

QUANTITY AND QUALITY: NATURALNESS IN METAPHYSICS

by

MAYA EDDON

A Dissertation submitted to the
Graduate School-New Brunswick
Rutgers, The State University of New Jersey
in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Theodore Sider

and approved by

New Brunswick, New Jersey

October 2009

ABSTRACT OF THE DISSERTATION

**QUANTITY AND QUALITY:
NATURALNESS IN METAPHYSICS**

By MAYA EDDON

Dissertation Director:

Theodore Sider

Ever since David Lewis argued for the indispensability of natural properties, they have become a staple of mainstream metaphysics. This dissertation is a critical examination of natural properties. What roles can natural properties play in metaphysics, and what structure do natural properties have? In the first half of the dissertation, I argue that natural properties cannot do all the work they are advertised to do. In the second half of the dissertation, I look at questions relating to the structure of natural properties. I argue that the metric structure of fundamental quantitative properties cannot be reduced to mereological structure, and I argue that the simplistic picture of natural properties as monadic must be abandoned in light of theories of fundamental physics.

ACKNOWLEDGMENTS

Many thanks to the members of my dissertation committee – John Hawthorne, Barry Loewer, Ted Sider, and Dean Zimmerman. All were incredibly patient and supportive throughout my years at Rutgers. I am especially grateful to Chris Meacham, who has spent countless hours reading my work and helping me improve it.

Quantity and Quality: Naturalness in Metaphysics

Table of Contents

Abstract.....	<i>ii</i>
Acknowledgements.....	<i>iii</i>
Chapter 1: Introduction.....	<i>1</i>
1.1 Natural Properties.....	<i>1</i>
1.2 Lewis’s Uses for Natural Properties.....	<i>1</i>
1.3 Assumptions.....	<i>11</i>
1.4 Roadmap.....	<i>13</i>
Part I: Too Much Work for a Theory of Universals.....	<i>15</i>
Chapter 2: Supervenience, Reduction, and Analysis.....	<i>15</i>
2.1 Three Notions of “Qualitative”	<i>15</i>
2.2 Uses for Qualitative _O : Supervenience.....	<i>20</i>
2.3 Uses for Qualitative _O : Reduction.....	<i>23</i>
2.4 Accounting for Qualitative _O	<i>26</i>
2.4.1 Divisions in Ontology vs. Divisions in Ideology.....	<i>27</i>
2.5 Analyzing Qualitative _O	<i>29</i>
2.5.1 Six Attempts.....	<i>30</i>
2.5.2 Analysis and Conceptual Possibility.....	<i>35</i>
2.5.3 Metaphysical Cointension.....	<i>38</i>
2.6 Rejecting Qualitative _O	<i>44</i>

2.6.1	The Ostrich Response.....	44
2.6.2	The Parsimony Response.....	53
Chapter 3: Humean Supervenience.....		54
3.1	Motivating Humean Supervenience.....	54
3.2	Formulating Humean Supervenience.....	62
3.3	Rejecting Either Qualitative _H or Qualitative _O	68
3.3.1	Rejecting Qualitative _H , Accepting Qualitative _O	69
3.3.2	Accepting Qualitative _H , Rejecting Qualitative _O	69
3.4	Rejecting Qualitative _H <u>and</u> Qualitative _O	72
3.4.1	An Escape Route?	75
3.4.2	Assessing the Costs.....	77
Chapter 4: Intrinsicity and Hyperintensionality.....		78
4.1	Introduction.....	78
4.2	Counterexamples and Methodology.....	79
4.3	Two Worries.....	80
4.4	The Qualitative Response.....	82
4.5	The Many Notions Response.....	87
4.6	The Spoils to the Victor Response.....	94
4.6.1	The Revised Spoils to the Victor Response.....	97
4.7	Analysis and Conceptual Possibility.....	100
4.8	Individuating Properties Hyperintensionally.....	103
4.8.1	Perfectly Natural, Qualitative _B , Qualitative _O , Qualitative _H	103

4.9	Assessing the Costs.....	106
Part II: The Structure of Fundamental Properties.....		109
Chapter 5: Temporary Intrinsic and Gauge Properties.....		109
5.1	The Problem.....	109
5.2	The Argument from Temporary Intrinsic.....	111
5.2.1	Endurantist Replies.....	112
5.3	The <i>Simpliciter</i> Argument.....	113
5.4	The <i>Simpliciter</i> Argument and the Paradigmatic Temporary Intrinsic..	114
5.5	Two Perdurantist Replies.....	116
5.5.1	Senses of “Shape”	116
5.5.2	Times and Having <i>Simpliciter</i>	119
5.6	The Fundamental <i>Simpliciter</i> Argument.....	120
5.7	Fundamental Relations Involving Objects and Times.....	122
5.7.1	Cannons, Arrows, and Quarks.....	122
5.7.2	A Perdurantist Response.....	126
5.8	Two Morals.....	128
5.8.1	Endurantism vs. Perdurantism.....	128
5.8.2	Humean Supervenience.....	129
Chapter 6: Quantitative Properties.....		130
6.1	Introduction.....	130
6.2	Resemblance as Partial Identity.....	131
6.3	The Metric Function.....	134

6.4	Point Particles.....	137
6.5	Positive and Negative Quantities.....	139
6.6	Vector Quantities.....	143
6.7	Resemblance in Different Respects.....	148
6.8	The “Fate of the Universals”	149
	Bibliography.....	151
	Vita.....	154

Chapter 1:

Introduction

1.1 Natural Properties

Since the publication of David Lewis's (1983*b*) paper, "New Work for a Theory of Universals," the distinction between natural and non-natural properties has become a staple of mainstream metaphysics. Natural properties have been employed in analyses of various notions central to philosophical inquiry. They have also been put to work in formulating metaphysical theses and clarifying metaphysical debates. For example, when we state metaphysical positions in terms of natural properties, we often find that many disputes come down to disputes about what the natural properties are like.

In this dissertation, I look at several issues regarding natural properties. I focus on the issues relevant to two questions. What work can the perfectly natural properties do? What structure do the perfectly natural properties have?

In this chapter, I lay the groundwork for these investigations. First, I describe some of the work natural properties are supposed to do. Then I lay out some assumptions and background that my discussion will presuppose. Finally, I provide a roadmap for the rest of the dissertation.

1.2 Lewis's Uses of Natural Properties

Some properties are perfectly natural, such as *mass* and *charge*. Other properties are somewhat natural, such as *cat*, *gold*, *mountain*, *table*, and so on. And others are so gruesomely gerrymandered that they are beyond our power to name.

We can capture these differences by introducing a primitive distinction between the properties that are perfectly natural and those that are not. We can then characterize the somewhat natural properties in terms of the perfectly natural. For example, we might characterize the degree to which a property is natural in terms of definitional length: the more complicated its definition in terms of predicates that refer to the perfectly natural properties, the less natural it is.¹ So the somewhat natural properties are those whose definitions are somewhat complex, while the gruesomely gerrymandered properties are those whose definitions are infinitely complex.

We have an intuitive idea of the distinction between more and less natural properties. But Lewis does not take this alone to provide sufficient justification for adopting the distinction. The full justification for positing perfectly natural properties comes from their usefulness in various areas of philosophy. Here is some of the work which natural properties are supposed to do for us:

1. *Resemblance*

What does it mean to say that two things resemble one another? It cannot merely be the fact that they share a property. According to Lewis, properties are *abundant* – there is a property corresponding to every set of possible individuals. So there is always some gerrymandered property that any two things will share, similar or not. Consider two electrons, *a* and *b*. These electrons share the property *being a or being b*. They also share the property of *having 1.6×10^{-19} coulombs charge*. But only the latter property seems to capture any objective similarity between the two.

¹ As Sider (1995) argues, definitional length may not be a good measure of naturalness. He proposes that we instead take *relative naturalness* as primitive, and then characterize *perfect naturalness* in terms of relative naturalness.

So the sharing of some properties makes for genuine resemblance, and the sharing of others does not. How can we distinguish between them? Enter natural properties. The perfectly natural properties make for objective similarity and difference. When two things share a perfectly natural property, they genuinely resemble one another in some respect.²

2. *Duplication*

We said that the sharing of perfectly natural properties makes for genuine similarity. The more perfectly natural properties two things share, the more similar they are. In the limiting case, they are perfect duplicates. So natural properties are ideally situated to characterize perfect duplication, or duplication *simpliciter*. Two mereologically simple objects are duplicates *simpliciter* iff they share all their perfectly natural properties. And we can extend this characterization to include mereologically complex objects: two objects are duplicates *simpliciter* iff they share all their perfectly natural properties and there is a one-one correspondence between their parts such that corresponding parts stand in the same perfectly natural relations and have the same perfectly natural properties.³

3. *Intrinsicity*

Intuitively, an intrinsic property is one an object has in virtue of itself. In contrast, an extrinsic property is one that an object has in virtue of its relations to other things. But how shall we capture this intuitive distinction?

² See Lewis ([1983b] 1999a, 13) and (1986b, 60). Although the sharing of perfectly natural properties captures *exact* resemblance in some respect, it's an open question whether this alone suffices to capture all the facts about resemblance. See discussion in Chapter (6).

³ See Lewis ([1983b] 1999a, 27) and (1986b, 61).

Above, we said that two objects are duplicates when each is, in itself, just like the other. And we said that an intrinsic property is one an object has in virtue of itself. So it seems that we can naturally characterize intrinsicity in terms of duplication: a property P is intrinsic *iff* for any pair of duplicates, either both have P or neither does.⁴

4. *Supervenience*

A supervenience claim is, as Lewis says, “a denial of independent variation”: A supervenes on B *iff* there can be no difference in A without a corresponding difference in B . (Lewis [1983*b*] 1999*a*, 29) Lewis takes there to be concrete possible worlds, inhabited by possible individuals. Given these assumptions, supervenience claims are usefully formulated in terms of *possibilia*. For instance, we may say A *individually supervenes* on B *iff* any pair of possible individuals alike with respect to B are alike with respect to A . And we may distinguish individual supervenience from global supervenience: A *globally supervenes* on B *iff* any pair of possible worlds alike with respect to B are alike with respect to A .

Not only are there different varieties of supervenience, but these varieties come in different strengths. In Chapter (4), the distinction between weak global supervenience and strong global supervenience will become relevant. So let us introduce these definitions here:

(WGS) A *weakly globally supervenes* on B *iff* for any pair of worlds, if there is a B -preserving isomorphism between these worlds then there is an A -preserving isomorphism.

⁴ See Lewis ([1983*b*] 1999*a*, 25-26) and (1986*b*, 62).

(SGS) *A strongly globally supervenes on B iff for any pair of worlds, every B-preserving isomorphism between these worlds is an A-preserving isomorphism.*⁵

When we formulate supervenience theses in terms of *possibilia*, we see that many supervenience claims make use of the notion of duplication *simpliciter*. For example, suppose one claims that the beauty of an object consists only in how that object is in itself, and has nothing to do with the maker of the object, or whether the object is a forgery, etc. This can be formulated as a supervenience claim cashed out in terms of duplication: any pair of objects that are duplicates *simpliciter* are duplicates with respect to how beautiful they are.

Note the implicit role of natural properties here. Without a notion of duplication *simpliciter*, we cannot formulate this kind of supervenience thesis. And to characterize duplication, we require natural properties.⁶

5. *Minimal Supervenience Base*

The perfectly natural properties comprise a *minimal supervenience base*.⁷ We can divide this claim into two parts. First, the perfectly natural properties form a supervenience

⁵ See Sider (1999) and Stalnaker (1996).

⁶ See Lewis ([1983*b*] 1999*a*, 30). Note that although I think natural properties are needed to formulate some supervenience claims, I do not believe they suffice to formulate *all* of the supervenience claims Lewis discusses in his 1983*b*. I explore this issue in Chapter (2).

⁷ The supervenience base is *global*, since it is characterized in terms of worlds: any two *worlds* alike in their supervenience base are duplicates. And we can say that this supervenience base is *strong*, since it is characterized in terms of duplication. (Every “supervenience base”-preserving isomorphism is a “duplication”-preserving isomorphism.)

base: worlds alike with respect to their arrangement of perfectly natural properties are perfect duplicates.⁸ This follows from the characterization of duplication given above.

Second, the perfectly natural properties form a *minimal* supervenience base. A supervenience base that is minimal has an additional feature: no property in the supervenience base supervenes on any other property in the supervenience base. So no perfectly natural property supervenes on any other: there are only just enough of the perfectly natural properties to “characterise things completely and without redundancy.” (Lewis 1986*b*, 60) This second part of the claim does not follow from anything we’ve said so far. But it fits the spirit of the theory, given that the perfectly natural properties are supposed to be sparse and selective.⁹

6. *Materialism*

Materialism is, roughly, the thesis that “our world is as physics says it is, and there’s nothing more to say. World history written in physical language is all of world history.” (Lewis [1983*b*] 1999*a*, 34) The sentiment is easy enough to grasp, but how shall we state it precisely?

It seems plausible that for any pair of worlds where Materialism obtains, if those worlds are alike physically then they are alike in every respect. This is a supervenience claim:

- (M) Worlds that are duplicates with respect to their physical properties are duplicates *simpliciter*.

⁸ See also Lewis ([1994*a*] 1999*a*, 225): “We may be certain *a priori* that any contingent truth whatever is made true, somehow, by the pattern of instantiation of fundamental properties and relations.”

⁹ See also Lewis ([1983*b*] 1999*a*, 12): “The world’s universals should comprise a minimal basis for characterising the world completely.”

But this only gets us halfway there. Materialism is a contingent thesis, but (M) is necessary. Perhaps our world is one where the world history written in a physical language is all of the world history, but that is not the case for all possible worlds. How shall we restrict the scope of (M)? Let us say that a perfectly natural property is *alien* to a world if it is not instantiated at that world. We can make (M) a contingent claim by restricting it to those worlds with no properties alien to the actual world. In other words, we can make (M) contingent by restricting it to the *inner sphere of possibility*:

(M*) Within the inner sphere of possibility – i.e. among worlds where no natural properties alien to our world are instantiated – any pair of worlds that are duplicates with respect to their physical properties are duplicates *simpliciter*.

(M*) seems to provide an adequate formulation of Materialism. Note again the role of natural properties. Natural properties enter into this formulation in two ways: in characterizing duplication, and in charactering the inner sphere of possibility.¹⁰

7. *Laws of Nature and Chance*

Lewis endorses a regularity analysis of laws of nature. Regularity accounts of lawhood identify the laws with certain regularities or patterns of events. But not just any regularity should count as a law. After all, some regularities are merely accidental. It might a true regularity that all solid lumps of gold have a volume of less than one cubic mile, but this regularity isn't a law.¹¹

¹⁰ See Lewis ([1983b] 1999a, 37).

¹¹ This example comes from Armstrong (1983, 17).

Lewis takes the laws to be those regularities that earn inclusion in the *best system* of the world. The best system of a world “must be entirely true; it must be closed under strict implication; it must be as simple in axiomatisation as it can be without sacrificing too much information content; and it must have as much information content as it can have without sacrificing too much simplicity.” (Lewis ([1983*b*] 1999*a*, 41))

But how shall we judge the simplicity and informativeness of candidate systems? We need to place some restrictions on the language they are stated in. Otherwise, a system may be maximally strong and simple, but only artificially so. Consider a language with the predicate F , where F applies to all and only the things at the world in question. Then the system whose only axiom is “for all x , x is an F ” will be both maximally simple and informative. This is not the result we want.¹²

Enter natural properties. The best system must be expressed in the language whose predicates refer only to the perfectly natural properties and relations. Since this language won’t include predicates like F , we avoid the worry above. So, the laws of nature are the regularities that appear in the system that, when stated in the language of natural properties, best balances simplicity and informativeness.¹³

What about chances? To account for chances, we need to consider systems that make claims not only about the history of the world in question, but also about the chances of events at that world. We then assess candidate systems with respect to an additional criterion: *fit*, a measure of the chance a system assigns to the history of the world in question. The best system is the one that achieves the best balance of simplicity,

¹² This example comes from Lewis ([1983*b*] 1999*a*, 42).

¹³ See Lewis ([1983*b*] 1999*a*, 39-43) and ([1994*a*] 1999*a*, 231-232).

informativeness, and fit. Then the laws of nature are the regularities that appear in the best system, and the chances are whatever the laws say they are.¹⁴

8. *Causation and Counterfactuals*

Natural properties play several roles in Lewis's analysis of causation. Lewis offers an analysis of causation in terms of counterfactuals; and he offers an analysis of counterfactuals that employs similarity relations among possible worlds. As we've seen, similarity between worlds is partly a matter of the perfectly natural properties they share. So Lewis's analyses of causation and counterfactuals both make use of perfectly natural properties.¹⁵

Lewis's account of causation makes use of perfectly natural properties in another way as well. Causation is a relation that holds between genuine events. But not just any goings-on should count as a genuine event. "Unless we distinguish genuine from spurious events," says Lewis, "we will be left with too many putative causes." ([1983*b*] 1999*a*, 43) Enter natural properties. The conditions of occurrence for genuine events are given by relatively natural properties; while those of spurious events are given by relatively non-natural ones.¹⁶

¹⁴ See Lewis ([1994*a*] 1999*a*, 234). This is not say that the chances are only a function of the laws. Lewis takes them to be a function of a history up to a time as well. But still, it is the laws that tell us *how* to assign the chances; i.e. they provide the history-to-chance conditionals.

¹⁵ His analysis of counterfactual dependence also makes use of laws. And as we've seen, Lewis's analysis of laws requires perfectly natural properties.

¹⁶ See Lewis (1973) and ([1983*b*] 1999*a*, 43-45).

9. *Content*

What determines the meanings of our words? We do, in part. But this constraint on reference is not enough to block massive indeterminacy in our language. For it seems that if “there is no semantic glue to stick our words onto their references, then reference is very much up for grabs.” ([1984] 1999*a*, 56)

If our language is to have a fairly determinate interpretation, as it seems it does, then we need some semantic glue. What could provide it? Enter natural properties. The meaning of our words is determined in part by our use and in part by the eligibility of that to which we refer. And the extent to which a referent is eligible depends on the naturalness of properties it instantiates.¹⁷

The tenability of Lewis’s accounts of laws of nature, chance, causation, and meaning has accounts has generated an extraordinary amount of literature. I have nothing new to add to this literature. My focus is more general: can one primitive notion do all the work Lewis wants?

1.3 Assumptions

Most of the discussion in this dissertation takes places against the backdrop of the metaphysical framework developed by David Lewis. I will assume much of this framework in what follows. I do so for two reasons. First, much of my discussion is a critical examination of some of Lewis’s claims, and a fair evaluation requires working within his framework. Second, these assumptions provide a convenient and familiar

¹⁷ See Lewis ([1983*b*] 1999*a*, 47) and Lewis (1984).

framework for approaching the issues I examine. That said, many of these assumptions are not crucial to my arguments.

I will not attempt to describe the totality of Lewis's metaphysical framework here. A comprehensive description can be found in Lewis's *On the Plurality of Worlds*. But I will briefly review some of the main Lewisian assumptions I employ.

I assume that there are possible worlds and possible individuals. Worlds are fusions of possible individuals;. Every possible individual is a part of some world, and no possible individual is part of more than one world.

I assume that there is a plurality of possible worlds. Every way the world could be corresponds to some possible world. These possible worlds and the individuals that inhabit them are "of a kind" with our own. In particular, possible individuals can be perfect duplicates of actual individuals, and possible worlds can be perfect duplicates of the actual world.¹⁸

Lewis takes there to be more than one conception of "property." For the purposes of this dissertation, I grant that there may be different conceptions of properties. But I assume that on any conception of properties relevant to the issues I discuss, properties are abundant: there is at least one property corresponding to every set of possible individuals. I shall remain neutral as to whether there are second-order properties as well as first-order properties.

I will use the following terminology. The *extension* of a property *P* at a world is the set of individuals that have *P* at that world. Properties *P* and *Q* are *coextensive* at a world *iff* they have the same extension at that world. We can think of the *intension* of a

¹⁸ Though if there are no indiscernible possible worlds, then the only duplicate of the actual world will be the actual world.

property P in a couple different ways. One way to think of the *intension* of a property P is as the function that, for every world, gives us P 's extension at that world.

Alternatively, following Lewis, we can identify the intension of a property with the set of possible individuals that have P . Properties P and Q are *cointensive* iff they have the same intension.

Consider a given conception of properties. If, according to this conception, properties are individuated by their intensions, then I will say that they are *intensionally individuated*. If properties are individuated more fine-grainedly than their intensions, then I will say that they are *hyperintensionally individuated*.

I assume there is a primitive distinction between properties that are perfectly natural and those that are not. I take the defining characteristic of the perfectly natural properties to be their role in characterizing duplication. I will not assume that perfectly natural properties have all of the other features that Lewis takes them to have.

If the perfectly natural properties characterize duplication, then it follows that the perfectly natural properties comprise a supervenience base. To say that they comprise a supervenience base is just to say that any pair of worlds alike with respect to the perfectly natural properties are duplicates. However, I shall not assume that the perfectly natural properties must comprise a *minimal* supervenience base. (In general, the issue of minimality won't have an impact on my arguments, though I briefly discuss a consequence of taking perfectly natural properties to comprise a minimal supervenience base in Chapter (4).)

Finally, I will at times want to consider the possibility of fundamental laws, primitive propensities, brute causal relations, and the like. So I will not assume that

Lewis's analyses of laws of nature, objective probabilities, causation, and so on are correct. Neither will I assume that these analyses are *not* correct. I simply leave the matter open.

1.4 Roadmap

In the first half of this dissertation, Chapters (2)-(4), I explore the work that natural properties can and cannot do for us. In the second half of this dissertation, Chapters (5)-(6), I look at questions relating to the structure of the perfectly natural properties.

In Chapter (2), I argue that perfectly natural properties cannot be used to formulate certain supervenience theses, particularly those supervenience theses that motivate Lewis's reductive analyses. Similarly, I argue that perfectly natural properties cannot be used to characterize certain kinds of reductive projects, nor can they be used to characterize certain kinds of Humean accounts of laws, chances, and the like. I show that this work can be done if we adopt an additional primitive notion. I then weigh the costs and benefits of adopting this additional primitive.

In Chapter (3), I examine a different kind of supervenience thesis: Humean supervenience. I argue that an adequate characterization of the doctrine of Humean supervenience requires a primitive notion as well. I then weigh the costs and benefits of adopting another primitive notion.

