: Oxford University Press.
- Lewis, D. (1983a) "Extrinsic Properties" reprinted in his *Papers in Metaphysics and Epistemology* (1999) Cambridge: Cambridge University Press.
- Lewis, D. (1983b) "New Work for a Theory of Universals" reprinted in his *Papers in Metaphysics and Epistemology* (1999) Cambridge: Cambridge University Press.
- Lewis, D. (1986a) *On the Plurality of Worlds*, Oxford: Blackwell.
- Lewis, D. (1986b) "Events" in his *Philosophical Papers*, vol. II (1986) Oxford: Oxford University Press.

- Lewis, D. (1986c) "Causal Explanation" in his *Philosophical Papers*, vol. II (1986) Oxford: Oxford University Press.
- Lewis, D. (1992) "Armstrong on Combinatorial Possibility" reprinted in his *Papers in Metaphysics and Epistemology* (1999) Cambridge: Cambridge University Press.
- Lewis, D. (1995) "Reduction of Mind" ed. S. Guttenplan, *A Companion to Philosophy of Mind* (1995) Oxford: Blackwell.
- Lewis, D. (1997) "Finkish Dispositions" reprinted in his *Papers in Metaphysics and Epistemology* (1999) Cambridge: Cambridge University Press.
- Lewis, D. (2001) "Truthmaking and Difference-Making" *Nous*, vol. 31.
- Lewis, D. (2003a) "Causation as Influence" forthcoming in *The Counterfactual Analysis of Causation*, eds. J. Collins, N. Hall and L. Paul, Massachusetts: The MIT Press.
- Lewis, D. (2003b) "Void and Object" forthcoming in *The Counterfactual Analysis of Causation*, eds. J. Collins, N. Hall and L. Paul, Massachusetts: The MIT Press.
- Linsky, B. (1999) *Russell's Metaphysical Logic*, America: CSLI Publications.
- Locke, J. (1975) *An Essay Concerning Human Understanding*, ed. P. H. Nidditch, Oxford: Clarendon Press
- Lombard, B. (1986) *Events: A Metaphysical Study*, London: Routledge and Kegan Paul.
- Lowe, E. (1989a) "What is a Criterion of Identity?" *The Philosophical Quarterly*, vol. 39.
- Lowe, E. (1989b) "Impredicative Identity Criteria and Davidson's Criterion of Event Identity" *Analysis*, vol. 49.
- Lowe, E. (1998) *The Possibility Of Metaphysics*, Oxford: Oxford University Press.
- Lowe, E. (2002) *A Survey of Metaphysics*, Oxford: Oxford University Press.
- Ludwig, K. (1995) "Shape Properties and Perception" in *Philosophical Issues*, vol. 7.
- MacDonald, C. (1989) *Mind-Body Identity Theories*, London: Routledge.
- MacDonald, C. (1995) "How To Be Psychologically Relevant" in *Philosophy of Psychology: Debates on Psychological Explanation*, eds. C. MacDonald and D. MacDonald (1995) Oxford: Blackwell.

- Mackie, J. (1973) *Truth, Probability and Paradox*, Oxford: Clarendon Press.
- Mackie, J. (1974) *The Cement of the Universe*, Oxford: Oxford University Press.
- Martin, C. (1994) "Dispositions and Conditionals" *The Philosophical Quarterly*, vol. 44.
- Martin, C. (1993) "Power for the Realists" in *Ontology Causality and Mind: Essays in Honour of D. M. Armstrong*, eds. J. Bacon, K. Campbell and L. Reinhardt (1993) Cambridge: Cambridge University Press.
- Martin, C. (1996) "Properties and Dispositions" in *Dispositions: A Debate*, ed. T. Crane (1996) London: Routledge.
- Mellor, D. and Oliver, (1997) "Introduction" in *Properties*, eds. D. Mellor and A. Oliver, Oxford: Oxford University Press.
- Mellor, D. (1974) "In Defence of Dispositions" reprinted in his *Matters of Metaphysics* (1991) Cambridge: Cambridge University Press.
- Mellor, D. (1991) "Properties and Predicates" in his *Matters of Metaphysics* (1991) Cambridge: Cambridge University Press.
- Mellor, D. (1995) *The Facts of Causation*, London: Routledge.
- Mellor, D. (2000) "The Semantics and Ontology of Dispositions" *Mind*, vol. 109.
- Mellor, D. (2003) "For Facts as Causes and Effects" forthcoming in *The Counterfactual Analysis of Causation*, eds. J. Collins, N. Hall and L. Paul, Massachusetts: The MIT Press.
- Menzies, P. (1989) "A Unified Account of Causal Relata" *Australasian Journal of Philosophy*, vol. 67.
- Menzies, P. (1996) "Probabilistic Causation and the Preemption Problem" *Mind*, vol. 105.
- Menzies, P. (1999) "Intrinsic versus Extrinsic Conceptions of Causation" in *Causation and Laws of Nature*, ed. H. Sankey (1999) Dordrecht: Kluwer Academic Publishers.
- Menzies, P. (2003) "Is Causation a Genuine Relation" in *Real Metaphysics: Essays in Honour of D. H. Mellor*, eds. H. Lillehammer and G. Rodriguez-Pereya (2003) London: Routledge.
- Moore, G. E. (1922) "Conception of Intrinsic Value" in his *Philosophical Studies* (1922) Routledge: London.