In Chapter (4), I assess Lewis's analysis of intrinsicity. I discuss the standard counterexamples to the account, and evaluate the various moves people have made to deflect these counterexamples. I argue that none succeed. I conclude by diagnosing the source of the problem: the notion of intrinsicity discriminates between properties

individuated hyperintensionally, while Lewis's account is not sensitive to these hyperintensional distinctions. Again, I weigh the costs and benefits of adopting another primitive notion.

Chapters (2), (3), and (4) together comprise an extended argument that natural properties cannot do all of the work they are advertised to do. We must decide between adopting less ambitious goals with respect to these projects, or adopting some further notions as primitive.

In Chapter (5), I look at the structure of properties that appear in our fundamental physical theories, and argue that this structure has implications for other metaphysical debates. In particular, I argue that Lewis's argument from temporary intrinsics is best understood as an argument about the structure of fundamental properties. I then evaluate this argument in light of the gauge properties of fundamental physics.

In Chapter (6), I examine quantitative properties. Here I focus on Armstrong's proposal that the fundamental quantitative properties are structural universals. Armstrong claims that the mereological structure of these universals accounts for the distinctive features of quantitative properties. I argue that the structure Armstrong proposes cannot account for these features. We need something more.

Part I: Too Much Work for a Theory of Universals

Chapter 2:

Supervenience, Reduction, and Analysis

2.1 Three Notions of “Qualitative”

In his (1983*b*), Lewis lists several supervenience claims that seem to require natural properties:

Interesting supervenience theses usually involve the notion of qualitative duplication that we have just considered. Thus we may ask what does or doesn't supervene on the qualitative character of the entire world, throughout all of history. Suppose that two possible worlds are perfect qualitative duplicates – must they then also have exactly the same distributions of objective probability, the same laws of nature, the same counterfactuals and causal relations? Must their inhabitants have the same *de re* modal properties?... We cannot even raise these questions of supervenience unless we can speak of duplicate worlds. And to do that, I have suggested, we need natural properties.

We might also ask whether qualitative character supervenes on anything less. For instance, we might ask whether global qualitative character supervenes on local qualitative character... If two worlds are local duplicates, then must they be duplicates *simpliciter*? Or could they differ in ways that do not prevent local duplication? (Lewis [1983*b*] 1999*a*, 30)

The term “qualitative” appears frequently in this passage. We can ask what supervenes on a world's “qualitative character,” both local and global; and we can ask if “perfect qualitative duplicates” need be alike in their objective probabilities, laws of nature, causal relations, etc. But what do we mean by “qualitative”?

Consider three different understandings of “qualitative.” On the first understanding, the qualitative properties are those shared by indiscernibles. Since this is the broadest of the three understandings of “qualitative,” call this sense “broad”-qualitative, or qualitative_B. The canonical qualitative_B properties are intuitively descriptive properties like *having a beard* or *being an uncle*, while the canonical non-

qualitative_B properties are haecceitistic properties like *being me*, *being David Lewis*, or *standing next to Howard*.¹⁹

Every perfectly natural property is qualitative_B, and any property that globally supervenes on the qualitative_B properties is itself qualitative_B; i.e. worlds alike with respect to their perfectly natural properties are alike with respect to their qualitative_B properties. Given Lewis's characterizations of *duplication* and *indiscernibility*, we may say that any two things that share all their intrinsic qualitative_B properties are duplicates *simpliciter*; if they share their extrinsic qualitative_B properties as well, they are indiscernible. In the special case of worlds, worlds that share all their qualitative_B properties (i.e. worlds that are duplicates *simpliciter*) are thereby indiscernible. The reason is that, according to Lewis, the perfectly natural properties and relations are instantiated by worlds and proper parts of worlds; there are no perfectly natural relations that hold *between* worlds.²⁰

The second understanding of "qualitative" is more restricted. Call this sense "occurrent"-qualitative, or qualitative_O. The canonical qualitative_O properties are intuitively occurrent or categorical – they include properties like *redness*, *squareness*, *father of*, and *five feet from*. Canonical non-occurrent properties include things like primitive propensities, brute causal dispositions, fundamental modal properties and the like.

¹⁹ In Chapter (4), I present some reasons to think that some haecceitistic properties are qualitative_B. But this complication needn't concern us here.

²⁰ See, for instance, pg. 30 of Lewis ([1983*b*] 1999*a*), where he says duplicate worlds are indiscernible. Also see pg. 225 of Lewis ([1994*a*] 1999*a*), where he writes: "if two possible worlds are discernible in any way at all, it must be because they differ in what things there are in them, or in how those things are. And 'how things are' is fully given by the fundamental, perfectly natural, properties and relations that those things instantiate."

Every perfectly natural property that is intuitively “occurrent” or “categorical” – such as *mass* or *charge* – is qualitative_O. Since the qualitative_O properties are a subset of the qualitative_B properties, it’s an open question whether worlds that are qualitative_O duplicates are thereby duplicates *simpliciter*; i.e. whether the qualitative_B properties globally supervene on the qualitative_O. If there are perfectly natural properties that are qualitative_B and not qualitative_O, then worlds may be qualitative_O duplicates but not duplicates *simpliciter*. But if all perfectly natural properties are qualitative_O, then all qualitative_O duplicates are duplicates *simpliciter*.

The third understanding of “qualitative” is the most restrictive of the three. It appears in the context of Humean supervenience, so call this sense “Humean”-qualitative, or qualitative_H. Humean supervenience is the doctrine that “all there is to the world is a vast mosaic of local matters of particular fact, just one little thing after another,” and these “local matters of particular fact” are called *qualities*. (Lewis 1986c, x)²¹

Every perfectly natural property that is intuitively “Humean” is qualitative_H, and any property that globally supervenes on the qualitative_H is itself qualitative_H. Exactly what sort of perfectly natural properties are qualitative_H is a substantive issue. Popular candidates include the fundamental properties of physics: properties like *mass* and *charge*, and maybe *flavor*, *spin*, and *quark color*. In Chapter (3), I suggest that the fundamental qualitative_H properties exclude things like qualia, spiritual properties, and emergent properties. (Such properties may be qualitative_O, however.) Among Humean worlds, worlds that are qualitative_H duplicates are duplicates *simpliciter*. But not all

²¹ The qualitative_H properties are all the qualitative_B properties instantiated at Humean worlds. In Chapter (3) we’ll see some reasons to think that the fundamental qualitative_H properties may include more than just the “qualities” – local intrinsic properties of points – that Lewis mentions.

worlds are Humean; among non-Humean worlds, qualitative_H duplicates need not be duplicates *simpliciter*.

Which of the three uses of qualitative does Lewis employ? Lewis frequently uses the term “qualitative” in the broadest sense. Here’s one example: Lewis describes perfect duplicates as intrinsically “qualitatively identical.” ([1983*b*] 1999*a*, 25) Here Lewis cannot mean qualitative_H, for one might believe that two things exactly alike in all their intrinsic qualitative_H properties need not be alike in other intrinsic ways – perhaps one has a soul and the other does not. Nor can he mean qualitative_O; one might believe that two things alike in all their intrinsic qualitative_O properties need not be alike in other intrinsic ways – perhaps one has the property of *possibly being a statue* and the other does not. So Lewis must mean qualitative_B: if two things are perfect duplicates, then they are alike with respect to all their intrinsic qualitative_B properties – *all* of them, no matter what they are.

A second example: In discussing similarity relations between worlds, Lewis writes, “Here is our world, which has a certain qualitative character. (In as broad a sense of ‘qualitative’ as may be required – include irreducible causal relations, laws, chances, and whatnot if you believe in them.)” (1986*b*, 22) Here too, “qualitative” is used in a broad sense. If two worlds are perfectly similar – i.e. they are duplicates – then they differ only in their haecceitistic, or non-qualitative_B, properties.

A third example: Langton and Lewis use “qualitative” in this broad sense when they introduce their account of intrinsicity, and stipulate that it applies only to “pure, or qualitative properties – as opposed to impure, or haecceitistic, properties.” ([1998] 1999*a*, 118)

But Lewis also uses the term “qualitative” in other ways. At one point, for instance, he writes that “worlds might differ in their qualitative character,” and adds the footnote:

Here I count causal relationships as a matter of ‘qualitative character’; it may or may not be that they supervene on qualitative character more narrowly defined, in particular on the point-by-point distribution of local qualitative character. (Lewis 1986*b*, 221)

If causal relationships are included in a world’s qualitative character whether or not they supervene on its qualitative character “more narrowly defined” (its qualitative_O or qualitative_H character), then such causal relationships must be qualitative_B. And when Lewis mentions the “point-by-point distribution of qualitative character,” it seems plausible that he means something like qualitative_H. So the first instance of “qualitative” is most plausibly understood as qualitative_B, the third instance of “qualitative” is plausibly understood as qualitative_H, and the second instance of “qualitative” can be understood as either qualitative_O or qualitative_H.

In any case, we run into trouble if we try to understand Lewis as always using “qualitative” in the broadest sense. Sometimes we would do better to understand “qualitative” in a more narrow sense – as qualitative_O or qualitative_H – allowing context and charity to indicate the appropriate meaning. Lewis’s (1983*b*) discussion of “interesting supervenience theses” is one such area. And we’ll see some other areas in the chapters to come.

There are certain kinds of work that this broad notion of qualitative cannot be employed to do, and for which we require a different notion instead. In this chapter, I focus on the need for a notion of qualitative_O; in the next chapter I focus on the need for a notion of qualitative_H. But the uses for these notions are related, and we’ll see that there is some overlap.

2.2 Uses for Qualitative_O: Supervenience

Lewis writes:

Interesting supervenience theses usually involve the notion of qualitative duplication that we have just considered... Suppose that two possible worlds are perfect qualitative duplicates – must they then also have exactly the same distributions of objective probability, the same laws of nature, the same counterfactuals and causal relations? Must their inhabitants have the same *de re* modal properties?
(Lewis [1983b] 1999a, 30)

Which notion of “qualitative” does Lewis have in mind in this passage? Let’s first suppose it is qualitative_B. Then, since qualitative_B duplicate worlds are duplicates *simpliciter*, we should understand Lewis’s “interesting supervenience theses” this way:

Suppose that two possible worlds are perfect duplicates *simpliciter* – must they then also have exactly the same distribution of objective probability, the same laws of nature, the same counterfactuals and causal relations? Must their inhabitants have the same *de re* modal properties?

If this is how we understand “qualitative,” however, then none of these supervenience claims are interesting. In fact, they’re essentially trivial.²²

Consider some different views about what laws of nature are like. Perhaps the laws of nature consist in fundamental causal powers (c.f. Shoemaker 1980). Perhaps they are necessary connections between universals (c.f. Armstrong 1983). Perhaps they are fundamental entities instantiated by worlds, or parts of worlds (c.f. Maudlin 2007). Or perhaps, as Lewis believes, the laws of nature encode regularities or patterns in the distribution of perfectly natural properties (c.f. Lewis 1983b). But however much

²² The only way two duplicate worlds may differ in their chances, laws, counterfactuals, causal relations, or modal properties is if any of these features are *external* relations that hold between worlds, or between worlds and some entities “outside” of any world. Although the possibility of these relations would make these claims non-trivial, this is clearly not what Lewis had in mind. In any case, Lewis’s conception of the perfectly natural properties as a minimal supervenience upon which all contingent truths supervene explicitly disavows the possibility of such relations. (See also footnote (3) of this chapter.)

disagreement there is over the metaphysical status of laws, there is little disagreement that if we duplicate an entire world, we duplicate the laws instantiated there.

This makes the supervenience claim in question – that qualitatively identical worlds must have the same laws – trivially true. If two worlds are qualitatively identical – “in as broad a sense of ‘qualitative’ as may be required” – then they are duplicates *simpliciter*. And since a world’s broad qualitative character includes “irreducible causal relations, laws, chances, and whatnot,” (Lewis 1986*b*, 22) any two worlds that are qualitative duplicates, or duplicates *simpliciter*, are alike in their laws.

Or consider some different views about what objective probabilities are like. Perhaps objective probabilities consist in patterns of frequencies (c.f. Lewis 1994*a*). Or perhaps they are fundamental dispositions or *propensities* (c.f. Fetzer 1982 and Miller 1994). But, again, however much disagreement there is over the metaphysical status of chances, there is little disagreement that if we duplicate an entire world, we duplicate the chances instantiated there. And again, this makes the supervenience claim in question – that qualitatively identical worlds must have the same chances – trivially true. If two worlds are qualitatively identical in the broad sense, then they are duplicates *simpliciter*, and any two worlds that are duplicates *simpliciter* are alike in their chances.

Similar things can be said about counterfactuals, causal relations, and modal properties. If the notion of qualitative in play is qualitative_B, all of these supervenience claims are essentially trivial.

Let’s next suppose the notion in play is not qualitative_B, but qualitative_O. If so, we should understand Lewis’s supervenience claims this way:

Suppose that two possible worlds are perfect qualitative_O duplicates – must they then also have exactly the same distribution of objective probability, the same laws of nature, the same counterfactuals and causal relations? Must their inhabitants have the same *de re* modal properties?

Understood this way, these supervenience claims are interesting. Suppose two worlds are alike with respect to all their qualitative_O properties; must they be alike in their laws of nature? Yes, if the laws of nature encode patterns in the arrangement of qualitative_O properties; no, if they reduce to fundamental causal powers or necessary connections. Now we have room for a substantive disagreement. It is an interesting question whether the laws of nature may vary independently of a world's qualitative_O character; i.e. whether qualitative_O duplicates must be alike in their laws.

Or suppose two worlds are alike with respect to all their qualitative_O properties; must they be alike in their objective probabilities? Yes, if objective probabilities reduce to frequencies in the pattern of qualitative_O properties; no, if they are primitive propensities. Again, we now have room for a substantive disagreement. It is an interesting question whether the objective probabilities may vary independently of a world's qualitative_O character; i.e. whether qualitative_O duplicates must be alike in their chances.

When it comes to interesting supervenience theses of the sort Lewis suggests, the notion of “qualitative” in play seems to be qualitative_O, not qualitative_B. (We could also try to understand these supervenience claims using the notion of qualitative_H. I address this in section (3.3.2).) Thus we can ask whether worlds alike in their qualitative_O character may differ in their laws, chances, causal relations, etc. or any other feature of

their qualitative_B character. And we can ask whether qualitative_O duplicates are duplicates *simpliciter*. But we cannot make sense of these questions if the notion of qualitative in play is qualitative_B.

2.3 Uses for Qualitative_O: Reduction

The supervenience theses Lewis discusses are interesting in their own right. But that is not the only reason why Lewis cares about them. Lewis is interested in certain reductive projects, and supervenience is a necessary condition for reduction. Thus Lewis writes:

A supervenience thesis seems to capture what the cautious reductionist wishes to say. Even if the reductionists ought to be less cautious and aim for translation, still it is a good idea to attend to the question of supervenience. For if supervenience fails, then no scheme of translation can be correct and we needn't go on Chisholming away in search of one. If supervenience succeeds, on the other hand, then some correct scheme must exist... (Lewis [1983*b*] 1999*a*, 29)

So claims about reduction often begin as claims about supervenience.

Lewis is especially interested in a particular kind of reduction: reduction to the qualitative_O. And the supervenience claims he discusses open the door to this sort of reduction. If two worlds are alike in their qualitative_O properties, must they be alike in their laws, chances, causal relations, and *de re* modal properties? If so, then “it makes sense to pursue such projects as a frequency analysis of probability, a regularity analysis of laws of nature, or a comparative similarity analysis of causal counterfactuals and *de re* modality.” (*ibid*, 30) These qualitative_O reductive accounts are exactly the sort of accounts that Lewis pursues.

But none of this will make sense if we take Lewis's claims about supervenience on and reduction to the qualitative to be claims about supervenience on and reduction to the qualitative_B. If two worlds are exactly alike in their perfectly natural properties, then they are duplicates – indeed, they are indiscernible. As we've already seen, indiscernible

worlds cannot differ in any way. So the supervenience of the laws, chances, causal relations, etc. on the perfectly natural properties is a trivial supervenience, one that gives us no reason to pursue any sort of reductive analyses whatsoever. The supervenience of the laws or chances on the perfectly natural is equally compatible with any of the following claims: the laws and chances are themselves perfectly natural, or the laws and chances reduce to perfectly natural dispositions or whatnot, or (Lewis's choice) the laws are regularities and the chances are frequencies. Supervenience itself gives us no reason to prefer one account over another.²³

Much of Lewis's desire for qualitative_O reduction stems from his endorsement of Humean supervenience, which I'll examine further in Chapter (3). But for now it is enough to articulate Lewis's interest in this sort of reduction:

When philosophers claim that one or another commonplace feature of the world cannot supervene on the arrangement of qualities, I make it my business to resist. Being a commonsensical fellow... I will seldom deny that the features in question exist. I grant their existence, and do my best to show how they can, after all, supervene on the arrangement of qualities. (Lewis 1986c, xi)

The "commonplace features of the world" that allegedly cannot supervene on the arrangement of qualities include laws of nature, objective probabilities, causal relations, and *de re* modal properties. And although in this context Lewis is speaking of Humean supervenience, the point generalizes. Lewis wants to show that the laws, chances, causal relations, and so on instantiated at a world *necessarily* supervene on the fundamental

²³ Of course, Lewis is correct to say that if supervenience fails, his reductive projects "are doomed from the start and we needn't look at the details of the attempts." ([1983b] 1999a, 30) But if supervenience fails, then it seems we'll have a much graver problem on our hands: since worlds with the same perfectly natural properties are duplicates by definition, then failure of supervenience entails that duplicate worlds are not duplicates! (In theory, one could escape this result if one believes that laws do not supervene on the perfectly natural, and so are not even qualitative_B. On such a view, worlds with the same perfectly natural properties but different laws of nature would be indiscernible. Nobody, to my knowledge, advocates this incredible claim.)

properties at that world.²⁴ Given any arrangement of perfectly natural properties – whether this-worldly or alien – the laws, chances, causal relations, etc. will emerge. In a sense, Lewis’s reductivist program is very grand: *every* possible qualitative_B property reduces to the qualitative_O.

Lewis is not alone in his desire for qualitative_O reduction. John Earman, for instance, argues for an *empiricist* account of laws of nature according to which, intuitively, the laws are fixed by the arrangement of occurrent (i.e. qualitative_O) properties. Earman writes:

The central empiricist intuition [is that] laws are parasitic on occurrent facts. Ask me what an occurrent fact is and I will pass your query on to empiricists. But in lieu of a reply, I will volunteer that the paradigm form of an occurrent fact is this: the fact expressed by the sentence $P(o, t)$, where ‘ P ’ is again a suitable kosher predicate, ‘ o ’ denotes a physical object or spatial location, and ‘ t ’ denotes a time. (Earman 1984, 195)

Earman freely acknowledges that he lacks an account of “occurrent” (“suitably kosher predicate” is not much of an improvement). But the reductive analysis of laws that he advocates is similar in most ways to the analysis Lewis offers.²⁵ It’s no mystery why: they are both motivated by the desire for qualitative_O reduction.

In their (2005), Earman and Roberts discuss the difficulty of characterizing what they call the “Humean base.” In essence, the properties that comprise the Humean base

²⁴ Why necessarily? In order to defend the tenability of Humean supervenience, Lewis wants to provide *analyses* of these notions that yield this kind of supervenience. Analyses, if true, are necessarily true.

²⁵ It’s an interesting question whether Lewis’s own analysis of laws requires the notion of qualitative_O. Lewis espouses a sophisticated regularity analysis of laws, according to which a regularity is a law *iff* it is a theorem of the best system of the world. The best system balances strength against simplicity, and is given in terms of predicates that express perfectly natural properties. (Lewis [1994a] 1999a, 232) This account is not obviously incompatible with there being perfectly natural properties that are qualitative_B and not qualitative_O. Perhaps there are worlds with, say, perfectly natural causal relations; if so, then according to Lewis’s analysis these causal relations will figure in the theorems of the best deductive system, and hence they will figure in the laws. So although Lewis himself does not believe there are any qualitative_B properties that are not qualitative_O, one might still think his account of laws is compatible with a very anti-Lewisian metaphysics.

are the fundamental qualitative_O properties. Earman and Roberts begin with the following intuitive idea behind the “Humean base”:

The *Humean base* may be characterized – very roughly – as the complete set of basic facts not offensive to those who are skeptical of non-logical, necessary connections in nature. The base includes particular facts about the existence of physical objects and their occurrent properties and spatiotemporal relations to one another. Excluded from the base are irreducibly general facts... and facts that involve laws of nature or other non-logical, natural modalities (e.g. facts about causal relations, counterfactuals, and irreducible dispositions). (Earman and Roberts 2005, 2)

Since Earman and Roberts are particularly interested in sophisticated regularity accounts of laws of the sort Lewis endorses, they offer the following description of how the laws are related to the Humean base on such a view:

It is the view that the laws of nature are not an independent metaphysical ingredient of the world, standing over and above the totality of more humble facts that they are supposed to govern; on the contrary, the more humble facts [the Humean base] exhaust what there is of the world, and a complete specification of them would settle everything there is to settle about what is a law of nature and what is not. (Earman and Roberts 2005, 2)

As they go on to note, this rough characterization of the Humean base and its relation to the laws is inadequate. What does Earman mean when he claims that “laws are parasitic on occurrent facts”? What does Lewis mean when he speaks of a world’s “qualitative character,” which may or may not include its laws, chances, causal relations, and so on? We cannot say. What we have is metaphor and rhetoric, but no answer to the crucial question: “What, exactly, is the criterion for deciding whether something belongs to the Humean base?” (*ibid*).

And until we have an account of “qualitative_O” (or “occurrent” or “Humean base” or “humble facts”) we can’t make sense of these reductive projects.

2.4 Accounting for Qualitative_O

We’ve seen some motivations for positing qualitative_O. First, it allows us to state interesting supervenience claims of the sort Lewis suggests. Second, we can make sense

of the general goal of occurrent or qualitative_O reduction that Earman and others explicitly endorse.

How should we meet this demand? We have three options. One option is to provide an analysis of qualitative_O in terms of other notions. Another option is to reject the notion of qualitative_O. The last option is to accept qualitative_O as primitive. The rest of this chapter is an examination of these options.

2.4.1 Divisions in Ontology vs. Divisions in Ideology

Before moving on, let me address a (perhaps) *prima facie* plausible reason why one might think that there's no pressure on the Lewisian to adopt qualitative_O as primitive. For the Lewisian, all qualitative_B properties are qualitative_O, and vice versa. So the distinction between qualitative_B and qualitative_O never arises. Offhand, one might think that if a distinction never divides anything in our ontology, then we needn't recognize it in our ideology. Since the Lewisian believes that qualitative_B and qualitative_O properties are coextensive, there's no need to recognize qualitative_O in our ideology.

This objection isn't plausible. Ideology and ontology often play off one another. Sometimes we can avoid positing ideological primitives by expanding our ontology, and sometimes we can avoid expanding our ontology by expanding our ideology instead. And in many cases there's pressure to recognize a distinction in our ideology even though it never divides anything in our ontology.

An example. Contrast the eternalist, who believes that past, present, and future exist, with the presentist, who believes that only the present exists. For the eternalist,

there are many things that exist but fail to exist in the present. For the presentist, however, what exists and what exists *now* never come apart.

The distinction between existing and presently existing never divides anything in the presentist's ontology. But surely that does not mean that the presentist can turn a blind eye to the distinction. He had better recognize the distinction somehow, even if that means positing it as primitive. And that, in fact, is what most presentists do. Even though the standard presentist does not believe the past and future exist, he includes primitive tense operators as elements of his ideology. Dinosaurs do not exist, but they *used to* exist. What the presentist lacks in ontology he makes up for in ideology.

Another example. Contrast the possibilist, who believes that non-actual worlds exist, with the actualist, who believes that only the actual world exists. For the possibilist, there are many things that exist but fail to exist in the actual world. For the actualist, however, what exists and what *actually* exists never comes apart.

The distinction between existing and actually existing never divides anything in the actualist's ontology. But surely that does not mean that the actualist can turn a blind eye to the distinction. She had better recognize the distinction somehow, even if that means positing it as primitive.

Some actualists do indeed posit at least some modal notions as primitive, while others try to analyze modal notions in terms of non-modal ones. But nobody thinks the actualist is off the hook because the distinction between what *could have been* and what *is* never divides anything in her ontology. So although the actualist denies that there are any things that possibly exist, or exist at other possible worlds, she does believe that there

are modal facts. There aren't any aliens, for instance, but there might have been. What the actualist lacks in ontology she makes up for in ideology.

The fact that, for the Lewisian, all qualitative_B properties are qualitative_O is not a reason to think that we needn't recognize a distinction between the two. There are costs to denying the distinction between qualitative_O and qualitative_B, and one isn't absolved of these costs just because the distinction never divides anything.

2.5 Analyzing Qualitative_O

For Lewis, one of the goals of philosophy is to produce theories with the greatest explanatory power and the fewest primitives. He writes: "We have the wherewithal to reduce the diversity of notions we must accept as primitive, and thereby to improve the unity and economy of the theory that is our professional concern – total theory, the whole of what we take to be true." (Lewis 1986*b*, 4)

But when must we accept a notion as primitive? The standard story goes like this. Consider all the notions you want to make use of in your theory. Each of these notions is either primitive or analyzable in terms of primitives. So we must accept a notion as primitive if we cannot provide an analysis of it in terms of the other primitive notions we accept.

This leaves a couple things up in the air. First, when considering the notions we make use of in our theory, how should we understand what it means to "make use of" a notion? Second, when providing analyses of notions, how should we understand "analysis"? I'll look more carefully at the question of what counts as an adequate analysis in section (2.5.2) and (2.5.3). In section (2.6.1), I'll return to the question of

what it means to “make use of” a notion. But let’s start by working with an intuitive notion of analysis, and examining a few different ways one might try to analyze qualitative_O.

2.5.1 Six Attempts

One sometimes sees “qualitative properties” described as “local intrinsic properties” (see Hawthorne *et al* [2004] 2006, 74). So we might try to analyze qualitative_O this way:

(Q_O1) A property p is qualitative_O iff p globally supervenes on the local properties (intrinsic qualitative_B properties instantiated at points) and spatiotemporal distance relations.

Upon reflection, however, (Q_O1) isn’t a plausible analysis of qualitative_O. First, not all local properties are qualitative_O. Primitive propensities instantiated at regions no larger than points are local properties, for instance, but they are not qualitative_O. Conversely, not all qualitative_O properties are local properties of points or spatiotemporal relations between points. If our world has fundamental quantum entanglement relations, for instance, then these relations are qualitative_O even though they are not spatiotemporal distance relations.

Second, this analysis makes trouble for the Earman (or empiricist)-style reductive accounts of laws, chances, and so on that make use of the notion of occurrent or qualitative_O.²⁶ Consider Earman’s account of laws, according to which the laws of nature are certain regularities in the pattern of fundamental qualitative_O properties. If “qualitative_O” is defined as “local,” then what should we say about possible worlds where

²⁶ Such accounts were mentioned in section (2.3).

the fundamental properties *aren't* local?²⁷ Perhaps they don't have laws. Or perhaps their laws just ignore fundamental non-local properties like quantum entanglement relations. Neither option is appealing.

Third, some of Lewis's other claims become puzzling in light of this analysis. Earlier, in section (2.2), we saw that when Lewis writes that "we may ask what does or doesn't supervene on the qualitative character of the entire world, throughout all of history," he seems to use "qualitative" to mean something like qualitative_O, not qualitative_{EB}. (Lewis [1983*b*] 1999*a*, 30) Immediately following that passage, Lewis asks the following question about supervenience on the qualitative: "We might ask whether global qualitative character supervenes on local qualitative character... If two worlds are local duplicates, then must they be duplicates *simpliciter*? Or could they differ in ways that do not prevent local duplication?" (Lewis [1983*b*] 1999*a*, 30) The notion of qualitative in play is the same as before: qualitative_O. But if we analyze qualitative_O in terms of locality, then a world's qualitative_O character simply *is* its local character. So it's trivially true that global qualitative character supervenes on local qualitative character.

Another attempt:

(Q_O2) A property *p* is qualitative_O iff *p*'s instantiation does not violate Lewis's *Principle of Recombination*, according to which "anything can coexist with anything else, at least provided they occupy distinct spatiotemporal positions." (Lewis 1986*b*, 88)²⁸

²⁷ Lewis acknowledges this possibility; see Lewis (1986*c*, *x*).

²⁸ This appears to be the characterization assumed by Bigelow *et al* (1993, 443).

This is better than (Q₀1). But the Principle of Recombination only gets us so far. First, (Q₀2) isn't available to those who would like to make use of qualitative₀ but reject Lewis's Principle of Recombination. Second, not everything compatible with recombination is thereby qualitative₀. Primitive propensities are not qualitative₀, for instance, although they don't violate recombination. To see why, first consider propensities that, intuitively, apply to the points at which they are instantiated. So consider the propensity instantiated by a point p , which assigns a chance of 0.2 to itself being occupied by a massive particle.²⁹ This propensity won't restrict recombination, because it doesn't place any constraints on what the points around p are like: a duplicate of p may exist at worlds where there are no massive particles, and at worlds where there are. Next, consider propensities that are instantiated at points and which intuitively bear on surrounding points. For instance, consider the propensity instantiated by a point p , which assigns a chance of 0.2 to there being a future neighboring point q which is occupied by a massive particle. This propensity also won't restrict combination, because it likewise doesn't place any constraints on what the points around p are like.³⁰ A duplicate of p may exist at a world where a future neighboring point is occupied by a massive particle, or a world where no future neighboring point is occupied (including worlds where there are no future neighboring points).

²⁹ Some accounts of chance, such as Lewis's (1980) account, prohibit such chances. Suppose the spacetime point is located at time t . According to Lewis, the chance at t of an event at t (the point's occupation) will always be 0 or 1. But this has no bearing on my argument.

³⁰ This is true even if the propensities are 0 or 1. Given standard probability theory, a chance of 1 does not entail that an event is inevitable, and a chance of 0 does not entail that the event won't occur. For example, the chance of a fair coin landing heads a countably infinite number of times is 0, but the event still may occur.

Let's try a different characterization. Borrowing from John Carroll (1994, 58) and Michael Tooley (1987, 29), we might try to characterize qualitative_O in terms of *nomic* properties:

(Q₀3) A property p is qualitative_O iff p globally supervenes on the fundamental non-nomic properties.

But this is unsatisfactory as an analysis of qualitative_O . First, it's unclear what counts as "nomic."³¹ Second, there seem to be properties that globally supervene on the fundamental non-nomic properties that are *not* qualitative_O . Suppose there are fundamental modal properties. If so, then these properties would not be qualitative_O . But not all modal properties are nomic; the modal property of *possibly being a hard-boiled egg*, for instance, seems to be a non-nomic property. If modal properties are fundamental, then properties like *possibly being a hard-boiled* are not qualitative_O , even though they are non-nomic.³²

Another characterization builds on a suggestion from Brian Ellis (2001, 4-5)

(Q₀4) A property p is qualitative_O iff p globally supervenes on the fundamental non-modal properties.

But what counts as a modal property? If properties are abundant, then every set of possible objects corresponds to at least one property. Which of these sets should count as "modal," and which should not? It's true that lawhood, chance, and causation are generally taken to be modal features, but it's unclear why this classification is deserved.

³¹ Is any property involved in the laws a *nomic* property? Then we must say that properties like *mass* and *charge* are nomic, in which case (Q₀3) entails that an Earman-style account of laws is self-undermining. Any property that appears in the laws is nomic, and therefore is *not* qualitative_O ; but if it's not qualitative_O then it cannot appear in the laws.

³² See also pg. 11-12 of Earman and Roberts (2005).

It seems we generally take a property to be modal if its analysis require an appeal to modality. But if that's our characterization of "modal," then (Q_O4) is trivial. Properties that have no analysis – that are fundamental – will count as non-modal, even if those properties are intuitively modal ones such as fundamental causal relations, laws, or chances. And if we take the modal properties to be those that constrain possibilities, we end up with something like (Q_O2) again. So it's unclear what distinguishes modal properties from non-modal ones. And if we cannot make this distinction, then (Q_O4) fails as an analysis of qualitative_O.

We might try to build on a suggestion from Barry Loewer (2004) to characterize qualitative_O:

(Q_O5) A property p is qualitative_O iff p globally supervenes on the *categorical* properties, and a property is categorical "just in case its involvement in the laws is not essential to it." (Loewer 2004, 1118)

This also seems unsatisfactory as an analysis of qualitative_O, since not all categorical properties supervene on the qualitative_O. Again, suppose the property of *possibly being a hard-boiled egg* is a fundamental modal property. Involvement in the laws of nature doesn't seem to be essential to *possibly being a hard-boiled egg*. So *possibly being a hard-boiled egg* is categorical in Loewer's sense, but it is not qualitative_O.

Let's try a different angle. Suppose we just try to characterize qualitative_O in this way:

(Q_O6) A property p qualitative_O iff p is qualitative_B.

We know how to analyze qualitative_B in terms of the perfectly natural: a property is qualitative_B iff it globally supervenes on the perfectly natural properties. So if we analyze

qualitative_O in terms of qualitative_B, we don't need to posit any primitive notion other than perfect naturalness.

What should we think of this analysis? There appear to be some obvious counterexamples. For instance, fundamental laws of nature, if there were any, would be qualitative_B and not qualitative_O. So the analysis cannot be right, since it's not true that a property is qualitative_O *iff* it is qualitative_B.

But maybe that was too quick. Suppose the proponent of this analysis responds in the following way. For Lewis, everything (everything non-haecceitistic, that is) ultimately reduces to qualitative_O properties. There are no fundamental laws of nature, chances, causal relations, modal properties, etc. So once we factor in all of Lewis's metaphysical commitments, we see that qualitative_B and qualitative_O properties are necessarily coextensive. Since they are coextensive across all metaphysically possible worlds, we can analyze one in terms of the other. So the Lewisian can invoke the notion of qualitative_O without positing it as primitive. (One might make similar moves for some of the other analyses above. For instance, one could defend (Q_O3) by claiming that fundamental modal properties are impossible, and thus that they cannot be raised as counterexamples to the proposed analysis.)

2.5.2 Analysis and Conceptual Possibility

To assess this response, we need to return to the issue of what counts as an analysis. For (Q_O6) to be tenable, we require a conception of analysis that is something along the following lines: if, for all metaphysically possible worlds, a notion *a* applies to all and only the things that are *b*, then we can analyze *a* in terms of *b*.

But this conception of analysis isn't plausible. For convenience, let's employ the following terminology. If it's metaphysically necessary that something is *F* iff it is *G*, then let us say that *F* and *G* are *metaphysically cointensive*. If it's conceptually necessary that something is *F* iff it is *G*, then let us say that *F* and *G* are *conceptually cointensive*.

In the next section, I argue that metaphysical cointension is not a sufficient condition for analysis. I also argue for the more controversial claim that metaphysical cointension is not a necessary condition for analysis, either. I maintain that the sense of possibility we should use when evaluating candidate analyses is *conceptual possibility*, not metaphysical possibility. If we can conceive of counterexamples to a proposed analysis, then that is a mark against the analysis – whether or not the counterexamples are metaphysically possible.

But this claim requires some caveats. First, not just any conceptually possible counterexample is a mark against an analysis. Following Chalmers (2002), we can distinguish *prima facie* conceivability from *ideal* conceivability. Not all *prima facie* possibilities survive ideal rational reflection; when they do, we can *ideally* conceive of those possibilities. It's the space of ideally conceivable possibilities that matters when evaluating analyses, not the space of *prima facie* conceivable possibilities.³³

Second, when assessing analyses, we need to allow for vagueness. In this I follow Lewis:

We should allow for semantic indecision: any interesting analysandum is likely to turn out vague and ambiguous. Often the best that any one analysis can do is to fall safely within the range of

³³ Chalmers (2002) discusses two other axes of conceivability: positive vs. negative and primary vs. secondary. A situation *S* is positively conceivable if we can “form some sort of positive conception of a situation in which *S* is the case,” and negatively conceivable if “*S* is not ruled out *a priori*, or when there is no (apparent) contradiction in *S*.” (Chalmers 2002, 150) A situation *S* is primarily conceivable when it is conceivable that *S* is actual, and secondarily conceivable when *S* conceivably might have been the case. (Chalmers 2002, 157) Because these two axes of conceivability do not play much of a role in the sorts of analyses I am interested in, I will not address them.

indecision. And we should allow for semantic satisficing: analysis may reveal what it would take to deserve a name perfectly, but imperfect deservers of the name may yet deserve it well enough. (Lewis [1994] 1999a, 298)

In the following discussion, I will generally leave issues involving vagueness aside.

Third, I take the force of conceptual counterexamples to come in degrees. An analysis of a notion *a* can be subject to such counterexamples if it doesn't apply in cases in which we can conceive of it applying. Likewise, an analysis can be subject to such counterexamples if it applies in cases in which we cannot conceive of it applying. But in either case, the more an analysis deviates from what is ideally conceptually possible with respect to the notion being analyzed, the greater the mark against the account. It's a matter of degree.

Most interesting analyses will deviate to some extent. For whenever it makes sense to ask whether a given analysis is correct, there is a sense in which we conceive of the analysis failing to be correct. Such analyses are *unobvious* analyses. As Lewis writes:

Whenever it is analytic that all A's are B's, but not obviously analytic, the Moorean open question – whether all A's are indeed B's – is intelligible. And not only is it intelligible in the sense that we can parse and interpret it... but also in the sense that it makes sense as something to say in a serious discussion, as an expression of genuine doubt. (Lewis [1989] 2000, 85-86)

I follow Lewis's method for evaluating analyses. In defending his unobvious analysis of "value" in terms of dispositions to value, Lewis asks, "Why think that it's analytic at all? – Because that hypothesis fits our practice... it does seem that if we try to find out whether something is a genuine value, we do try to follow – or rather, approximate – the canonical method. We gain the best imaginative acquaintance we can, and see if we then desire to desire it." (Lewis [1989] 2000, 87) The canonical method aims for the "fullest possible imaginative acquaintance that is humanly possible." (*ibid*, 77) In the case of value, we begin by imagining every conceptually possible case of

valuing that we can – we imagine “vividly and thoroughly how it would be if these putative values were realised (and perhaps how it would be if they were not)...” (*ibid*, 77) So the canonical method asks us to imagine what it would be like if we were *not* disposed to value the things that are valuable. If the analysis we ultimately arrive at says that something is a value *iff* we are disposed to value it, then the canonical method asks us to imagine not just what is metaphysically possible, but also what is metaphysically *impossible*.

So when it comes to constructing analyses, we do not merely consider the space of metaphysical possibility; we consider the space of conceptual possibility. And the greater the extent to which the analysis diverges from what’s ideally conceivable, the greater the mark against the account.

Note: I am not claiming that divergence from what’s ideally conceivable is the *only* requirement for an adequate analysis. There may be other desirable features; for instance, we may desire an analysis that is pleasingly simple, or avoids commitment to mysterious entities, etc. Similarly, we may not desire an analysis that matches our intuitive notion in a trivial way by, say, analyzing a notion *a* as “whatever we can ideally conceive of as satisfying *a*.” But divergence from conceptual possibility is an important requirement, and the one I focus on here.

2.5.3 Metaphysical Cointension

Let’s first see why metaphysical cointension is not a sufficient condition for analysis. Suppose I offer the following analysis of *being a human*: *x* is a human *iff* *x* is mortal and rational. *Prima facie*, there might be Martians who are mortal and rational, although they

are not human – a counterexample to my proposed account. But suppose that for independent reasons I believe that God necessarily exists and it is necessarily true that the only mortal and rational beings God would allow to exist are humans. Then this counterexample no longer applies. (There are other counterexamples, of course, such as humans that are not rational. But put that aside. It's the tenability of this kind of defense, not the particular analysis, which we're concerned with.)

Given this restriction on metaphysical possibility, it's true that something is a human *iff* it is mortal and rational; i.e. there aren't any metaphysically possible counterexamples to the proposed analysis.³⁴ But surely the analysis is inadequate. So metaphysical cointension is not a sufficient condition for analysis.

What exactly is wrong with this analysis? The problem is that we can ideally conceive of things that are mortal and rational but not human – things like Martians or Venusians or talking donkeys or hobbits. Given the space of ideal conceptual possibility, the set of things that the analysis calls “human” includes many things that are *not* human. Because the degree of divergence is so large, the analysis fails.

Another example. Suppose I offer an analysis of duplication as the sharing of physical properties of the actual world. There seems to be an obvious counterexample to my analysis: *prima facie*, there could be duplication at possible worlds where the physical properties differ from those at the actual world. Now suppose that, for independent reasons, I believe that there are no worlds with different physical properties. (Perhaps I believe God would not have been so cruel as to allow cases of duplication that cause trouble for my analysis.) By restricting the space of metaphysical possibility in this way,

³⁴ Again, putting other kinds of counterexamples aside.

I needn't worry about counterexamples arising from non-physical duplication. But, again, this is not an adequate analysis.

Why not? Because we can ideally conceive of cases of duplication that do not involve the physical properties of our world; thus we can ideally conceive of counterexamples to the proposed account.³⁵ In the previous case, the analysis included things in the conceptual intension of "human" that did not belong. In this case, the analysis excludes things in the conceptual intension of "duplicate" that *do* belong. But the problem is the same. This analysis fails because it deviates to a large extent from what we ideally conceive of as *duplication*.

So far, we've seen that metaphysically necessary coextension is not a sufficient condition for analysis. Next, let's see why it's also not a necessary condition.

Consider Langton and Lewis's (1998) account of intrinsicity according to which a "property *p* is intrinsic" is analyzed as "any pair of duplicates are alike with respect to *p*." "Duplication" is in turn analyzed as "two things that share all their basic intrinsic properties," and "basic intrinsic property" is analyzed as "(1) independent of accompaniment or loneliness of other contingent objects, (2) not disjunctive, and (3) not a negation of disjunctive properties."

Langton and Lewis discuss two counterexamples to their account. The first arises from strong laws, or laws that necessarily obtain. If laws are necessary, then the space of metaphysical possibility is the same as the space of nomological possibility. Given this restriction, their account of intrinsicity appears to deliver incorrect results. One example they discuss is the apparently intrinsic property of *being an ellipsoidal star*. If

³⁵ Those who worry that we cannot conceive of such cases of duplication should remember that people who know nothing about the world's physical properties have no trouble understanding the notion of duplication. See Lewis ([1983*b*] 1999*a*, 27).

the laws are such that a star can be ellipsoid only when accompanied by some other massive object, then the property of *being an ellipsoidal star* is not intrinsic on their account, because it is not independent of accompaniment or loneliness. So if we think laws are strong, and we think that *being an ellipsoidal star* is intrinsic, then we have trouble: when assessed with respect to metaphysical (i.e. nomological) possibility, we find that some apparently intrinsic properties are *not* independent of accompaniment or loneliness of other contingent objects. In other words, we find that some apparently intrinsic properties are not metaphysically cointensive with any of the properties picked out by Langton and Lewis's analysis of intrinsicity. So either we reject their account of intrinsicity, or we reject the claim that metaphysical cointension is a necessary condition for a satisfactory analysis.

Langton and Lewis grant that if one thinks that laws are strong and *being an ellipsoidal star* is intrinsic, then what's intrinsic isn't metaphysically cointensive with the properties picked out by Langton and Lewis's analysis. But that doesn't mean that their account is incorrect, say Langton and Lewis. Rather, it means that metaphysical cointension isn't a necessary condition for an analysis. When assessing the tenability of a candidate analysis, metaphysical possibility isn't the relevant notion of possibility:

If a theory of strong laws is to be credible, it had better provide not only a sense of 'possible' in which violations of laws are impossible, but also another sense in which violations of laws are possible... Friends of strong laws might think it a hoked-up, artificial sense. But no harm done, provided they acknowledge the possibility of lonely ellipsoidal stars, or whatnot, in some sense or other... It is this sense of possibility, whatever it may be, that a friend of strong laws should use in defining 'intrinsic'.
(Langton & Lewis [1998] 1999a, 122)

I agree. The sense of possibility that a friend of strong laws should use in defining "intrinsic" isn't metaphysical possibility – it's something like conceptual possibility. We can ideally conceive of violations of laws, and we often do so when considering interesting counterfactual claims. So even if violations of laws aren't metaphysically

possible, they are conceptually possible. And Langton and Lewis will maintain that intrinsic properties are *conceptually* cointensive with those picked out by their account.³⁶ So strong laws are not a problem for their account.