- Moreland, J. (1985) *Universals, Qualities, and Quality-Instances*, New York: University Press of America.
- Mulligan, K., Simons, P. and Smith, B. (1984) "Truthmakers" *Philosophy and Phenomenological Research*, vol. 44.
- Mumford, S. (1998) *Dispositions*, Oxford: Oxford University Press.
- Neale, S. (1990) *Descriptions*, Massachusetts: The MIT Press.
- Neale, S. (2001) *Facing Facts*, Oxford: Clarendon Press.
- Noordhof, P. (1998) "Critical Notice: Causation, Probability and Chance" *Mind*, vol. 107.
- O'Neill, E. (1993) "Influxus Physicus" in *Causation in Early Modern Philosophy*, ed. S. Nadler (1993) Pennsylvania: Pennsylvania State University Press.
- Owens, D. (1992) *Causes and Coincidences*, Cambridge: Cambridge University Press.
- Peacocke, C. (1985) "Imagination, Experience and Possibility" in *Essays On Berkeley*, eds. J. Foster and H. Robinson (1985) Oxford: Oxford University Press.
- Place, U. T. (1996) "Dispositions as Intentions" in *Dispositions: A Debate*, ed. T. Crane (1996) London: Routledge.
- Plato (1970) *The Dialogues of Plato*, vol. 2, trans. B. Jowett, London: Sphere Books Limited.
- Prior, A. (1949) "Determinables, Determinates and Determinants" *Mind*, vol. 58.
- Prior, E. (1985) *Dispositions*, Aberdeen: Aberdeen University Press.
- Putnam, H. (1967) "Psychological Predicates" in *Art, Mind and Religion*, eds. W. Captain and D. Merrill (1967) Pittsburgh: University Of Pittsburgh Press.
- Quine, W. (1960) *Word and Object*, Massachusetts: MIT Press.
- Quine, W. (1963) *Set Theory and Its Logic*, Massachusetts: Harvard University Press.
- Quinton, A. (1979) "Objects and Events" *Mind*, vol. 88.
- Radner, D. (1978) *Malebranche*, Amsterdam: Van Gorcum.
- Ramsey, F. (1931) *The Foundations of Mathematics and Other Essays*, ed. R. Braithwaite, London: Routledge.