The second counterexample Langton and Lewis consider arises from the existence of a (metaphysically) necessary God. If God necessarily exists, then the apparently extrinsic property *being created by God* is independent of accompaniment or loneliness of any contingent being, and so is intrinsic on their account. So if we think God necessarily exists, and we think that the property of *being created by God* is extrinsic, then we have trouble: when assessed with respect to metaphysical possibility, we find that some apparently extrinsic properties *are* independent of accompaniment or loneliness of other contingent objects. In other words, we find that some extrinsic properties are not metaphysically cointensive with any of the properties picked out by Langton and Lewis's analysis of extrinsicality. So either we reject their analysis of intrinsicity, or we reject the claim that metaphysical cointension is a necessary condition for a satisfactory analysis.

Again, Langton and Lewis's response amounts to rejecting the claim that metaphysical cointension is a necessary condition for a satisfactory analysis. If one thinks that God exists as a matter of metaphysical necessity and *being created by God* is extrinsic, then what's extrinsic isn't metaphysically cointensive with the properties that are not independent of accompaniment or loneliness of other contingent objects. But that doesn't mean that their account is incorrect, say Langton and Lewis. Rather, it means that metaphysical cointension isn't a necessary condition for an analysis. Again, when assessing the tenability of a candidate analysis, metaphysical possibility isn't the relevant

³⁶ I offer some reasons to doubt this claim, however, in Chapter (4).

notion of possibility: “Perhaps God’s existence may be supposed to be necessary in some sense. Yet in a second sense, it still might be contingent... it is the second sense and not the first that should be used in defining ‘intrinsic.’” (Langton & Lewis [1998] 1999a, 123) We can ideally conceive of God not existing. So even if his nonexistence isn’t metaphysically possible, it is conceptually possible. And Langton and Lewis will maintain that intrinsic properties are conceptually cointensive with those picked out by their account. So a metaphysically necessary God is not a problem for their account.

In the cases above, the space of metaphysical possibility is constrained in certain ways so that it’s metaphysically impossible that God created Martians, or that there are different physical properties, or that laws may be violated, or that God might not exist. As we saw, however, we can nonetheless conceive of things that are metaphysically impossible – that Martians could exist, that physical properties could differ, that laws may be violated, or that God might not exist. And it’s conceptual possibility that we should use when evaluating the merits of an analysis, not metaphysical possibility.

With this in mind, let’s return to the analysis of qualitative_O as (Q_{O6}): a property is qualitative_O *iff* it is qualitative_B. Since a property is qualitative_B *iff* it globally supervenes on the perfectly natural, we can rephrase (Q_{O6}) in the following way: a property is qualitative_O *iff* it globally supervenes on the perfectly natural properties.

Given Lewis’s other metaphysical commitments (such as his regularity account of laws, etc.), (Q_{O6}) turns out to be true. There are no qualitative_O properties that are not qualitative_B. But that doesn’t make (Q_{O6}) an *analysis* of qualitative_O. As we’ve seen, metaphysically necessary coextension is neither a sufficient nor a necessary condition for analysis. More importantly, the analysis of qualitative_O as (Q_{O6}) is widely divergent

from what we ideally conceive of as qualitative_O. We can conceive of fundamental laws, primitive propensities, brute causal dispositions, and the like. So (Q_O6) is not an adequate analysis.

2.6 Rejecting Qualitative_O

In section (2.5.1), we saw that a number of attempts to analyze qualitative_O – (Q_O1)– (Q_O5) – fail. If my claims regarding the correct ways to analyses in this context, then (Q_O6) fails as well: we cannot analyze qualitative_O in terms of the perfectly natural properties. Given this, the prospects for analyzing qualitative_O look grim.

Let's look at some other options. What if we reject the notion of qualitative_O? Those who reject the notion of qualitative_O may take one of two routes. First, they may claim that, contrary to appearances, the Lewisian incurs no cost by rejecting the notion of qualitative_O. Call this the “Ostrich Response.” Second, they may argue that rejecting qualitative_O incurs a cost, but this cost is outweighed by the cost of positing an additional primitive notion. Call this the “Parsimony Response.” In the next section, I argue that the Ostrich Response is untenable. In the section after that, I assess the Parsimony Response.

2.6.1 The Ostrich Response

According to the Ostrich Response, the various uses we've seen for qualitative_O aren't important to the Lewisian. The Lewisian doesn't need to make any claims about supervenience on the qualitative_O or reduction to the qualitative_O, nor does he need to provide analyses in terms of the qualitative_O. After all, given the Lewisian's ontology,

what else could anything supervene on or reduce to *other* than the qualitative_O? The philosophers who need to make such claims are those whose ontologies *don't* settle the matter. The philosophers who require the notion of qualitative_O are those who believe that not all qualitative_B properties supervene on the qualitative_O. But the Lewisian is not in this position, and so he does not need to recognize qualitative_O.

In assessing the Ostrich Response, we need to return to the second question posed in section (2.5): when do we need to “make use of” a notion? Let us say that advocates of a given view make use of a notion if they employ it in any analysis, argument, or claim they want to make. Those who offer the Ostrich Response claim that there is no analysis, argument, or claim that the Lewisian wants to make that requires qualitative_O. In this section, I argue that this is false. We will see that if the Lewisian rejects qualitative_O, his theory is severely limited in its expressive power.

To get a feel for the cost of rejecting qualitative_O, let's first consider a different case – a case where someone does not understand modal notions like “actual” and “possible.” Suppose we have three philosophers. The first philosopher, Possibilia, is a possibilist. Possibilia believes that there are possible worlds distinct from the actual world. The second philosopher, Actualia, is an actualist. Actualia believes that the actual world is all there is. Possibilia and Actualia have a substantive disagreement: they disagree about whether possible things exist.

The third philosopher, Puzzledono, does not understand modal locutions like “actual” and “possible.” As a result, we cannot say where Puzzledono stands in the debate. For suppose Possibilia and Actualia ask Puzzledono whether unicorns exist, and Puzzledono says he thinks that they do. How should we interpret this response? One

could argue that Puzzledono is a possibilist: since unicorns don't actually exist, anyone who says that they exist must think they exist at some other possible world. But one could also argue that Puzzledono is an actualist: since only actual things exist, anyone who says that unicorns exist must mistakenly believe that unicorns are actual (or perhaps Puzzledono knows something we don't). Which is it? We cannot say.

At first glance, it might seem that some further questions could settle the matter. It might seem, for example, that we can settle the issue by asking Puzzledono whether unicorns are spatiotemporally related to us. But that won't help. For suppose Puzzledono says that unicorns are not spatiotemporally related to us. This claim is compatible with both possibilism and actualism. A possibilist could believe that these unicorns exist at a different possible world. An actualist could believe that unicorns exist in "island universes." I.e. she might believe that parts of a single world need not bear any fundamental relations to one another, and she might believe that these other parts of the world are where unicorns exist (e.g. see Bricker 2001). And we cannot rule out this possibility by asking Puzzledono whether the areas inhabited by unicorns belong to the actual world or a possible world, since he does not understand what "actual" and "possible" mean. Nor can we ask him whether island universes are possible, since the distinction between an island universe and a possible world requires a grasp of "actual" and "possible."

Now suppose it turns out that Puzzledono's ontology matches the ontology of Actualia, who takes modality as primitive and who believes there are no island universes. If so, can we say that, according to Puzzledono, possible things don't exist – even though he cannot make that claim himself? No. For even if Puzzledono's ontology matches that

of Actualia, it's still compatible with possibilism. For instance, there's nothing in Puzzledono's ontology that rules out the following possibilist view: Puzzledono is a solipsist, and everything else only possibly exists. So we still cannot classify him as a possibilist or an actualist.³⁷

Here is the upshot: because Puzzledono does not understand modal locutions, his theory lacks expressive power. As a result, he cannot coherently take a position regarding the structure of modal reality. So, for example, he cannot say whether he is an actualist or a possibilist. And neither can we.

Furthermore, Puzzledono cannot provide certain kinds of principled reasons for accepting the ontology that he does. To see why, let's suppose that everyone involved (who understands such talk) accepts a particular principle of recombination with respect to possibility – one which, given what they believe actually exists, entails that there could

³⁷ One might think that there are still some cases in which *we* can justifiably classify Puzzledono as a possibilist or an actualist. Suppose Puzzledono thinks that the only things that exist are spatiotemporally related to him. If we make the further assumption that two things belong to the same world if they bear fundamental relations to one another, then it seems we can classify Puzzledono as an actualist (even if he cannot say so himself).

First, even if this is true, it doesn't affect my main point: Puzzledono will suffer from lack of expressive power, and an inability to offer certain kinds of principled justifications for his beliefs.

Second, if this assumption is analytically true, then it seems right that we can classify him as an actualist. But there is reason to doubt that it's analytic that things that bear fundamental relations to one another have to belong to the same world. For example, a trans-world identity theorist who takes identity to be a perfectly natural relation (not a logical relation) will think that there are fundamental relations between possible worlds. And whatever the merits of this view, it doesn't seem that the view is *analytically* false.

Third, if the assumption is not analytic, then we cannot justify classifying Puzzledono as an actualist. It's the assumption – that two things belong to the same world if they bear fundamental relations to one another – that does all the work. And we cannot tell whether Puzzledono would accept this assumption.

Consider an analogy. Suppose that we think that all bunny rabbits are actual. And suppose Puzzledono thinks that the only things that exist are bunny rabbits. Then it seems we can classify Puzzledono as an actualist. But unless we can justify attributing the assumption to Puzzledono in some way, we cannot justify this classification. After all, there's always *some* assumption we can cook up that will allow us classify Puzzledono as a possibilist or an actualist. (For that matter, we could just assume that anyone who claims not to understand modal locutions is an actualist!)

The way we justify our assumptions about individuating possible worlds is by using our pretheoretic understanding of what's possible and what's actual. But Puzzledono doesn't have any pretheoretic understanding of this sort. So there's no way to adjudicate which assumptions about possibility he would accept *were* he to understand modal locutions. And thus, there's no way to justify using this assumption to classify him as an actualist.

have been non-actual things like talking donkeys and blue swans. Further suppose that everyone involved (insofar as they understand such talk) agrees that there are no actual talking donkeys or blue swans, and agrees that there are actual ordinary donkeys and white swans.

Now suppose we ask each philosopher – Possibilia, Actualia, and Puzzledono – whether he or she believes talking donkeys exist, and why. Possibilia has a principled answer: Given the things that actually exist and the principle of recombination, it follows that talking donkeys could have existed. To say that something could have existed is to say that it exists at some possible world. Since anything that exists at a possible world exists, talking donkeys exist. Actualia also has a principled answer: Talking donkeys don't actually exist. Since all and only things that exist actually exist, talking donkeys don't exist.

But what can Puzzledono say? Suppose he thinks there are no talking donkeys. Whereas Actualia can give a principled justification for this belief, Puzzledono can offer no similar justification. He can only restate his belief, perhaps with more emphasis. Likewise, suppose Puzzledono thinks that talking donkeys *do* exist. Whereas Possibilia can give a principled justification for this belief, Puzzledono has no similar justification. He simply accepts some things and rejects others, and it seems he does so in an entirely *ad hoc* fashion.

I doubt that many would find Puzzledono's position appealing. His theory may be internally coherent, but it lacks the expressive power to say anything about possibility.

Now let's return to the original question: is the Lewisian who rejects qualitative_o any worse off than the Lewisian who does not? This time, let's consider three

philosophers arguing about which properties there are.³⁸ The first philosopher, Largessia, believes that there are qualitative_B properties that are not qualitative_O. The second philosopher, Diminutia, believes that all qualitative_B properties are qualitative_O. Largessia and Diminutia have a substantive disagreement: they disagree about what sorts of properties there are.

The third philosopher, Perplexedono, does not understand “qualitative_O”. As a result, we cannot say where Perplexedono stands in the debate. And it’s no use asking Perplexedono whether he thinks there are qualitative_B properties that are not qualitative_O, since he doesn’t know which properties are which.

Furthermore, we cannot settle the matter by asking Perplexedono further questions, such as questions about particular properties. Suppose Largessia and Diminutia ask Perplexedono whether he thinks primitive propensities exist. And suppose Perplexedono says that he thinks they do. How should we interpret this response? One could argue that Perplexedono agrees with Largessia: since primitive propensities are qualitative_B but not qualitative_O, anyone who says they exist must think that not all qualitative_B properties are qualitative_O. But one could also argue that Perplexedono agrees with Diminutia: since all qualitative_B properties are qualitative_O, anyone who says there are primitive propensities must (mistakenly) believe that they are qualitative_O. Which is it? We cannot say.

One might think that it’s true in virtue of the meaning of “qualitative_O” that primitive propensities are not qualitative_O. If so, one might object that if Perplexedono believes there are primitive propensities, then, whether he knows it or not, he is

³⁸ I’m assuming that what properties there are bears on one’s ontology, whether these properties are universals, tropes, classes, or sets. Those who deny this can rephrase the discussion in a nominalistically acceptable way.

committed to there being qualitative_B properties that are not qualitative_O. But this objection is mistaken. It's true that in introducing the distinction between qualitative_B and qualitative_O, I used primitive propensities as a likely example of something that is qualitative_B and not qualitative_O. But whether a property is qualitative_O is a substantive question, akin to the question of whether a property is perfectly natural.

Compare: when Lewis introduces the notion of perfect naturalness, he often cites *mass* and *charge* as likely examples of perfectly natural properties (for instance, see Lewis [1984] 1999a, 66). But it is no part of the meaning of perfectly natural that *mass* and *charge* are perfectly natural. They might not be perfectly natural. After all, physics is full of surprises, and even if physics sticks with *mass* and *charge*, we've no guarantee that our physics has got it right, or that our world is one where *any* physical properties are perfectly natural.

One can believe that there are perfectly natural properties, but not know whether any particular property is perfectly natural. Likewise, one can believe that some things possibly exist and some things actually exist, but not know whether any particular thing is possible or actual. And finally, one can believe that some properties are qualitative_O and some are not, but not know whether any particular property is qualitative_O.

So, as with the case above, there is nothing that Perplexedono can say that will not be compatible with both Largessia's view and Diminutia's view.³⁹

³⁹ One might worry that the implicit assumption that all parties can state what's in their ontology, coupled with the possibility of having qualitative_B and qualitative_O as parts of one's ontology, will make trouble for this claim. It won't, though the details depend on how we treat second-order predicates such as qualitative_B, qualitative_O, perfectly natural, and so on. First, suppose we take qualitative_B, qualitative_O, perfectly natural, and so on to be parts of our ideology, not ontology. Then qualitative_B, qualitative_O, perfectly natural properties, etc. are like primitive modal operators, and the situation is exactly analogous to the case involving Possibilia, Actualia, and Puzzledono discussed earlier. And for analogous reasons, there will be nothing that Perplexedono can say that will not be compatible with both Largessia's view and Diminutia's view.

Because Perplexedono does not understand “qualitative_O,” his theory lacks expressive power. As a result, he cannot say whether he believes all qualitative_B properties supervene on the qualitative_O. And neither can we.

Furthermore, he cannot provide certain kinds of principled reasons for believing that some properties are instantiated and not others. To see why, let’s suppose that everyone involved (who understands such talk) agrees that qualitative_O properties satisfy some empiricist criterion that philosophers like Earman endorse. Further suppose that everyone involved (insofar as they understand such talk) agrees that primitive propensities do not satisfy that criterion, and so are qualitative_B but not qualitative_O. Does Diminutia believe there are primitive propensities? No – and she can give a principled justification for her belief. Primitive propensities don’t satisfy the empiricist criterion, so they would be qualitative_B and not qualitative_O. Since all qualitative_B properties are qualitative_O, there are no primitive propensities.

Now suppose that Perplexedono agrees with Diminutia that there are no primitive propensities. Unlike Diminutia, he can give no similar justification for his belief. He can only restate it, perhaps with more emphasis.

The same goes for brute modal properties, causal relations, and other things that are qualitative_B and not qualitative_O. In each case, Diminutia can give a principled reason

Second, suppose we take qualitative_B, qualitative_O, perfectly natural properties, etc. to be second-order properties that appear in our ontology. If we think there are second-order properties, then presumably we will take them to be abundant (just as we are assuming abundance for first-order properties). But then, what second-order properties one thinks there are will just depend on what first-order properties one thinks there are. And if we cannot classify Perplexedono by appealing to the first-order properties he adopts, we won’t be able to classify him by appealing to second-order properties either, since the second-order properties just comprise every possible permutation of the first-order ones. (If Perplexedono adopts abundantly many second-order properties, won’t he think that qualitative_B, qualitative_O, perfectly natural properties, etc. exist as well? In a sense. But he won’t know which second-order properties they correspond to, nor how to make sense of the claim that a particular second-order property corresponds to, say, the qualitative_O properties. So this won’t be any help.)

for why she accepts certain properties and rejects others. But Perplexedono can give no similar reason. He simply accepts some things and rejects others, and it seems he does so in an entirely *ad hoc* fashion.

If the Lewisian rejects the notion of qualitative_O , then he is in the same position as Perplexedono. And although it seems that the Lewisian would want to ally himself with those who believe that qualitative_O and qualitative_B never come apart, he lacks the ability to do so if he rejects the notion of qualitative_O . True, he can say that the chances, laws, causal relations, modal properties, etc. supervene on the perfectly natural properties. But in that he agrees with everyone who accepts an egalitarian conception of properties – including those who believe in fundamental propensities, laws of nature, causal relations, and the like.

In sum, the Lewisian who rejects the notion of qualitative_O is in the same boat as the philosopher who rejects the notion of “possibility.” Neither one has the expressive resources to state his own view. Consequently, neither one can distinguish his view from (apparently) rival views, and neither one can provide principled justifications – of the sort available to those whose theories are expressively adequate – for accepting the ontology that he does.

The Ostrich Response holds that the Lewisian has nothing to gain by making use of qualitative_O . I’ve shown that this claim is false – the Lewisian has plenty to gain. So the Ostrich Response is untenable.

2.6.2 The Parsimony Response

Those who reject qualitative_O are left with one option. They may argue that rejecting qualitative_O incurs a cost, but this cost is outweighed by the cost of positing an additional primitive notion. This is the Parsimony Response.

This is a tenable position. But we must be careful not to diminish the costs involved. As we've seen, the Lewisian who adopts the Parsimony Response lacks the expressive resources to state his own view. He can neither distinguish his view from rival views nor provide the kind of principled justifications for accepting the ontology that he does that are available to those whose theories are expressively adequate.

This is a significant cost for the Lewisian. One feature that unites philosophers in the Lewisian tradition is a general commitment to the sort of reductive project that Lewis spent much of his philosophical career advancing: reduction to the qualitative_O. But if we reject the notion of qualitative_O, then we cannot even state what this reductive project is.

A further cost of rejecting any primitive notion besides perfect naturalness will emerge in Chapter (3). That cost has to do with Humean supervenience. This is the topic I turn to next.

Chapter 3:

Humean Supervenience

3.1 Motivating Humean Supervenience

Another area where we run into trouble if we try to understand Lewis as using “qualitative” in the broad sense concerns Humean supervenience – the thesis, roughly, that “all there is to the world is a vast mosaic of local matters of particular fact, just one little thing after another.” (Lewis 1986*c*, *ix*)

Humean supervenience is often tied to the doctrine of Materialism, the endorsement of “the truth and descriptive completeness of physics more or less as we know it.” (Lewis 1986*c*, *x*) Lewis cites an adherence to Materialism as one motivation for adopting Humean supervenience: “The point of defending Humean Supervenience is not to support reactionary physics, but rather to resist philosophical arguments that there are more things in heaven and earth than physics has dreamt of.” (Lewis [1994*a*] 1999*a*, 226) Others have reiterated this motivation. Sider (2001), for instance, says that “the whole point of defending Humean Supervenience is to avoid being pushed by philosophical arguments into making posits beyond those forced by physics.” (Sider 2001, 230)

Materialism is one motivation for Humean supervenience, but it is not the only one. For if we want to avoid positing entities beyond those “forced by physics,” then we need only adopt Materialism. But Humean supervenience is not identical to Materialism. Although Lewis believes both Materialism and Humean supervenience are true, he is careful to distinguish the two.

The doctrine of Humean supervenience is both less and more demanding than the doctrine of Materialism. It is less demanding than Materialism since “it might just be that Humean supervenience is true, but our best physics is dead wrong in its inventory of the qualities.” (Lewis 1986*c*, *x*) For example, we might end up with two empirically equivalent physical theories that posit different fundamental properties, and physics might settle on the wrong one. In such a case, Humean supervenience might be true but Materialism false. So Materialism is not a necessary condition for Humean supervenience. In some ways, however, Humean supervenience is more demanding than Materialism; for although our best physics appears to posit certain sorts of entities – such as laws of nature or objective probabilities – it does not give much of an account of what these entities are like. Humean supervenience, on the other hand, places strict restrictions on what entities such as laws and chances can be like. And similar restrictions apply to accounts of features that don’t seem to appear in fundamental physics, such as causation. So Materialism is not a sufficient condition for Humean supervenience, either.

Although the desire for Materialism is one motivation for the Humean view, Materialism is neither necessary nor sufficient for Humean supervenience. But if Humean supervenience isn’t Materialism, then what is it? Lewis describes Humean supervenience in several ways. One characterization he provides is this: Humean supervenience is “the thesis that the whole truth about a world like ours supervenes on the spatiotemporal distribution of local qualities,” (Lewis [1994*a*] 1999*a*, 224) where local qualities are “perfectly natural intrinsic properties of points.” (*ibid*, 226) Call this Locality:

Locality: At worlds like ours, the perfectly natural properties are instantiated at points, and the only perfectly natural relations are spatiotemporal relations.

Within the inner sphere of possibility (i.e. among worlds where no perfectly natural properties alien to the actual world are instantiated), any two worlds that are duplicates with respect to their arrangements of perfectly natural properties and spatiotemporal relations are duplicates *simpliciter*.

However, there are several reasons to think that *Locality* does not capture the spirit of Humean supervenience. One reason is that *Locality* seems to count intuitively non-Humean worlds as Humean. If one of the motivations for adopting Humean supervenience is to resist arguments that there are more things in heaven and earth than physics has dreamt of, then worlds with things undreamt of by physics should not count as Humean. But if Humean supervenience is just *Locality*, then many such worlds are Humean worlds. Here are some examples:

1. Suppose there is a possible world physically just like ours, but where fundamental mental properties are instantiated at regions no larger than points. It seems that Humean supervenience should not obtain at such a world, although the world satisfies *Locality*.
2. Suppose there is a possible world physically just like ours, but where angels dance on the heads of point-sized pins. Again, it seems that Humean supervenience should not obtain at such a world, although the world satisfies *Locality*.
3. Suppose God is fundamental and is located at every point. God is everywhere; he is pervasive. If so, a world where such a being exists does not seem to be Humean, although it may satisfy *Locality*.

Locality counts worlds with irreducible mental properties instantiated at points, angels-on-pins, or pervasive deities as Humean. But if part of our motivation to adopt Humeanism is to resist positing entities beyond those demanded by physics, then such worlds should not be Humean worlds.

Additional reasons to think that Locality does not capture the spirit of Humean supervenience come from Lewis himself. For instance, at one point Lewis characterizes Humean supervenience as the view that everything supervenes on the “spatiotemporal arrangement of local qualities throughout all of history.” (*ibid*, 226) But it seems that not just *any* fundamental property instantiated at a point counts as a “local quality” of the sort the Humean would share a beer with. In his discussion of chances, for instance, Lewis says that primitive propensities instantiated at points, if there were any, would only be “technically Humean.” (*ibid*, 240) If Humean supervenience were just Locality, it would sound odd to call local propensities only “technically Humean,” just as it sounds odd to call an armchair only “technically furniture.” Either the armchair is furniture or it isn’t; and if it isn’t, we’d better get straight on what we mean by “furniture.” The fact that Lewis takes these propensities to be only “technically” Humean suggests that Locality doesn’t really capture the spirit of Humean supervenience.

Likewise, in his (1986*c*) introduction of Humean supervenience, Lewis describes Humean supervenience as the “doctrine that all there is to the world is a vast mosaic of local matters of particular fact,” but then makes the following remark: “But it is no part of the thesis that these local matters are mental.” (1986*c*, *ix*) Why the additional remark? If all it takes to be a local quality is instantiation at a point, then a world where all properties are instantiated at points is Humean, whether or not these properties are

mental. The remark seems oddly out of place if we take Humean supervenience to just be Locality. But it makes more sense if there is something in addition to Locality that makes a world Humean.

Yet further reasons to think that Locality doesn't capture the spirit of Humean supervenience arise when we consider how various physical theories are taken to bear on Humean supervenience. Consider quantum mechanics. Quantum mechanics strongly suggests Locality doesn't obtain at our world. (Arntzenius 2003) What does this suggest about the tenability of Humean supervenience? Although Lewis states that Humean supervenience was not designed with quantum mechanics in mind ("I am not ready to take lessons in ontology from quantum physics as it now is" (1986*c*, *xi*)), he does say the following: "If I defend the philosophical tenability of Humean Supervenience, that defence can doubtless be adapted to whatever better supervenience thesis may emerge from better physics." (Lewis [1994*a*] 1999*a*, 226)

But if we think that Humean supervenience is nothing more than Locality, then this statement is bizarre. Suppose some better physics proposes a non-local theory. To adapt Humean supervenience to this theory, we will need to reject Locality. If Humean supervenience is just Locality, then there is nothing left to adapt! Again, this suggests that Locality does not capture all there is to Humean supervenience.

One might concede that Locality does not capture the spirit of Humean supervenience, but argue that there's a quick fix: all we need is to modify Locality in the appropriate way. How might we modify Locality? One suggestion comes from Barry Loewer (1996, 104), who suggests that the fundamental properties are instantiated not at points of spacetime, but at points of configuration space. Another modification is to

extend Locality to include other fundamental relations in addition to the spatiotemporal relations, such as quantum entanglement relations. So we can amend Locality to accommodate physics in various ways.

But these modifications are problematic. First, if we don't place any principled restrictions on how we can amend Locality, "Locality" runs the danger of becoming vacuous. Consider Loewer's suggestion that Locality be applied not to spacetime, but to configuration space. If it's legitimate to move to a higher-dimensional space like configuration space to preserve the letter of Locality, then it's unclear why we cannot *always* move to a higher-dimensional space whenever we want to posit relations that appear problematic for Humeanism. Any relation whatsoever can be represented in a higher-dimensional space as a "local" property. For instance, we could employ a space isomorphic to the space of possible worlds, and take the fundamental properties to be purely local properties of points in this space. On this picture, even fundamental laws instantiated by entire worlds would count as local! Surely this notion of "locality" is not the one that Humeans are interested in.

Second, these modifications make Humean supervenience non-empirical. Consider the suggestion that we extend Locality to include whatever fundamental relations are posited by ideal physics. If this move is generally legitimate, then anything physics reveals will be compatible with Locality. If nothing physics may uncover can be incompatible with Locality, then Locality – and thus Humean supervenience – is not an empirical issue. But Lewis explicitly states otherwise: "Humean supervenience is a contingent, and therefore empirical, issue." (Lewis 1986*c*, *xi*)

Finally, even if we amend Locality in either of these ways, Locality still counts worlds with fundamental mental properties instantiated at points, angels-on-pins, and pervasive deities as Humean. So it seems that Humean supervenience is more than just Locality, or even modified versions of Locality.

More problems with understanding Humean supervenience as Locality arise when we consider classical mechanics. Lewis writes that the doctrine of Humean supervenience is “inspired by classical physics,” and so, presumably, is compatible with classical physics. ([1994a] 1999a, 226) But Locality doesn’t seem to be compatible with classical physics. Classical mechanics posits fundamental vector properties such as electric field values that don’t seem to be intrinsic properties of points. Arntzenius (2003) writes that the best we can hope for is that such vector properties are “neighborhood-local,” where a neighborhood-local property of a point p is an intrinsic property of an arbitrarily small region around p .⁴⁰

What should we conclude? Perhaps we should modify Locality to include neighborhood-local properties. But note that if we make this move, it’s not because it’s *obvious* that vector properties should count as local. Rather, it’s because it seems that however we construe Humean supervenience, classical mechanics should be compatible with that construal, and classical mechanics involves vector properties. This is the line of reasoning we see in Lewis’s (1999b) discussion of vectors: “Let’s grant that a vector quality associated with a spacetime point (or a point-sized bit of matter) shall count as local. Otherwise classical electromagnetism would be a problematic case for Humean supervenience, and we wouldn’t want that.” (Lewis 1999b, 209)

⁴⁰ More precisely: a neighborhood-local property is an intrinsic property of any open neighborhood of the point at which it is instantiated. (Arntzenius 2003, 9)

If we think that Humean supervenience is nothing more than Locality, then the dialectic here is exceedingly peculiar. Of course Lewis wouldn't *want* electromagnetism to be a problematic case for Humean supervenience – but so what? When considering whether vectors are local, the only relevant consideration should be whether vectors really are intrinsic properties of points. And any implications this has regarding the compatibility of Humean supervenience and classical mechanics is beside the point.

But if we think that Humean supervenience is more than just Locality, then the dialectic may not be so strange. Suppose we take Humean supervenience to be whatever thesis ends up best satisfying various desiderata – which may include Locality as well as compatibility with classical mechanics. Then we might interpret Lewis as follows: Lewis thinks there is clearly some way to spell out Humean supervenience that is compatible with electromagnetism and largely satisfies these other desiderata.⁴¹ And it's *this* thesis that we're interested in. If something along these lines is how we think of Humean supervenience, then we won't find Lewis's remarks puzzling. But this line of thought seems reasonable only if we don't assume that Humean supervenience is just Locality.

Perhaps some will conclude that classical mechanics and quantum mechanics provide ample evidence that there is no hope of defending Humean supervenience or related supervenience claims. But it would be better to take them as evidence that Humean supervenience is more than Locality. Many, if not most, elements of the Humean project can be endorsed whether or not our world is Local: a general optimism that our best physics will provide a complete inventory of the fundamental properties, a commitment to a regularity analysis of laws or a frequency analysis of chance, and so on. And since neither classical nor quantum mechanics seem to seriously threaten these

⁴¹ Perhaps it will satisfy something very close to Locality, such as Neighborhood-Locality.

features of Humeanism, it would be altogether too hasty to conclude that Humean supervenience – or whatever related supervenience claim that emerges from better physics – is untenable.

So we've seen that Humean supervenience is not Locality, and we've seen that it's not Materialism. But then how should we characterize Humean supervenience?

3.2 Formulating Humean Supervenience

In Chapter (2), I introduced the notion of qualitative_O. The qualitative_O properties exclude things like primitive propensities, fundamental laws of nature, brute modal properties, and so on. Such properties are nowhere instantiated because, according to Lewis, all reduce to the fundamental qualitative_O. Can we use qualitative_O to help characterize Humean supervenience? Yes, but only if we're willing to give up on some of Lewis's claims about Humean supervenience.

Let's see how qualitative_O might be used to characterize Humean supervenience. Here's an attempt:

(HS₁) Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all qualitative_O. Within the inner sphere of possibility, any two worlds that are duplicates with respect to their fundamental qualitative_O properties are duplicates *simpliciter*.

But that doesn't seem right. Given Lewis's metaphysical framework, the laws, chances, causal relations, etc. all reduce to the qualitative_O. So it's true that if we're Lewisian, we believe that the perfectly natural properties and relations at our world are all qualitative_O. But if we're Lewisian, we believe that the perfectly natural properties and relations at

every world – not just Humean worlds – are qualitative_O. So (HS₁) doesn't distinguish Humean worlds from non-Humean worlds. For although a world whose perfectly natural properties are qualitative_O is free of things like primitive propensities, fundamental laws of nature, brute modal properties, etc., it is not necessarily free of other “unHumean” stuff. Lewis writes:

Two worlds might indeed differ only in unHumean ways, if one or both of them is a world where Humean supervenience fails. Perhaps there might be extra, irreducible external relations, besides the spatiotemporal ones; there might be emergent natural properties of more-than-point-sized things; there might be things that endure identically through time or space and trace out loci that cut across all lines of qualitative continuity. It is not, alas, unintelligible that there might be suchlike rubbish. Some worlds have it. And when they do, it can make differences between worlds even if they match perfectly in their arrangements of qualities. (Lewis 1986c, x)

If there are worlds with non-spatiotemporal external relations or emergent natural properties of more-than-point-sized things, these worlds would not be Humean. But such unHumean properties are nonetheless qualitative_O, since they are part of the reductive base for the laws, chances, causal relations, etc. at these strange worlds. So qualitative_O is not enough to characterize Humean supervenience; we need to restrict the properties instantiated at Humean worlds even further.

Perhaps we could restrict the properties at Humean worlds to those that are *local* and qualitative_O:

(HS₂) Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all local and qualitative_O. Within the inner sphere of possibility, any worlds that are local and qualitative_O duplicates are duplicates *simpliciter*.

But that doesn't seem quite right either. One of the motivations for adopting Humean supervenience is Materialism, and worlds with irreducible mental properties instantiated at points, or angels dancing on the heads of point-sized pins, or pervasive deities

instantiated at every point are not compatible with Materialism.⁴² So it seems they should not be compatible with Humean supervenience either. But such properties are local and qualitative_O; given (HS₂) worlds that have these properties are Humean worlds. So locality and qualitative_O are not enough to characterize Humean supervenience; we need to restrict the properties instantiated at Humean worlds even further.

Perhaps we could restrict the properties at Humean worlds to those that are *physical*, local, and qualitative_O. Let's call a property *physical* if it appears in our best physics, or supervenes on the properties that appear in our best physics. Then we could attempt to characterize Humean supervenience as:

(HS₃) Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all physical, local, and qualitative_O. Within the inner sphere of possibility, any worlds that are physical, local, qualitative_O duplicates are duplicates *simpliciter*.

(HS₃) rules out worlds with irreducible mental properties instantiated at points, angels-on-pins, or pervasive deities as Humean. In fact, I think it comes tantalizingly close to the characterization of Humean supervenience that many philosophers – perhaps even Lewis – have in mind.⁴³ But (HS₃) has some drawbacks.

⁴² I.e. these properties are not compatible with the claim that physics, more or less as we know it, provides a complete and comprehensive description of the world.

⁴³ Those who think that it's a datum that Humean supervenience is compatible with classical mechanics will not be satisfied with (HS₃). They may want to modify (HS₃) in the following way:

(HS₃*) Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all physical, neighborhood-local, and qualitative_O. Within the inner sphere of possibility, any worlds that are physical, neighborhood-local, qualitative_O duplicates are duplicates *simpliciter*.

First, it's incompatible with Lewis's (1986c) description of Humean supervenience, where he writes:

Is [Humean supervenience] materialism? – no and yes. I take it that materialism is metaphysics built to endorse the truth and descriptive completeness of physics more or less as we know it; and it just might be that Humean supervenience is true, but our best physics is dead wrong in its inventory of the qualities. (Lewis 1986c, x)

Recall the example mentioned in the previous section: suppose we end up with two empirically equivalent theories that posit different fundamental properties, and physics settles on the wrong one. In this example, there are fundamental properties that are *not* physical – i.e. there are fundamental properties that do not appear in our best physics. In such a case, Lewis would want to say that Humean supervenience could be true even though Materialism is false. But if we adopt (HS₃), then Humean supervenience entails Materialism, and we must accept that both Materialism and Humean supervenience are false in this case.⁴⁴

Second, the notion of “physical” is difficult to cash out precisely. Are the physical properties those that appear in our present-day physics? That seems to make “physical” unacceptably dependent on contingent, historical factors. Are the physical properties those that appear in the *correct* physics? That seems question-begging: the correct physics is the theory that correctly identifies the physical properties, and the physical properties are just those identified by the correct physics. Lewis says that the physical properties are “those that are mentioned in the language of physics,” where the language of physics is taken to be something “not too different from present-day physics

⁴⁴ One might think the following: when Lewis says that Humean supervenience may be true but our physics “dead wrong” about its inventory of qualities, he meant to include cases where there are perfectly natural properties like point-sized spirits, mental properties, etc. So Lewis thinks that Humean supervenience *is* compatible with point-sized spirits, mental properties, etc. – and we should never have been tempted to include “physical” in our characterization of Humean supervenience in the first place. But this thought leads back to the claim that Humean supervenience is simply Locality. And as I’ve shown in the previous section, there are a number of reasons to think this is false.

– but presumably somewhat improved.” ([1983*b*] 1999*a*, 33-34) But this is more of a gesture towards what it takes to be “physical” than an account. Those who want to adopt (HS₃) need to give some sort of account of what they mean by “physical.”

These are the costs of using qualitative_O to characterize Humean supervenience. If one is unwilling to accept these costs, however, then (HS₃) is not a satisfactory characterization of Humean supervenience. Indeed, it’s hard to see how *any* characterization of Humean supervenience in terms of qualitative_O can capture the spirit of Humeanism while keeping Materialism and Humean supervenience independent, and avoiding all mention of “physical.” So if we want to capture the spirit of Humeanism without using qualitative_O, it seems we’ll need another primitive notion: “Humean”-qualitative, or qualitative_H. This notion was introduced in section (2.1) to refer to those properties found at Humean worlds. Now we can formulate Humean supervenience using qualitative_H and locality:

(HS₄) Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all local and qualitative_H. Within the inner sphere of possibility, any worlds that are local and qualitative_H duplicates are duplicates *simpliciter*.⁴⁵

The absence of “physical” from (HS₄) gives this formulation two advantages over (HS₃). First, (HS₄) is compatible with Lewis’s claim that Humean supervenience does not entail Materialism. Again, suppose we have two empirically equivalent theories, and physics

⁴⁵ As before, those who think that it’s a datum that Humean supervenience is compatible with classical mechanics will not be satisfied with (HS₄), and may want to modify it using neighborhood-locality instead:

(HS₄*) Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all neighborhood-local and qualitative_H. Within the inner sphere of possibility, any worlds that are neighborhood-local and qualitative_H duplicates are duplicates *simpliciter*.

settles on the wrong one. Then Materialism will be false. If we adopt (HS₄), and the correct theory only invokes local qualitative_H properties, then Humean supervenience will be true. So (HS₄) allows for the possibility that Materialism and Humean supervenience may come apart.

How much of an advantage is this? I do not think it is overwhelming. Although Lewis believes Humean supervenience is distinct from Materialism, he also believes it's unlikely that the two will come apart: "Most likely, if Humean supervenience is true at all, it is true more or less the way that present physics would suggest." (Lewis 1986c, x-xi) So I do not think that (HS₄) is a dramatic improvement over (HS₃). True, (HS₄) keeps Materialism and Humean supervenience separate, and (HS₃) does not; but since it seems likely that either both are true or neither is, (HS₄) does not have an overwhelming advantage over (HS₃) on this point.

The second advantage of adopting (HS₄) over (HS₃) is that we don't need to resort to physical properties to rule out mental properties instantiated at points, angels dancing on the heads of point-sized pins, pervasive deities, or other unHumean rubbish. We can employ qualitative_H properties instead, and thus avoid worries about cashing out "physical."

One might complain that replacing "physical" – a vague notion in need of an account – with "qualitative_H" – a somewhat obscure primitive – is not much of an improvement. I am sympathetic to this complaint. Still, it is *some* improvement. At least (HS₄) is upfront about the primitives it employs.

Note that both (HS₃) and (HS₄) are easily adaptable to accommodate physical theories that posit fundamental relations besides the spatiotemporal ones, such as

quantum entanglement relations. All we need to do is remove the stipulation that the perfectly natural properties at Humean worlds are local. Let's call the "better supervenience claim that emerges from better physics" *Awesome Humean supervenience*. If one prefers formulating Humean supervenience in terms of qualitative_O , then one can give the following characterization of Awesome Humean supervenience:

(AHS₁) Awesome Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all physical and qualitative_O . Within the inner sphere of possibility, any worlds that are physical and qualitative_O duplicates are duplicates *simpliciter*.

If one prefers formulating Humean supervenience in terms of qualitative_H , then one can give the following characterization of Awesome Humean supervenience:

(AHS₂) Awesome Humean supervenience is the thesis that at worlds like ours, the perfectly natural properties and relations are all qualitative_H . Within the inner sphere of possibility, any worlds that are qualitative_H duplicates are duplicates *simpliciter*.

3.3 Rejecting Either Qualitative_H or Qualitative_O

What do we lose if we reject qualitative_H ? It depends on whether we posit qualitative_O . In this section, I explore the consequences of rejecting one of qualitative_H or qualitative_O , but not both. In the following section, I explore the consequences of rejecting both qualitative_H and qualitative_O .

3.3.1 Rejecting Qualitative_H, Accepting Qualitative_O

If we accept just qualitative_O – and reject qualitative_H – then we cannot characterize Humean supervenience as (HS₄), and must instead choose (HS₃). The costs we incur are those mentioned in the previous section. First, it seems that any plausible formulation of Humean supervenience that does not involve qualitative_H will entail the truth of Materialism, an entailment that Lewis explicitly denies. Second, advocates (HS₃) need to provide some account of “physical.” These costs must be weighed against the cost of positing qualitative_H as primitive.

3.3.2 Accepting Qualitative_H, Rejecting Qualitative_O

If we accept just qualitative_H – and reject qualitative_O – then we can characterize Humean supervenience as (HS₄). That way we can avoid formulating Humean supervenience in a way that entails Materialism, and thus avoid conflict with Lewis on this point.

But this option incurs several costs. Consider Lewis’s “interesting supervenience claims” – claims that the laws, chances, etc. supervene on the qualitative properties – discussed in Section (2.2). We saw that these supervenience claims are trivially true if we understand Lewis’s use of “qualitative” as qualitative_B. We avoid these problems if we understand “qualitative” as qualitative_O, instead. Using qualitative_O, we can formulate supervenience claims such as the following:

1. Worlds that are duplicates with respect to their qualitative_O properties are duplicates with respect to the laws of nature.

2. Worlds that are duplicates with respect to their qualitative_O properties are duplicates with respect to the distributions of objective probability.
3. Worlds that are duplicates with respect to their qualitative_O properties are duplicates *simpliciter*.

Philosophers working within a broadly Lewisian framework tend to endorse each of these claims, and philosophers rejecting that framework tend to deny them (c.f. Maudlin 2007).

Likewise, we saw that if we posit qualitative_O, we can make sense of the reductive projects corresponding to these supervenience claims – projects motivated by a general commitment to qualitative_O reduction. Again, those working within a broadly Lewisian framework tend to endorse these reductive claims, and philosophers rejecting that framework tend to reject them.

But if we reject qualitative_O, then how should we understand these supervenience claims and reductive projects? We might try to replace “qualitative_O” with “qualitative_H.” For the supervenience claims just described, this gives us:

- 1*. Worlds that are duplicates with respect to their qualitative_H properties are duplicates with respect to the laws of nature.
- 2*. Worlds that are duplicates with respect to their qualitative_H properties are duplicates with respect to the distributions of objective probability.
- 3*. Worlds that are duplicates with respect to their qualitative_H properties are duplicates *simpliciter*.

Although these supervenience claims are not trivial, they are poor replacements for (1)-(3). First, Lewis would not want to endorse these supervenience claims. Suppose two worlds are duplicates with respect to their qualitative_H properties. If they are non-

Humean, they might still differ with respect to their laws, chances, or other non-qualitative_H properties. Since Lewis takes Humean supervenience to be contingent, he will deny that these supervenience claims obtain.⁴⁶

Second, Lewis would not want to endorse any of the reductive projects corresponding to these supervenience claims. Let's first suppose that everything – everything non-haecceitistic, that is – reduces to the qualitative_H. Then certain *prima facie* possible worlds turn out to be impossible. It would be metaphysically impossible for there to exist an irreducible God, and it would be metaphysically impossible for there to be irreducible emergent properties of macroscopic objects, and so on. But Lewis believes there could be such things. He just believes there aren't such things at the actual world.

Third, even if we abandon wholesale reduction to qualitative_H, we still may want to pursue particular reductions – reduction of the laws, chances, etc. As we saw in section (2.3), Earman-style analyses of laws, chances, and the like require the notion of qualitative_O. And if we substitute qualitative_H in place of qualitative_O, we won't get the right results. For instance, suppose there is a possible world with nothing but God and the angels of the Seven Heavens. They all live in a perfect, well-ordered harmony, and their activities can be described in a superbly simple and informative way. But this world instantiates no qualitative_H properties. So if we analyze laws in terms of the qualitative_H properties, then we must say that this world has no laws. If one were to look only at the

⁴⁶ True, among *Humean* worlds, supervenience does not fail. Within the inner sphere of possibility, worlds that are duplicates with respect to their qualitative_H properties are duplicates *simpliciter*. And so we can try to pursue the corresponding reductive projects as long as we restrict their scope to Humean worlds: we can ask whether, among Humean worlds, the laws reduce to regularities in the pattern of qualitative_H properties, or the chances reduce to frequencies, or the causal relations reduce to similarities between Humean worlds, and so on. But even if they do, we cannot give an *analysis* of laws or chances or causation. Analyses are necessary, and Humean supervenience is contingent.

laws of nature of every possible world, one could not distinguish this well-ordered paradise from a hellish world where chaos reigns.