- Reiber, S. (2002) "Causation as Property Acquisition" *Philosophical Studies*, vol. 109.
- Robb, D. (1997) "The Properties of Mental Causation" *The Philosophical Quarterly*, vol. 47.
- Rosenberg, A. (1984) "Mackie and Shoemaker on Dispositions and Properties" *Midwest Studies Philosophy*, vol. 9.
- Russell, B. (1918) "The Philosophy of Logical Atomism" in *Logic and Knowledge*, ed. R. Marsh (1956) London: George Allen & Unwin.
- Russell, B. (1925) *Principia Mathematica*, 2nd ed. Cambridge: Cambridge University Press.
- Ryle, G. (1949) *The Concept of Mind*, London: Hutchinson.
- Salmon, W. (1997) "Causality without Counterfactuals" *Philosophy of Science*, vol. 64.
- Sainsbury, R. M. (1979) *Russell*, London: Routledge.
- Schaffer, J. (2001a) "Causation, Influence and Effluence" *Analysis*, vol. 61.
- Schaffer, J. (2001b) "The Individuation of Tropes" *Australasian Journal of Philosophy*, vol. 79.
- Searle, J. (1959) "Determinables and the Notion of Resemblance" *Proceedings of the Aristotelian Society*, suppl. vol. 33.
- Shoemaker, S. (1980a) "Causality and Properties" reprinted in his *Identity, Cause and Mind* (1984) Cambridge: Cambridge University Press.
- Shoemaker, S. (1980b) "Properties, Causation and Projectibility" in *Aspects of Inductive Logic*, eds. L. Cohen and M. Hesse (1980) Oxford: Clarendon Press.
- Shoemaker, S. (1981) "Some Varieties of Functionalism" reprinted in his *Identity, Cause and Mind* (1984) Cambridge: Cambridge University Press.
- Shoemaker, S. (1998a) "Causal and Metaphysical Necessity" *Pacific Philosophical Quarterly*, vol. 79.
- Shoemaker, S. (1998b) "Realisation and Mental Causation" in *Physicalism and its Discontents*, eds. C. Gillett and B. Loewer (2001) Cambridge: Cambridge University Press.
- Shoemaker, S. (1999) "Self, Body and Coincidence", *Proceedings of the Aristotelian Society*, vol. 73.

- Sidelle, A. (2002) "On the Metaphysical Contingency of Laws of Nature" in *Conceivability and Possibility*, eds. T. Gendler and J. Hawthorne (2002) Oxford: Clarendon Press.
- Smith, B. ed. (1982) *Parts and Moments: Studies in Logic and Formal Ontology*, Munich: Philosophia Verlag.
- Simons, P. (1987) *Parts: A Study in Ontology*, Oxford: Clarendon Press.
- Sokolowski, R. (1968) "The Logic of Parts and Wholes in Husserl's Logical Investigations" *Philosophy and Phenomenological Research*, vol. 28.
- Steward, H. (1997) *The Ontology of Mind*, Oxford: Clarendon Press.
- Stout, G. (1921) "The Nature of Universals and Propositions" in his *Studies in Philosophy and Psychology* (1930) London: Macmillan.
- Strawson, P. F. (1985) "Causation and Explanation" in *Essays on Davidson: Actions and Events*, eds. B. Vermazen and M. Hintikka (1985) Oxford: Oxford University Press.
- Swinburne, R. (1980) "Reply to Shoemaker" in *Aspects of Inductive Logic*, eds. L. Cohen and M. Hesse (1980) Oxford: Clarendon Press.
- Swoyer, C. (1982) "The Nature of Natural Laws" *Australasian Journal of Philosophy*, vol. 60.
- Theil, U. (1997) "Epistemologism and Early Modern Debates about Individuation and Identity" *British Journal for the History of Philosophy*, vol. 5.
- Tooley, M (1987) *Causation: A Realist Approach*, Oxford: Clarendon.
- Tooley, M. (1990) "The Nature of Causation: A Singularist's Account" *Canadian Journal of Philosophy*, suppl. vol. 16.
- Vallentyne, P. (1996) "Intrinsic Properties Defined" *Philosophical Studies*, vol. 88.
- Van Fraassen, B. (1989) *Laws and Symmetry*, Oxford: Oxford University Press.
- Wiggins, D. (1980) *Sameness and Substance*, Oxford: Blackwell.
- Williams, D. (1953) "On the Elements of Being" reprinted in *Properties*, eds. D. Mellor and A. Oliver (1997) Oxford: Oxford University Press.
- Williamson, T. (1986) "Criteria of Identity and the Axiom of Choice" *Journal of Philosophy*, vol. 83.
- Williamson, T. (1990) *Identity and Discrimination*, Oxford: Blackwell.

Wright, C. (1983) *Frege's Conception of Numbers as Objects*, Aberdeen: Aberdeen University Press.

Yablo, S. (1987) "Identity, Essence and Indiscernibility" *The Journal of Philosophy*, vol. 84.

Yablo, S. (1992a) "Cause and essence" *Synthese*, vol. 93.

Yablo, S. (1992b) "Mental Causation" *Philosophical Review*, vol. 101.

Yablo, S. (1993) "Is Conceivability a Guide to Possibility" *Philosophical and Phenomenological Research*, vol.53.

Yablo, S. (1999) "Intrinsicness" *Philosophical Topics*, vol. 26.