Furthermore, if we substitute qualitative_H for qualitative_O in Earman-style analyses of laws, then we must rule out *actual* physical theories in which non-qualitative_H properties play a role in the laws. For example, we must say that certain quantum mechanical theories, such as Wigner’s (1961) “consciousness” collapse theory and Albert and Loewer’s (1988) “many minds” theory, in which irreducible consciousness facts play a role in the laws of nature are impossible – not nomologically impossible, but *metaphysically* impossible. Such theories are admittedly bizarre, but given the pedigree of those who have proposed them we should be hesitant to rule them out by fiat.⁴⁷

3.4 Rejecting Qualitative_H and Qualitative_O

In sections 2.6 and 3.3.2, we saw that those who reject qualitative_O face difficulties when it comes to stating certain claims about supervenience and reduction. Those who reject both qualitative_O and qualitative_H face an additional difficulty: they cannot provide a satisfactory characterization of Humean supervenience.

I’ve argued that a satisfactory formulation of Humean supervenience requires an additional primitive notion: either qualitative_O or qualitative_H. If we reject both qualitative_O and qualitative_H, what should we do? Let’s consider two avenues.

First, one might try to characterize Humean supervenience in some way that avoids appeal to any other primitive notion. The most promising characterization along

⁴⁷ What if we replace qualitative_O with qualitative_B instead of qualitative_H in this Lewisian sort of analysis of lawhood? The result will be compatible with accounts of laws that Earman and many Lewisians would want to reject. For instance, if a world has fundamental modal properties, then these properties may appear in the theorems of the best system of the world. But proponents of this kind of analysis do not generally believe that fundamental modal properties play any role in the laws.

these lines seems to be Locality. But we've seen in section (3.1) that characterizing Humean supervenience as Locality has some highly undesirable consequences. For instance, it seems that quantum mechanics demonstrates that Locality, and thus Humean supervenience, is false. If Humean supervenience is just Locality, then Humean supervenience is *completely* falsified. And it's hard to see why we should care about a completely falsified theory. This makes the project of constructing theories compatible with Humean supervenience about as interesting and worthwhile as the project of constructing theories that are compatible with phlogiston.

Second, one might give up on characterizing Humean supervenience at all. This is a serious cost for those with a stake in these issues. Jonathan Schaffer (2003), for example, argues that we should favor accounts of chance that are compatible with Humean supervenience. Since propensity theories of chance are not compatible with Humean supervenience, this is a mark against such accounts. But if we don't have any way of assessing whether an account is Humean, then we cannot say that propensity theories are *not* Humean, and use that as a reason to reject propensity accounts.⁴⁸ More generally, one often finds philosophers who take compatibility with Humean supervenience to be a reason in favor of a metaphysical position.⁴⁹ And if we cannot formulate any of the Humean's central tenets, then this move is illegitimate.⁵⁰

⁴⁸ Schaffer says that on a Humean account of chance, the chances supervene on the "arrangement of occurrent facts," where an occurrent fact is "a categorical, intrinsic quality of a region (or a field, or a portion of matter, or whatnot)." (Schaffer 2003, 29) But I do not know what it means to say that a property is "occurrent" or "categorical," unless we understand this in terms of qualitative_O or qualitative_H.

⁴⁹ Example: Sider (2001) writes that any metaphysical picture that "requires giving up Humean supervenience... should be avoided if possible." (Sider 2001, 230)

⁵⁰ We can put the point another way. As we saw earlier (section (2.6.1)), if one wants to use a notion in any argument or claim, one must either provide an account of that notion or accept it as primitive. By assumption, we've given up on trying to provide an account. And presumably the proponent of this response will not want Humean supervenience as a primitive – if we're going to resort to adopting a primitive to characterize Humean supervenience, we might as well adopt qualitative_O or qualitative_H. So if

One might object that even though we cannot formulate Humean supervenience, this does not mean we have no grip on the doctrine, and thus does not mean we cannot cite an adherence to Humean supervenience as a reason for rejecting or accepting various metaphysical claims. After all, even if we cannot state what Materialism is, it seems we can nonetheless cite an adherence to Materialism as a reason to reject, say, fundamental properties of consciousness. But note that if one believes that Materialism rules out fundamental consciousness properties, then one needs to say *why* such properties are incompatible with Materialism. And that requires giving an account of what Materialism is.⁵¹ And if it turns out that no characterization can be given, then we face a choice: either the distinction between fundamental physical and non-physical properties is primitive, or Materialism is not a coherent doctrine.

Similarly, if we believe that Humean supervenience entails that there are no primitive propensities, then we must say something about what Humean supervenience is. And if it turns out that no characterization can be given, then we face a choice: either the distinction between Humean and non-Humean properties is primitive, or Humean supervenience is not a coherent doctrine.

we adopt this response, then it's illegitimate to use Humean supervenience in arguments in favor of or against metaphysical claims.

⁵¹After all, we can imagine formulations of Materialism that will not rule out such properties. Suppose we formulate Materialism as a supervenience claim along the following lines:

(M) Any two worlds alike in their physical properties are alike in their mental properties.

If mental properties supervene on physical ones, then (M) is true. But (M) does not rule out the possibility that mental properties are fundamental. So if we want to maintain that Materialism entails that there are no fundamental mental properties, we must say something about what Materialism is.

3.4.1 An Escape Route?

Perhaps things are not so dire. One might think that even though we cannot provide a satisfactory characterization of Humean supervenience without some primitive notion of qualitative_O or qualitative_H, we can nonetheless provide a satisfactory characterization of some of the Humean's individual theses. And if we can characterize some individual Humean theses, then we can still say that this or that view is incompatible with Humean supervenience because it violates one of these claims.

Let's use the laws of nature as a case study, and try to characterize a Humean account of laws. Here's a first pass attempt:

(L₁) The laws of nature are Humean *iff* they are not perfectly natural.

This won't work, for (L₁) is compatible with perfectly *non*-natural laws – laws which, like haecceitistic properties, fail to supervene on the perfectly natural ones. But perfectly non-natural laws are not Humean! If we want to rule out this possibility, we need modify

(L₁):

(L₂) The laws of nature are Humean *iff* they are not perfectly natural, and they supervene on the perfectly natural.

That's better, but not by much. Suppose I believe that the laws of nature reduce to fundamental causal relations. Such an account is compatible with (L₂), but would not be Humean.

Perhaps we could modify (L₂) in the following way:

(L₃) The laws of nature are Humean *iff* they are not perfectly natural, and they supervene on the perfectly natural, and there are no perfectly natural causal relations.

This rules out accounts of laws that involve fundamental causal relations, thus getting us one step closer to a Humean account. But (L₃) is also too weak. For I might believe that the laws of nature reduce to fundamental modal properties. Such an account is compatible with (L₃), but would not be Humean.

Of course, we could modify (L₃) to exclude fundamental modal properties as well. But then we could raise similar problems for primitive propensities, dispositional powers, motivational essences, and so on.⁵² And we could go on in this way indefinitely. To get around this problem, we need a way of specifying what the undesirable properties are. But this is exactly what we give up if we reject qualitative_O and qualitative_H. And we cannot get around this problem by trying to state the Humean's tenets as individual theses.⁵³

⁵² And for some worlds – like ours – we may want to add mental properties, emergent properties, ghosts, deities, genidentity relations, and so on to the list of properties that are not perfectly natural.

⁵³ Suppose I believe that among the perfectly natural properties are “motivational essences,” upon which the laws, chances, causal relations, etc. supervene. Is this a Humean view? If it depends on whether motivational essences appear on the list of Humean-compatible properties, then Humean supervenience is in trouble. For whether or not a property appears on the list of Humean-compatible properties ought to depend on whether the property *really is* compatible with Humeanism, not the other way around. If there is nothing in virtue of which motivational essences are or are not Humean, then the decision either way is entirely arbitrary. We have no grounds for declaring such a view unHumean, and no grounds for declaring it Humean, either. So this move doesn't help us formulate individual Humean supervenience claims, and it doesn't help us say what makes an analysis compatible with Humean supervenience.

3.4.2 Assessing the Costs

We've seen that rejecting both qualitative_O and qualitative_H incurs a number of costs. We cannot state any of the Lewisian's claims about supervenience on the qualitative_O. We cannot state any of the Lewisian's claims about reduction to the qualitative_O. We cannot provide analyses of laws, chances, etc. which employ qualitative_O. And we cannot provide a satisfactory characterization of Humean supervenience. In the end, the Lewisian cannot state his own view in a way that distinguishes it from rival views – views that posit, say, fundamental laws or primitive propensities or any other number of things that the Lewisian finds offensive.

These costs are not insignificant. But neither are they decisive. If the Lewisian wants to avoid these consequences, then he needs to add some additional primitives to his ideology: qualitative_O, qualitative_H, or both. And that addition incurs a cost as well: the more primitives a theory has the less appealing that theory becomes. One might think that, all things considered, the cost of positing additional primitives outweighs the other costs. But it's a matter on which reasonable people may diverge.

Chapter 4: Intrinsicity and Hyperintensionality

4.1 Introduction

Lewis argues that we should accept natural properties because of their great many philosophical uses. In the previous two chapters, we've seen that natural properties are not well-suited to play some of the roles Lewis desired. Given only natural properties, we are unable to formulate certain sorts of supervenience and reductive claims; certain accounts of laws, chances, and the like; and we are unable to provide a satisfactory characterization of the doctrine of Humean supervenience.

Another area where Lewis puts natural properties to work is in analyzing intrinsicity. In the present chapter, I argue that here, too, natural properties are ill-suited to do the work required. The reason is this: the intrinsic/extrinsic distinction discriminates between cointensive properties, but Lewis's account in terms of natural properties cannot.

The plan is as follows: I begin by briefly outlining some ways that advocates of a given account may try to deflect apparent counterexamples. I then introduce two sorts of properties that appear to cause trouble for Lewis's analysis of intrinsicity: identity properties and necessary properties. I examine three different ways defenders of this account have tried to block counterexamples that stem from identity properties and necessary properties, and argue that none succeed. I conclude by diagnosing the source of the problem: *intrinsic* is a hyperintensional notion.

Throughout most of the discussion, I assume the space of possibility we are

working with is metaphysical possibility. In section (4.8), I explore the (lack of) consequences of moving to the space of conceptual possibility instead. Finally, I discuss how the conclusions I reach are tied to the conclusions reached in Chapters (2) and (3).

4.2 Counterexamples and Methodology

There are several options available when a philosophical account appears susceptible to counterexamples. One option is to reject the account. Another option is to bite the bullet, and grant that the account yields counterintuitive results. A third option is to argue that any apparent counterintuitive consequences are merely apparent; understood appropriately, the alleged “counterexamples” are not counterexamples at all.

The third option can be pursued in various ways. One move is to restrict the scope of the account in order to exclude counterintuitive cases.⁵⁴ Another move is to claim that there are several different concepts in play, and the account in question successfully tracks one of these concepts.⁵⁵ Finally, one may argue that our intuitions about alleged counterexamples are inconsistent or untrustworthy, and therefore should not carry any weight.⁵⁶

Each of these moves has been deployed in defense of Lewis’s (1983*b*) account of intrinsicity. I will argue that, in this context, none of these moves are legitimate. While I focus on Lewis’s original proposal, my criticisms apply *mutatis mutandis* to any account that is not sensitive to fine-grained or hyperintensional distinctions among

⁵⁴ An example: Kim (2005) argues that a functional characterization of mental states need not extend to qualia.

⁵⁵ An example: Hall (2004) argues that we have two distinct concepts of causation, and that counterfactual accounts of causation track only one of them.

⁵⁶ An example: Lewis (1973) argues that counterintuitive results of his counterfactual analysis of causation are acceptable because they depend on unreliable intuitions.

properties.⁵⁷ Only a hyperintensional framework has the structure required for an adequate account of intrinsicality.

4.3 Two Worries

According to Lewis's (1983*b*) account of intrinsicality, a property *P* is intrinsic *iff* for any two possible duplicates, either both have *P* or both lack it. Lewis then analyzes duplication in terms of natural properties: two objects are duplicates *iff* they share all their perfectly natural properties, and their parts can be put into correspondence in such a way that corresponding parts have the same perfectly natural properties, and stand in the same perfectly natural relations. (Lewis 1983*b* and 1986, 61) Let us call this account of intrinsicality the Duplication Account.

Two standard worries arise with respect to the Duplication Account. The first concerns identity properties.⁵⁸ An identity property is the property of being a particular individual. *Being David Lewis* is an identity property, as are the properties *being me* and *being you*.⁵⁹ There are two kinds of identity properties. One kind corresponds to the singleton set of a world-bound individual; the other kind corresponds to the set containing a world-bound individual *p* and all of *p*'s counterparts. Both cause trouble for the Duplication Account. For simplicity, I will focus on the former. Consider the identity property *being me*. Intuitively, *being me* is intrinsic; whether or not I instantiate it has nothing to do with what other things there are or how they are. Likewise, *being David Lewis* seems intrinsic, and so does *being you*. But since identity properties divide

⁵⁷ This includes the account suggested by Langton and Lewis (1998).

⁵⁸ The term comes from Sider (1996).

⁵⁹ Identity properties are a species of *haecceitistic* properties, but not all haecceitistic properties are identity properties. For example, *being five feet from Abraham Lincoln* is a haecceitistic property, but it is not an identity property. *Being Abraham Lincoln* is both a haecceitistic property and an identity property.

duplicates, the Duplication Account classifies them as extrinsic. That seems mistaken. (I am not yet claiming that it *is* mistaken, just that it seems that way.)⁶⁰

The second worry concerns necessary properties. Necessary properties are properties shared by all possible individuals; *a fortiori* they are shared by all duplicates. According to the Duplication Account, all necessary properties are intrinsic. Again, that seems mistaken. Suppose, for example, that for every possible amount of mass, there is some possible object that instantiates that amount of mass. There is no upper bound on mass, so every possible individual instantiates the property *being such that there possibly exists something greater in mass*. Intuitively, this property is extrinsic, but not according to the Duplication Account.

These worries have a common source: our intuitions about intrinsicity are fine-grained, but the Duplication Account is not. Consider two different ways of thinking about properties. A rough way to individuate properties is *intensionally* – by their possible instances. On an intensional conception of properties, properties that are cointensive – have the same instances across possible worlds – are identical. A more fine-grained way to individuate properties is *hyperintensionally*. On a hyperintensional conception, cointensive properties may be distinct.

The worries above arise because cointensive properties may intuitively diverge with respect to their intrinsicity. But on the Duplication Account, cointensive properties cannot diverge in this way. Here's why. The Duplication Account says that a

⁶⁰ The second kind of identity properties raises similar problems for the Duplication Account. Suppose, for example, that under some intuitive counterpart relation, anything born to counterparts of my parents around the time that I was actually born is a counterpart of me. Then, I may have a counterpart that is quite different from me: this counterpart might be handicapped, or might differ in gender, or might grow up to be a famous football player. Since it is a counterpart of me, it instantiates the identity property *being me* (using the “counterpart” notion of an identity property). But it will have duplicates that are not my counterparts – a famous football player created in a vat, say – and so do not instantiate *being me*. Therefore, the Duplication Account will classify *being me* as extrinsic.

property is intrinsic when it doesn't divide duplicates. A set of duplicates is a set of possible individuals. So if two properties correspond to the same set of possible individuals, then they are alike with respect to whether they divide duplicates. Thus, for the Duplication Account, cointensive properties are alike with respect to their intrinsicity. So even if one is open to the possibility of distinct cointensive properties – as Lewis is (see Lewis 1986*b*, 56) – the Duplication Account cannot capture differences in intrinsicity between them. How might one defend the Duplication Account from these worries? Each of the three moves discussed in section (4.2) has been offered on behalf of the Duplication Account. In what follows, I examine each response in turn, and argue that none succeed.

4.4 The Qualitative Response

Call the first response to these worries the “Qualitative Response.” According to the Qualitative Response, Lewis's account applies only to *qualitative* properties; and since the alleged counterexamples all involve non-qualitative properties, they have no force against the account.

But what do we mean by “qualitative”? In the context of the Qualitative Response, “qualitative” is used to mean something like *non-haecceitistic*. Thus, says this response, the Duplication Account applies only to non-haecceitistic properties – properties like *having a beard* – and does not apply to haecceitistic properties – properties like *being David Lewis*.

Let's begin with the worry involving identity properties. In “Defining ‘Intrinsic’,” Langton and Lewis appeal to the Qualitative Response in defending their

account against counterexamples involving identity properties (or *haecceities*):

A first qualification is that the proposed definition, and likewise all that follows, is to be understood as restricted to pure, or qualitative, properties – as opposed to impure, or haecceitistic, properties... Our proposal is offered as a way of distinguishing amongst the pure, or qualitative properties, those which are intrinsic, and those which are extrinsic. Impure properties are set aside as falling outside the scope of the present discussion. (Langton and Lewis [1999a] 1998, 335)

The same response can be made in defending the Duplication Account: since identity properties like *being me* are impure, or non-qualitative, they fall outside the scope of Lewis's analysis. Therefore, the response goes, identity properties are not counterexamples to the Duplication Account.

One worry with this response is that it seems somewhat arbitrary. Our only motivation for restricting the account to qualitative properties seems to be the desire to avoid counterexamples. Consider an analogy. Jaegwon Kim (1982) proposed an account of intrinsicity according to which a property *P* is intrinsic *iff* a lonely object – an object that is the only inhabitant of a world – may have *P*. Lewis (1983a) objects that Kim's account incorrectly classifies the extrinsic property being lonely as intrinsic, and concludes that Kim's account was untenable. Now suppose a defender of Kim's account responds thus: the account is not intended to apply to "lonely" properties, or properties an object has if and only if it is unaccompanied. Since "lonely" properties fall outside the scope of Kim's account, they cannot be raised as counterexamples. The trouble is that this response seems unmotivated: the only reason to exclude lonely properties is that the success of Kim's account requires it.

A deeper worry with the Qualitative Response is the difficulty of spelling out the distinction between qualitative (or non-haecceitistic) and non-qualitative (or haecceitistic) properties. The standard strategy is to characterize the qualitative properties as those that

globally supervene on the perfectly natural properties and relations.⁶¹ (See Bricker (2007, fn 24), Langton and Lewis (1998, 344), Lewis (1986*b*, 62-63), McDaniel (2007, 250), Teller ([1985] 1999, 18).) Thus, the standard strategy is to take “qualitative” to mean what I earlier called *broad-qualitative*, or qualitative_B. But if we take qualitative to mean qualitative_{E_B}, then the Qualitative Response does not work.

We can spell this out in two different ways, depending on whether we allow for duplicate worlds. (Lewis himself remains neutral on whether there are duplicate worlds (Lewis 1986*b*, 87).) Let’s first assume that there are no duplicate worlds. (Let’s also assume, as we ordinarily think, that our world isn’t a symmetrical world.⁶²) If no distinct worlds are duplicates, then no distinct worlds are alike with respect to their perfectly natural properties and relations. It follows that my identity property, *being me*, supervenes on the perfectly natural: any world with the same distribution of perfectly natural properties and relations as the actual world has the same distribution of the property *being me* (since the only world with the same distribution of perfectly natural properties and relations as the actual world *is* the actual world). Since *being me* supervenes on the perfectly natural, it is qualitative_B; since it is qualitative_B, it falls within the scope of the Duplication Account, which counterintuitively classifies it as extrinsic. But the Qualitative Response was supposed to safeguard the Duplication Account from these kinds of counterexamples by saying that identity properties are not qualitative_B, and

⁶¹ There are different kinds of global supervenience. Consider two: strong and weak. Following Sider (1999), we can characterize these notions this way: *A strongly globally supervenes on B* iff for any worlds w_1 and w_2 , every *B*-preserving isomorphism between w_1 and w_2 is an *A*-preserving isomorphism. *A weakly globally supervenes on B* iff for any worlds w_1 and w_2 , if there is a *B*-preserving isomorphism between w_1 and w_2 , then there is an *A*-preserving isomorphism. The difference between strong and weak global supervenience won’t matter for present purposes. See also footnote (10) of this chapter.

⁶² A non-symmetrical world is a world where the only one-one function that maps the domain of the world to itself is the identity map.

thus are not classified by the Duplication Account.⁶³

Next, assume that there are duplicate worlds. Suppose that the actual world has exactly one duplicate, w_1 , and that my duplicate at w_1 instantiates *being me*₁. *Being me* and *being me*₁ seem intrinsic, and so does their disjunction, *being me or being me*₁. But *being me or being me*₁ supervenes on the perfectly natural: any world with the same distribution of perfectly natural properties and relations as the actual world has the same distribution of the property *being me or being me*₁. Since *being me or being me*₁ supervenes on the perfectly natural, it is qualitative_B; since it is qualitative_B, it falls within the scope of the Duplication Account, which counterintuitively classifies it as extrinsic. Again, the Qualitative Response was supposed to safeguard the Duplication Account from these kinds of counterexamples by saying that identity properties (and disjunctions thereof) are not qualitative_B, and thus are not classified by the Duplication Account.

The Qualitative Response aims to block counterexamples to the Duplication Account that arise from haecceitistic properties by restricting the scope of the account to qualitative, or non-haecceitistic, properties. But how shall we distinguish the haecceitistic properties from the non-haecceitistic ones? As we've seen, characterizing "qualitative" as qualitative_B won't work. However, in Chapter (2) I introduced two other notions of "qualitative" in addition to qualitative_B: qualitative_O and qualitative_H. Could we use either one of these notions to capture the distinction desired by advocates of the Qualitative Response?

⁶³ Are there *any* identity properties that do not globally supervene on the perfectly natural, and so are non-qualitative on this characterization? If we understand "qualitative" in terms of strong global supervenience, then yes. The identity properties of objects at symmetrical worlds – worlds where there are multiple one-one mappings from the domain of the world onto itself that preserve the perfectly natural properties and relations – will not strongly globally supervene. But the identity properties of objects that are not located at symmetrical worlds will still strongly globally supervene on the perfectly natural properties, and so are still qualitative_B.

No, we cannot. If we say that the Duplication Account applies only to qualitative_O properties, we still include properties in the scope of the account that intuitively should not be included: haecceitistic properties like *being me*. And we exclude properties from the scope of the account that intuitively should not be excluded: properties like *being a world such-and-such fundamental laws*. Similarly for qualitative_H. If we say that the Duplication Account applies only to qualitative_H properties, we again include haecceitistic properties within the scope of the account; and we exclude even more properties from the scope of the account that intuitively should not be excluded: properties like *being an angel* or *being in so-and-so irreducible phenomenal state*.

So we cannot capture the distinction desired by advocates of Qualitative Response by appealing to qualitative_B, qualitative_O, or qualitative_H. Still, the intuitive distinction that we want is clear enough: properties such as *having a beard* should be qualitative or non-haecceitistic, and properties such as *being David Lewis* should be non-qualitative or haecceitistic. Perhaps we could simply posit the distinction as primitive. But if we do this, the Qualitative Response loses its appeal. If one must posit a primitive distinction in order to get Lewis's account to apply to the right properties, one might as well take the intrinsic/extrinsic distinction as primitive, and be done with it.

Moreover, this version of the Qualitative Response seems to require a conception of properties that allows for cointensive properties to be distinct. Assume, for simplicity, that there are no duplicate worlds. Consider the property *having such-and-such features and so-and-so relations to other things* which is cointensive with *being me*. The former property is intuitively extrinsic, the latter is intuitively intrinsic. But on the Duplication

Account, both are extrinsic. The Qualitative Response attempts to avoid this counterintuitive result by restricting the Duplication Account to qualitative properties like *having such-and-such features and so-and-so relations to other things*. But this restriction is intelligible only if one believes that cointensive properties – like *being me* and *having such-and-such features and so-and-so relations to other things* – may be distinct. Conceding that cointensive properties may be distinct undermines the Duplication Account. Once one is willing to grant that properties are individuated hyperintensionally, why would one desire an account of intrinsicity that cannot recognize hyperintensional distinctions?

Finally, let's consider the worry involving necessary properties. The Qualitative Response fails here as well. All necessary properties supervene on the perfectly natural properties, so all necessary properties are qualitative_B, and fall within the scope of the Duplication Account. And even if we were to adopt a primitive distinction between qualitative and non-qualitative that aligns with an intuitive classification, many of the necessary properties in question – such as *being such that there possibly exists something greater in mass* – are intuitively qualitative. So the Duplication Account would still classify them as intrinsic, even though they are intuitively extrinsic. No matter how we get the distinction between the qualitative and non-qualitative properties, the Qualitative Response is not a plausible defense of the Duplication Account.

4.5 The Many Notions Response

Like the Qualitative Response, the Many Notions Response restricts the scope of the analysis. But it does so in a different way. Advocates of this response claim that there

are multiple notions of intrinsicity, all of which are interesting and important, and Lewis's account successfully tracks one of these notions.

The Many Notions Response may seem compelling when we reflect on how difficult it is to pinpoint exactly what intrinsic means. As Sider (1996) notes, philosophers offer many different glosses on what it is to be intrinsic, not all of which are compatible with one another (see also Lewis 1983*a*, 117). And this apparent disagreement gives us reason to think we employ more than one notion of intrinsicity.

Let's use Sider's (1996) version of the Many Notions Response as an example. Sider claims that we have two notions of intrinsic: *qualitative intrinsic* and *non-qualitative intrinsic*. A qualitative intrinsic property is one that an object has in virtue of the *way* it is, while a non-qualitative intrinsic property is one that an object has in virtue of *what* it is. (Sider 1996, 4)⁶⁴ How is this distinction supposed to overcome the worry involving identity properties? Consider the property corresponding to my singleton set: *being me*. According to our qualitative notion, this property is extrinsic, because it is a property I have in virtue of the way I am and the way I am related to other things. According to our non-qualitative notion, *being me* is intrinsic, because it is a property I have solely in virtue of who I am. So, says Sider, when we assess the intrinsicity of a given property, our judgments shift depending on the notion of intrinsic we employ. And since Lewis's account is intended as an account of qualitative intrinsic, the fact that identity properties seem intrinsic according to our non-qualitative notion of intrinsicity does not tell against the analysis.

It is not always legitimate to appeal to the Many Notions Response. If it were,

⁶⁴ Clearly, Sider has in mind the intuitive haecceitistic/non-haecceitistic distinction, not the distinction between properties that are qualitative_B and those that are not.

even an implausible account of intrinsicity could deflect counterexamples by claiming there are different notions in play. Again, consider Kim's (1982) proposal. Kim's account (a property *P* is intrinsic *iff* a lonely object may have *P*) had the counterintuitive result that being lonely is intrinsic. But a defender of the account could respond thus: we have two notions of intrinsicity, *lonely intrinsic* and *non-lonely intrinsic*, and Kim's account is only intended to capture lonely intrinsic. According to our notion of lonely intrinsicity, *being lonely* is intrinsic; so the property *being lonely* is not a counterexample to the analysis.

It should not be this easy to defend an account against counterexamples. Sider acknowledges that there must be constraints on when the Many Notions Response is legitimate. According to Sider, it is legitimate to invoke the Many Notions Response in defense of an account of intrinsicity only when the account in question provides a characterization that "(i) does the work we require of it, and (ii) fits the intuitive gloss we use to pick out intrinsicity in the first place." (Sider 1996, 5)

Does the Duplication Account do the work we require? Although it has counterintuitive results, it's plausible that the account can play a useful philosophical role in many contexts. But the same can be said of Kim's account.⁶⁵ Does the Duplication Account fit the intuitive gloss we use to pick out intrinsicity? Again, although it has counterintuitive results, it's plausible that the account captures our intuitive judgments in many cases. But here, too, the same can be said of Kim's account.⁶⁶

⁶⁵ And even if the Duplication Account is not able to do the work we require, there is an easy solution, says Sider. We can "simply introduce a disjunctive notion – a property is intrinsic in the broader sense *iff* it is intrinsic as defined by [the Duplication Account] or it is an identity property" and use this disjunctive notion whenever the Duplication Account isn't quite adequate (Sider 1996, 6). Obviously, this move is available to a defender of Kim's account, too.

⁶⁶ Sider says that the Duplication Account satisfies criterion (ii) because "the intuitive glosses of 'intrinsic' are capable of qualitative and non-qualitative interpretation"; since the Duplication Account captures the

So even with Sider's two constraints, the Many Notions Response makes it too easy for an account to evade counterexamples. If the Many Notions Response is to be a legitimate defense of the Duplication Account, its advocates need to offer substantive reasons for thinking both (1) that we have multiple notions of intrinsicity, and (2) that one of these notions corresponds to the Duplication Account.

Advocates of the Many Notions Response haven't provided such reasons. Moreover, it's unclear that they *could*. There are two reasons why. First, there doesn't seem to be any way to understand the two notions of intrinsicity that Sider proposes – qualitative intrinsic and non-qualitative intrinsic – that will support this defense of the Duplication Account. Second, the Duplication Account doesn't seem to correspond to any intuitive notion of intrinsicity. Let's consider how advocates of the Many Notions Response might respond to each of these worries in turn.

First consider the non-qualitative notion of intrinsicity, which is supposed to capture the intuition that being me is intrinsic. How should we characterize this notion? We have two options.

One option is to take non-qualitative intrinsic to be an *intensional* notion; i.e. one that applies to properties individuated intensionally. For simplicity, assume there are no duplicate worlds (although nothing hangs on this).⁶⁷ Suppose *being me* is cointensive *with having such-and-such features and so-and-so relations to other things*. Since our non-qualitative notion of intrinsicity classifies *being me* as intrinsic, then – since it's an

qualitative interpretation, it fits an intuitive gloss on intrinsicity (Sider 1996, 6). A defender of Kim's account can make a parallel move. He may say that our intuitive glosses of intrinsic are capable of lonely and non-lonely interpretation; since Kim's account tracks our lonely notion of intrinsicity, his account fits an intuitive gloss on intrinsicity.

⁶⁷ To accommodate duplicate worlds, change the example in the following way. Suppose there is only one duplicate of the actual world, and at this world my duplicate instantiates the identity property *being me₁*. Replace *being me* with *being me or being me₁*, and run the example.

intensional notion – it must also classify *having such-and-such features and so-and-so relations to other things as intrinsic*. But there is no intuitive sense in which *having such-and-such features and so-and-so relations to other things* is intrinsic!⁶⁸ So, if non-qualitative intrinsic is an intensional notion, it cannot correspond to an intuitive notion of intrinsicity, and cannot play a role in this defense of the Duplication Account.

Another option is to take non-qualitative intrinsic to be a *hyperintensional* notion; i.e. one that applies to properties individuated hyperintensionally. Then, cointensive properties may diverge with respect to whether they are non-qualitative intrinsic. Even though *being me* is cointensive with *having such-and-such features and so-and-so relations to other things*, one may say, the latter property is non-qualitative extrinsic while the former is non-qualitative intrinsic. Once one is willing to grant that we have a hyperintensional notion of intrinsicity, however, it's unclear why one would believe we also have an intensional notion. The intuitions captured by the intensional notion are captured by the hyperintensional notion as well; and the hyperintensional notion captures intuitions that the intensional notion cannot accommodate. So the hyperintensional notion is strictly better at capturing our intuitive judgments. Given this, the only reason to insist on both notions is to defend the claim that the Duplication Account captures some intuitive sense of intrinsicity. But that is clearly question-begging, since whether the Duplication Account really does capture an intuitive notion of intrinsicity is one of the points at issue.

Let's turn to the second worry – does the Duplication Account capture *any* intuitive notion of intrinsicity? Even if one could make a convincing case for there being qualitative and non-qualitative notions of intrinsicity, this isn't enough to

⁶⁸ Fine (1994, 7) makes a similar observation with respect to different notions of *essence*.

vindicate the Duplication Account. One must also show that one of these notions corresponds to the Duplication Account. Advocates of the Many Notions Response often turn to identity properties to support their case in the following way: The Duplication Account classifies identity properties like *being me* as extrinsic, but clearly there is some sense in which such properties are intrinsic. So, they say, we must have two intuitive notions of intrinsicity.

The trouble is that there is no intuitive sense in which *being me* is extrinsic. Indeed, once we notice this, it's natural to wonder how we could have thought otherwise. I think this happens when we forget that our intuitive notion of a property is quite separate from the formal apparatus of possible worlds. When one is already familiar with the device of identifying properties with sets of *possibilia*, and realizes that *being me* is cointensive with a purely descriptive property like *having such-and-such features and so-and-so relations to other things* – then, and only then, one can slip into thinking that *being me* is intuitively extrinsic. But once we step back from this metaphysically loaded picture, the belief is hard to maintain.

In sum, the Many Notions Response holds little promise as a defense against the worry about identity properties. Proponents of the Many Notions Response have not provided substantive reasons for believing that we have multiple notions of intrinsicity. And even if we did have multiple notions, the Duplication Account doesn't seem to track one of them.⁶⁹

Finally, how does the Many Notions Response address the worry involving necessary properties? The response does not fare well here, either. The Duplication

⁶⁹ Of course, the Duplication Account may very well track a *useful* theoretic notion. But that does not mean the notion it tracks deserves the name *intrinsic*.

Account classifies necessary properties as intrinsic, but many of the necessary properties in question are qualitative extrinsic (given Sider's characterization of our qualitative notion of intrinsicity). Consider the necessary property *being such that there possibly exists something greater in mass*. According to the Many Notions Response, a property is qualitative intrinsic if an object has it in virtue of the way it is, and qualitative extrinsic if an object has it in virtue of the way it is and the way other things are. An object instantiates the property *being such that there possibly exists something greater in mass* in virtue of the way it is and the way other things are; so, *being such that there possibly exists something greater in mass* is qualitative extrinsic. But the Duplication Account counterintuitively classifies this property as intrinsic. Again, even if one could make a plausible case for there being qualitative and non-qualitative notions of intrinsicity in the manner that Sider suggests, the Duplication Account doesn't track our qualitative notion.

(One might try to spin the Many Notions Response in an alternate way, according to which we have two notions of *property*, not two notions of *intrinsic*. This variation encounters problems as well, although the problems vary depending on which notions of property we employ. If the two notions of property are qualitative and non-qualitative, then this response effectively collapses into the Qualitative Response, which we've already seen is untenable. If the two notions are intensional and hyperintensional, then many of the same objections to the original Many Notions Response apply here as well: there is no evidence that we have a notion of intrinsicity that applies to intensional properties in addition to one that applies to hyperintensional properties, and even if there were, it's unclear the Duplication Account tracks any uniform and intuitive notion of

intrinsicity.)

I submit that the best explanation for our intuitive assessments is not that we have qualitative and non-qualitative notions of intrinsicity, but that ‘intrinsic’ applies to properties individuated hyperintensionally. Conflict arises when we attempt to account for our intuitive judgments of intrinsicity within a framework that conflates intuitively distinct properties. This is not evidence that we have multiple notions of intrinsic or anything else; it is evidence that the framework behind the Many Notions Response lacks the structure necessary to capture our intuitive judgments. *Being me* seems intrinsic and *having such-and-such features and so-and-so relations to other things* does not; and unless we employ a framework that allows for a distinction between the two, our intuitive judgments will continue to clash with whatever account of intrinsic we propose.

4.6 The Spoils to the Victor Response

Finally, consider the “Spoils to the Victor Response.” According to this response, the alleged counterexamples to the Duplication Account rest on inconsistent intuitions. No account can be expected to accommodate inconsistent intuitions, so these “counterexamples” have no force.

The Spoils to the Victor Response generally begins with the observation that a property is either intrinsic or extrinsic. If property P is identical to Q , then P cannot be intrinsic if Q is not intrinsic, or vice versa. If P is intuitively intrinsic and Q is intuitively extrinsic, and P and Q are identical, then one of our intuitive assessments is mistaken – which one? Lewis writes, “When common sense falls into indecision or controversy... then theory may safely say what it likes. Such cases can be left as spoils to the victor.”

(Lewis 1973, 194) So it is with intrinsic, says the proponent of this response. If identical properties differ in their apparent intrinsicity, then we can claim whichever result is compatible with our theory. *To the victor go the spoils.*

Several people have employed the Spoils to the Victor Response in defense of the Duplication Account. Sider writes that an intuitively extrinsic property like “*being such that Socrates is either wise or not wise* is identical to *being round or not round*; since the latter seems intrinsic, the former is as well.” (Sider 1996, 11) Lewis stated in correspondence with Dunn (1990) that his account correctly classifies the intuitively extrinsic *property being a perfect duplicate of b* as intrinsic, because it “amounts to only an infinite conjunction of intrinsic properties, and hence is itself intrinsic.” (Dunn 1990, fn. 7)

Note that the moves Sider and Lewis suggest won’t work unless we assume an intensional conception of properties. If properties are individuated hyperintensionally, then the fact that *being such that Socrates is wise or not wise* is cointensive with *being round or not round* gives us no reason to think that these properties are alike with respect to their intrinsicity. So for the sake of argument, let’s grant that properties are individuated intensionally.

How does the Spoils to the Victor Response address the worry involving identity properties? Consider the property *being me*. Suppose *being me* and *having such-and-such features and so-and-so relations to other things* are cointensive. Therefore, they are identical. But the former property seems intrinsic, and the latter property seems extrinsic. Our intuitions conflict, so we defer to theory: according to the Duplication Account, the property in question is extrinsic.

This move is problematic. If we allow this sort of move, it's too easy to defend oneself from counterexamples. Again, let's consider Kim's (1982) account of intrinsicity (a property P is intrinsic iff a lonely object may have P) and Lewis's (1983a) objection (*being lonely* is intrinsic on Kim's account). We can apply the Spoils to the Victor Response to Kim's account just as easily as to the Duplication Account. A defender of Kim's account could respond thus: It's true that *being lonely* seems like an extrinsic property. But the set of all lonely objects is identical to the property being x_1 or x_2 or x_3 or ... x_n (where the x 's are lonely objects). Since the latter property is intrinsic, the former is as well. Therefore, Kim's account correctly classifies *being lonely* as intrinsic.

And this is just the tip of the iceberg. We can *always* find some intuitively extrinsic property cointensive with an identity property *being x* . Simply tack on *and being located at world w* to *being x* (where w is the world at which x is located), and now we have an extrinsic property that picks out x 's singleton set. We can also do the trick the other way around – every extrinsic property corresponds to some intuitively intrinsic property. Take the extrinsic property *being five feet from a desk*, which is the set of all things that are five feet from a desk. We can pick out the same set by naming each one of the individuals in this set: *being a_1 or a_2 or a_3 or ... a_n* . Now we have an intrinsic property identical to *being five feet from a desk*.

So every set of possible individuals can be identified with both an intuitively intrinsic property and an intuitively extrinsic one. If that's all we need to have a case of "spoils to the victor," then our intuitions place no constraints on an account of intrinsicity. While this defense frees the Duplication Account from intuitive objections,

it does so by making our intuitions irrelevant. If we make this move, we have no reason to prefer the Duplication Account to one according to which all properties are intrinsic, or one according to which all properties are extrinsic, or one which assigns intrinsicity in an entirely arbitrary way. The price of this defense is too dear.

What about the worry involving necessary properties? The same problems arise here. Every necessary property is cointensive with both an intuitively intrinsic property and an intuitively extrinsic property. If the Spoils to the Victor strategy is legitimate, then our intuitions place no constraints on our account of intrinsicity, and we have no reason to prefer the Duplication Account to any other.

4.6.1 The Revised Spoils to the Victor Response

Perhaps we can amend the Spoils to the Victor Response to avoid this consequence. The quotes from Lewis and Sider above suggest an alternative. A set of possible individuals can sometimes be expressed in an intuitively non-qualitative way (*being such that Socrates is wise or not wise*) and an intuitively qualitative way (*being round or not round*). When our intuitions about the property's intrinsicity diverge depending on how we express it, we should adjudicate in favor of the intuitions evoked by the qualitative way.

Let's spell this out in more detail. First, note that this version of the Spoils to the Victor Response requires some sort of qualitative/non-qualitative (or haecceitistic/non-haecceitistic) distinction. And, as with the other responses, it's unclear how the Spoils to the Victor Response can ground this distinction. But let's grant the distinction and see where the response leads. Second, note that we will want to draw the qualitative/non-

qualitative distinction at the level of predicates, not properties. Here's why. Advocates of the Spoils to the Victor Response claim that our intuitions conflict: we judge *being me* as intrinsic and *having such-and-such features and so-and-so relations to other things* as extrinsic, even though these two properties are identical. Of course, they cannot say that the non-qualitative property *being me* is identical to the qualitative property *having such-and-such features and so-and-so relations to other things*, for that would be contradictory – a property cannot be both qualitative and non-qualitative. Instead, they can say that the property corresponding to my singleton set can be expressed with both qualitative and non-qualitative predicates. Expressed by the non-qualitative predicate “being me,” my singleton set seems intrinsic; expressed by the qualitative predicate “having such-and-such features and so-and-so relations to other things,” my singleton set seems extrinsic. Our intuitions conflict.

Here is how the revised version of the Spoils to the Victor Response resolves the conflict. When we can express a property with both qualitative and non-qualitative predicates, and our intuitions about the property's intrinsicness diverge depending on which predicate we use, we should favor the intuitions evoked by the qualitative predicate. In the example above, my singleton set seems extrinsic when expressed by the qualitative predicate “having such-and-such features and so-and-so relations to other things,” and intrinsic when expressed by the non-qualitative predicate “being me.” Since it's the intuitions elicited by the qualitative predicate that should constrain theory, the Duplication Account correctly classifies the property being me as extrinsic. According to this response, Lewis's objection to Kim's account of intrinsicness goes through. The set of all lonely objects can be expressed using the qualitative predicate “being lonely” or the

non-qualitative predicate “being a_1 or a_2 or a_3 or ... a_n ” where each “ a ” is the name of a lonely object. Using the former predicate, the property seems extrinsic; using the latter, it seems intrinsic. It’s the intuitions elicited by the qualitative predicate that our theory should respect, so being lonely should be extrinsic.

The revised version of the Spoils to the Victor Response looks more promising. And it may well be what Sider and Lewis had in mind. But it runs into trouble as well. Consider identity properties. This response does well enough in cases where an identity property may be expressed with both a qualitative and a non-qualitative predicate. But not every identity property can be expressed with both a qualitative predicate and a non-qualitative predicate. And it’s in these cases that the Spoils to the Victor strategy breaks down.

The way in which it breaks down depends on whether we allow duplicate worlds. First, let’s assume there are duplicate worlds. Suppose the actual world has a duplicate, at which there exists a duplicate of me. Both my duplicate and I share the *property having-such-and-such features and so-and-so relations to other things*; but my duplicate does not instantiate *being me*. So, the property *having-such-and-such features and so-and-so relations to other things* is not identical to the property *being me*. In fact, the property corresponding to my singleton set cannot be expressed by any purely qualitative predicate, for any qualitative predicate that picks out me will also pick out my duplicate. Since my singleton set cannot be expressed with any qualitative property, advocates of the Spoils to the Victor Response cannot dismiss the intuition *that being me* is intrinsic.

Next, suppose we deny the existence of duplicate worlds. Consider a symmetrical world with just two things in it: Thing 1 and Thing 2. Thing 1 and Thing 2 are duplicates

that bear the same relations to everything else. Both Thing 1 and Thing 2 share the *property having thing-ish features and thing-ish relations to other things*. But only Thing 1 has the property *being Thing 1*, and only Thing 2 has the property *being Thing 2*. The property corresponding to Thing 1's singleton set cannot be expressed with any qualitative predicate, for any qualitative predicate that picks out Thing 1 will also pick out Thing 2. Likewise for the property corresponding to Thing 2's singleton set. Again, the Spoils to the Victor strategy cannot be applied, and advocates of the Spoils to the Victor Response cannot dismiss the intuition that *being Thing 1* and *being Thing 2* are intrinsic. So this version of the Spoils to the Victor Response doesn't work either. Even if we grant an intensional conception of properties, and the distinction between qualitative and non-qualitative predicates, the response fails to undermine our intuitions that identity properties are intrinsic.

What about the worry involving necessary properties? Here the revised Spoils to the Victor Response fares even worse. The set containing every possible individual can be expressed by many different qualitative predicates, some which suggest the set is intrinsic and others which suggest it is not. For example, "being such that there possibly exists something greater in mass" is a qualitative predicate that intuitively expresses an extrinsic property; while "being round or not round" is a qualitative predicate that intuitively expresses an intrinsic property. Since both predicates are qualitative, we cannot employ the revised Spoils to the Victor strategy to adjudicate between them.

4.7 Analysis and Conceptual Possibility

So far, I have been assuming that the space of possibility is metaphysical possibility. But

in Chapter (2), I argued that we should evaluate a candidate analysis of a notion against the backdrop of *conceptual possibility*, not metaphysical possibility. Sometimes an account fares worse when we move to the space of conceptual possibility, because there are conceptually possible counterexamples that are not metaphysically possible. Sometimes an account fares better, because a counterexample that arises when the space of possibility is restricted in certain ways disappears when these restrictions are lifted.

One might object that if the Duplication Account appears to face trouble, it's only because we've been assessing it with respect to the wrong kind of possibility. So let's evaluate the account with respect to conceptual possibility, not metaphysical possibility. When we move to the space of conceptual possibility, do these troubles disappear?

It doesn't seem so. First consider identity properties. Even when we move to the space of conceptual possibility, it is still the case that the property *being me* will not be shared by all of my conceptually possible duplicates. So the Duplication Account still incorrectly classifies *being me* as extrinsic.

Next, let's look at necessary properties. Consider the property *being such that there possibly exists something greater in mass*. Can we ideally conceive of something that does not instantiate this property? This might be *prima facie* conceivable, but arguably it is not ideally conceivable. Given a unit of measurement, the possible mass values correspond to the positive real numbers. There is no greatest real number. So no matter how massive an object is, it's always possible that something else is *more* massive. Thus it's not conceptually possible for an object to lack this property. Because *being such that there possibly exists something greater in mass* does not divide conceptually possible duplicates, the Duplication Account incorrectly classifies this

property as intrinsic.

(One might argue that we *can* ideally conceive of something lacking the property *being such that there possibly exists something greater in mass*. Perhaps it is ideally conceivable that the possible mass values correspond to the extended positive real numbers (the positive real numbers extended to include positive infinity). If so, then there are possible objects with infinite mass, and these objects do not have the property *being such that there possibly exists something greater in mass*. In general, however, this move won't help the Duplication Account. Let us define *mass** to be the property of having a mass equal to some positive real number of grams. There are no possible objects with infinite *mass**, and so there are no possible objects that lack the property *being such that there possibly exists something greater in mass**. Since this property never divides conceptually possible duplicates, the Duplication Account incorrectly classifies it as intrinsic.)

Likewise, consider the property *being such that Socrates is wise or not wise*, which, as Sider (*op. cit.*) notes, seems to be an extrinsic property. This property is shared by all conceptually possible duplicates, and thus is intrinsic on the Duplication Account. But that appears to appear to be the wrong result.

Of course, we could try to finesse our intuitions about these properties by appealing to one of the moves I've examined in this chapter. As we've seen, however, these moves are not tenable. So these properties stand as counterexamples to the Duplication Account, even when we assess the analysis using the space of conceptual possibility.

4.8 Individuating Properties Hyperintensionally

Although these responses do not succeed in defending the Duplication Account, they help pinpoint what is wrong with it.

All three responses acknowledge that *being me* seems intrinsic and *having such-and-such features and so-and-so relations to other things* does not – even though the two properties correspond to the same singleton set. *Prima facie*, this acknowledgment suggests that our intuitive notion of intrinsicness is sensitive to hyperintensional differences among properties. If this is so, then no account of intrinsicness that is unable to accommodate the relevant hyperintensional differences will be able to accommodate our intuitive beliefs.

And that is the problem with the Duplication Account. While the account works well enough in cases where hyperintensional distinctions aren't relevant, it breaks down exactly where we would expect: in cases where properties that correspond to the same set of individuals fall on different sides of the intrinsic/extrinsic divide.

4.8.1 Perfectly Natural, Qualitative_B, Qualitative_O, Qualitative_H

We've seen reason to think that *intrinsic* should be applied to properties individuated hyperintensionally: the intrinsic/extrinsic distinction discriminates between properties that are intensionally equivalent. Given this, it's natural to ask whether the other distinctions between properties discussed in previous chapters – naturalness, qualitative_B, qualitative_O, and qualitative_H – also discriminate between intensionally equivalent properties.

At least some intuitions support the claim that they do. Consider the perfectly

natural properties. Suppose *mass* is a perfectly natural property that corresponds to the set of all things with mass. This property is cointensive with the property *being a_1 or a_2 or a_3 or... a_n* – where the *a*'s name all the possible objects with mass. But the latter property, *being a_1 or a_2 or a_3 or... a_n* , does not seem to be perfectly natural. If one has the intuition that the former is perfectly natural while the latter is not, then one has some reason to take naturalness to discriminate between intensionally equivalent properties.

(In fact, if naturalness applies to properties individuated hyperintensionally, and if the perfectly natural properties comprise a minimal supervenience base, then “perfectly natural” *must* discriminate between intensionally equivalent properties. Take the pair of cointensive properties considered above: *mass* and *being a_1 or a_2 or a_3 or... a_n* . Suppose both were perfectly natural. Then the perfectly natural properties would fail to form a minimal supervenience base: we could subtract *mass* or *being a_1 or a_2 or a_3 or... a_n* and still have the same supervenience base as before. So either *mass* or *being a_1 or a_2 or a_3 or... a_n* is redundant. Since only one may be perfectly natural, it follows that naturalness discriminates between cointensive properties.)

Now consider the kind of properties we're naturally inclined to associate with qualitative_B , qualitative_O , and qualitative_H . *Mass*, for instance, seems to be qualitative in all three senses. But *being a_1 or a_2 or a_3 or... a_n* – the property cointensive with *mass* – does not. Again, if one has these intuitions, then one has reason to take qualitative_B , qualitative_O , and qualitative_H to discriminate between intensionally equivalent properties.

Those who are swayed by these intuitions will not be satisfied by the characterizations of qualitative_B , qualitative_O , and qualitative_H that I gave in Chapter (2). According to those characterizations, anything that supervenes on the qualitative_B ,

qualitative_O, or qualitative_H is itself qualitative_B, qualitative_O, or qualitative_H. But characterizations in terms of supervenience cannot discriminate between cointensive properties. That reason is that for any property P , if P supervenes on X then anything cointensive with P also supervenes on X .

More carefully: Suppose P and Q are cointensive properties. And suppose that P supervenes on X . Does Q also supervene on X ? Suppose for *reductio* that it does not. Then there is a pair of possible worlds alike with respect to X and different with respect to Q . Since P and Q are cointensive, it follows that there is a pair of possible worlds alike with respect to X and different with respect to P . But that entails that P does not supervene on X , *contra* our supposition. So if P supervenes on X , then anything cointensive with P must also supervene on X .

So we've seen that the characterizations of qualitative_B, qualitative_O, and qualitative_H I gave in Chapter (2) will not allow these notions to discriminate between cointensive properties. Those who think that these notions discriminate between cointensive properties should replace the role played by supervenience in these characterizations with an appropriately discriminating surrogate.⁷⁰ Since nothing I've said hangs on whether these notions can discriminate between cointensive properties, I shall remain neutral as to which characterization is best.⁷¹

⁷⁰ There are several ways to do this. The easiest is to replace the notion of supervenience with a primitive notion of "grounding," "in virtue of," or "ontological dependence." (c.f. Kim 2005 and Schaffer 2009)

⁷¹ In the interest of explicitness, here's why taking these notions to discriminate between cointensive properties has no impact on my assessment of the Duplication Account. First, even if naturalness discriminates between cointensive properties, this won't help the Duplication Account. The Duplication Account takes the intrinsic properties to individually supervene on the perfectly natural, but as we've seen, supervenience doesn't allow for discrimination between cointensive properties.

What if the qualitative notions discriminate between cointensive properties? One might think that the Qualitative Response is in a better position if we take qualitative_B to discriminate between cointensive properties. For then we can restrict the properties to which the Duplication Account applies to just the qualitative_B properties; and if qualitative_B is appropriately discriminating, then properties like *being me*

4.9 Assessing the Costs

To sum up: there is no way to defend the Duplication Account against counterexamples arising from identity properties and necessary properties.

What to do? One option is to give up on *intrinsic* altogether. Given that the distinction between intrinsic and extrinsic properties plays a major role in many areas of philosophy – broad and narrow content, intrinsic value, epistemic internalism, change over time, etc. – this option is not appealing.

Another option is to accept the Duplication Account despite its counterintuitive results. Those who do so need to provide a justification for thinking that the notion picked out by Duplication Account (1) fits our use well enough to deserve the name “intrinsic,” and (2) is the best deserver of the name “intrinsic.”

A third option is to adopt some other analysis of intrinsicity. The most promising analyses employ a further primitive notion (e.g. see Trogon 2009). A final option is to take *intrinsic* itself as primitive. Such accounts have an obvious drawback, of course: they require an additional primitive notion. Nonetheless, it is worse to settle for an account that is false.

will not be qualitative_B even though the property cointensive with *being me – having such-and-such properties and so-and-so relations to other things* – is qualitative_B. But the Qualitative Response is not in a better position. First, if we take qualitative_B to be appropriately discriminating, then we cannot characterize it in terms of supervenience on the perfectly natural. And the natural alternatives require some further primitive (see footnote (17)). If we need to posit an additional primitive notion to get the Duplication Account to apply to the right properties, then we might as well take *intrinsic* as primitive and be done with it. Second, the charge of arbitrariness still applies: our only motivation for restricting the account to qualitative properties seems to be the desire to avoid counterexamples. Third, the charge of undermining still applies. The Qualitative Response requires a conception of properties that allows for cointensive properties to be distinct. But once one is willing to grant that properties are individuated hyperintensionally, why would one desire an account of intrinsicity that cannot recognize hyperintensional distinctions?

Part II: The Structure of Fundamental Properties

In the first part of this dissertation, I explored what the fundamental properties can and cannot do for us. In the second part of this dissertation, I will explore what the fundamental properties are and are not like. Before I asked: what work can we do with the perfectly natural properties? Now I ask: what structure do these perfectly natural properties have?

In Chapter (5), I examine the topic of temporary intrinsics. Temporary intrinsics are apparently intrinsic properties that an object may have at one time and not another. Lewis (1986*b*) argues that endurantism – the view, roughly, that objects do not have temporal parts – cannot account for temporary intrinsics. I examine two different versions of the argument from temporary intrinsics, and show that neither has much force. I present a new version of the argument, which I believe to be a more promising articulation of the perdurantist’s objection. The key premise of this argument is this: other than the spatiotemporal relations, the perfectly natural properties at our world are monadic. But if we think that fundamental physics is a guide to the perfectly natural properties instantiated at our world, then we should think that the structure of the perfectly natural properties is in part an empirical matter. I argue that the key premise is false, and the argument is unsound. I conclude that no version of the Argument from Temporary Intrinsics gives us a compelling reason to favor one theory of persistence over another.

In Chapter (6), I examine quantitative properties. Quantitative properties have three distinctive features. First, they are gathered into *families*; second, quantities within

families are *ordered*; third, quantities within families are certain *distances* from one another. What is the metaphysical basis for these features? In this chapter, I focus on Armstrong's proposal that the fundamental quantitative properties are structural universals. These universals have a distinctive structure akin to mereological structure. Armstrong claims that the distinctive structure of these universals accounts for the features of quantitative properties. I argue that that Armstrong's proposal fails. The structure he proposes cannot account for the distinctive features of quantitative properties. We need something more.

Chapter 5:

Temporary Intrinsic and Gauge Properties

5.1 The Problem

I am bent at one time and straight at another. But I cannot be both bent and straight, since then I would instantiate contradictory properties. So what underwrites this change?

Prima facie, an object undergoes change when it gains or loses a certain kind of property. These properties are often called *temporary intrinsics*. The term *temporary intrinsics* can be misleading, however, since whether such properties are intrinsic will be one of the questions at issue. To avoid confusion, I will use the term *ephemera* to refer to those properties and relations involved in intuitively ordinary matters of change.

Two popular theories of persistence – perdurantism and endurantism – provide different accounts of change. According to the perdurantist, objects have temporal parts, each of which may instantiate different properties. An object changes when one of its temporal parts has an *ephemeron* that another one lacks. In the case above: one of my temporal parts instantiates *being bent* and another instantiates *being straight*. Since different temporal parts of me instantiate different ephemera, I undergo change.

The endurantist gives a different response. According to her, objects are not conglomerations of temporal parts. Objects endure: they are wholly present at every time at which they exist. How can an object – the whole of it – instantiate *being bent* as well as the *being straight*? By instantiating these properties at different times. For the endurantist, an object's ephemera are two-place relations that hold between objects and

times. In the case above: I instantiate *being bent* relative to time t_1 and *being straight* relative to time t_2 . Since I am bent at t_1 and straight at t_2 , I undergo change.⁷²

David Lewis has famously argued that this result constitutes a *reductio* of endurantism. (Lewis 1986*b*, 202-204) When an object undergoes change, this has to do with how the object is in itself and not how it is relative to anything else. But this isn't the case on the endurantist's account. According to endurantism, an object's change has to do with how it is in itself *and* how it is in relation to something else – a time. Since the endurantist fails to respect the observation that an object's change does not involve anything besides the object itself, endurantism is not a tenable position. This argument is known as the Argument from Temporary Intrinsic.

In response to the Argument from Temporary Intrinsic, endurantists have constructed a number of ways to make their account of change, and hence their account of ephemera, more intuitively acceptable. And, in response to these moves, Lewis has proposed a more sophisticated version of the Argument from Temporary Intrinsic (what I call the Argument from Having *Simpliciter*). I show that neither the standard argument nor the more sophisticated version has much force.

But that is not the end of the story. There are other ways to understand the Argument from Temporary Intrinsic. The most promising way, I believe, is to understand it as a debate about the nature of the fundamental properties. Drawing on lessons from gauge theory, I show that even this version fails. I conclude that no version of the Argument from Temporary Intrinsic gives us a compelling reason to favor one theory of persistence over another.

⁷² In this chapter, I assume that the truth about the world can be stated in a timeless language, and so I will not engage with those who take tense as fundamental.

5.2 The Argument from Temporary Intrinsic

The standard way of formulating the perdurantist's objection begins with the notion of *intrinsic*. Intuitively, ephemera are intrinsic properties, or properties that an object has in virtue of itself alone. There are a number of characterizations of intrinsic, but since I will focus on evaluating Lewis's argument against endurantism, I will employ Lewis's characterization. This is the Duplication Account presented in Chapter (4): a property is intrinsic *iff* it never differs between duplicates. (Lewis 1983*b* and 1986*b*, 61-62) I may have a duplicate that is five feet from a dog, or has a nephew, or differs in any number of extrinsic ways from myself. But no duplicate of me may differ in sex, lack an appendix, or have an extra limb.

Now we can flesh out Lewis's objection to the endurantist. Suppose I am bent at time t_1 . According to the perdurantist, this consists in my having a bent temporal part that exists at time t_1 . Any duplicate of that temporal part likewise instantiates *being bent*. Since the property *being bent* never varies between duplicates, it is intrinsic.

Not so for the endurantist. If I am bent at time t_1 , then, according to the endurantist, I instantiate the property *being-bent-at- t_1* . But not every duplicate of me instantiates this property. I may have a future duplicate that is bent at time t_2 instead of t_1 , and this duplicate will instantiate *being-bent-at- t_2* instead of *being-bent-at- t_1* . So properties such as *being-bent-at- t_1* and *being-bent-at- t_2* may vary between duplicates, and therefore are extrinsic.⁷³ If endurantism is true, none of my ephemera are intrinsic.

This is the Argument from Temporary Intrinsic:

P1: Ephemera are intrinsic.

⁷³ It's true the property *being bent at some time* is intrinsic for the endurantist, since it never varies between duplicates. But the Argument from Temporary Intrinsic addresses a worry about change, and hence temporary properties. *Being bent at some time* is not a temporary property, even if it is intrinsic.

P2: If endurantism is true, then ephemera are not intrinsic.

Conclusion: Endurantism is false.

5.2.1 Endurantist Replies

One reply available to the endurantist is to reject the definition of *intrinsic* in play. Since Lewis's characterization fails to render the paradigmatic ephemera intrinsic, one could argue that his characterization is unsatisfactory. The endurantist might propose a different account; perhaps something along the following lines will do:

Intrinsic_E: A property is intrinsic_E iff it never divides duplicates located at the same time.

Suppose I am bent at time t . Any duplicate of me located at t will also be bent. So, given the endurantist's account of persistence, *being bent* is intrinsic_E – as it intuitively should be.⁷⁴

Alternatively, the endurantist can implement her picture by relativizing the *instantiation* relation, rather than relativizing the ephemera themselves. The more familiar two-place instantiation relation is replaced with a three-place relation that holds between objects, properties, and times. If I am bent at t_1 , then the instantiation relation holds between me, *being bent*, and t_1 . If there is a bent duplicate of me at t_2 , the instantiation relation holds between it, *being bent*, and t_2 . All of my duplicates instantiate *being bent*, though they may instantiate it at different times from me. So, on this account, ephemera such as *being bent* are monadic and intrinsic, just as they seem.

⁷⁴ See also Haslanger (1989, 123) for an endurantist-friendly characterization of *intrinsic*.

Or, the endurantist can follow Haslanger (1989) in making the truth-value of propositions relative to times. For example, if I am bent at t_1 and straight at t_2 , then the proposition that I am bent is true at t_1 and false at t_2 . Again, on this account, ephemera such as *being bent* are monadic and intrinsic.⁷⁵

But these responses feel less than satisfying. The motivation behind the Temporary Intrinsic Argument is the conviction that an object has its ephemera in and of itself alone. The endurantist may contrive a sense in which her ephemera are intrinsic and monadic, but nonetheless these properties seem unacceptably relational. The intuitive objection still stands, and the perdurantist needs a better argument to capture it.

5.3 The *Simpliciter* Argument

Enter “having *simpliciter*,” a notion invoked by David Lewis to better capture this intuitive worry about endurantism.⁷⁶ When object a has property F *simpliciter*, this has nothing to do with anything besides a and F and simple two-place instantiation. For the perdurantist, the properties involved in intuitively intrinsic matters of change are had *simpliciter*. The endurantist relativizes these properties to times (either by turning them into relations or relational properties, or by temporally relativizing the instantiation relation), and as a result they are not had *simpliciter*.

Call this the *Simpliciter* Argument:

P1: Ephemera are had *simpliciter*.

P2: If endurantism is true, then ephemera are not had *simpliciter*.

⁷⁵ See Caplan (2005), Haslanger (1989, 120), Hinchliff (1996), Lewis (1988, 65-66 fn. 1) and (2002), and van Inwagen (1990) for discussion of the latter two moves.

⁷⁶ See Lewis (1986b, 52-54); also see Lewis (1988) and (2002).

Conclusion: Endurantism is false.

This argument is meant to capture the ingredient missing from the standard Argument from Temporary Intrinsic. The endurantist may contest certain tendentious glosses on “intrinsic” and may make instantiation a three-place relation, but still her ephemera are not had *simpliciter*. Why? Because when an object has a property *simpliciter*, this consists only of a thing, a property, and two-place instantiation.

5.4 The *Simpliciter* Argument and the Paradigmatic Temporary Intrinsic

Although the endurantist cannot respond in the same way to the *Simpliciter* Argument as she could to the Temporary Intrinsic Argument, there are other reasons to worry about the *Simpliciter* Argument.⁷⁷

The force of the argument is proportional to the force of our intuitions about the nature of canonical ephemera. Intuitively, as Merricks (1994) says, it isn't the case “that all of the properties that an object seems to gain or lose are really relations to times or time indexed. A short list of those properties which are not... includes shape, color, size, and mass.” (Merricks 1994, 528)

That *seems* right. But is it? Arguably, none of these paradigmatic ephemera are had *simpliciter*. Consider *mass*. Given the advent of relativity, there are two viable notions of mass: *rest mass* and *relativistic mass*. Unlike rest mass, relativistic mass is dependent on reference frame: an object that has a relativistic mass of two grams in one

⁷⁷ There are other responses to the *Simpliciter* Argument that I am not considering here. Wasserman (2003) offers several different responses on behalf of the endurantist. In particular, it seems that the perdurantist can no more accommodate our intuitive beliefs about ordinary objects than the endurantist. Both must deny that ordinary objects have their ephemera *simpliciter* (for the perdurantist, a persisting object may have a part that is bent *simpliciter*, but it is not itself bent *simpliciter*). And, as Wasserman shows, just about any move the perdurantist may make in an attempt to mitigate this counterintuitive result is available to the endurantist as well.

frame may have a relativistic mass of two hundred grams in another. So objects don't have relativistic mass *simpliciter*; they have relativistic mass *relative to a reference frame*.⁷⁸ If, as Field (1973) argues, our ordinary use of "mass" does not determinately pick out rest mass as opposed to relativistic mass, it's unclear whether objects have mass *simpliciter*.

Or consider *color*. Many philosophers have argued that color properties are relational in some way. For instance, Jackson (1998) argues that the color of an object depends on "a certain kind of creature and circumstances of viewing." (Jackson 1998, 95) If so, objects do not have color *simpliciter*; they have their colors relative to a creature and circumstances of viewing.

Or consider *shape* and *size*. The spatial shape and size of an object supervenes on the spatial relations between the object's parts. But in a relativistic world, these relations are frame-dependent. A sphere in one reference frame may be an ovoid in another; a gallon jug in one frame may hold a half-gallon in another. Objects don't have their spatial shapes and sizes *simpliciter*; they have them *relative to a reference frame*.⁷⁹

In sum, it's arguably the case that none of the paradigmatic ephemera are had *simpliciter*, regardless of one's preferred theory of persistence. This doesn't rule out the possibility that some ephemera are had *simpliciter* (if endurantism is false). But the further we move away from the canonical ephemera, the weaker the argument becomes. The *Simpliciter* Argument is only as strong as the strength of the intuitions it captures; and while we seem to have firm intuitions about canonical ephemera like *shape* – "if we

⁷⁸ On some views of mass, an object's rest mass is nothing more than its relativistic mass relativized to its own frame of reference. In that case, there is no notion of *mass* according to which it is had *simpliciter*. (I thank an anonymous referee for this point.)

⁷⁹ Skow (2007) offers a different reason to think that shape is not intrinsic.

know what shape is, we know it is a property, not a relation” (Lewis 1986*b*, 204) – our intuitions about other ephemera are more tenuous. At the very least, the tentative nature of the argument gives us motivation to look for something better.

More importantly, the fact that the argument is vulnerable to the sort of piecemeal rebuttal given above suggests that it isn’t getting at the right intuition. The debate over persistence shouldn’t depend, it seems, on a miscellany of unrelated issues. The perdurantist’s objection is more general: even if canonical ephemera like *color* and *shape* aren’t had *simpliciter*, then surely there is something underlying these properties which is. It is this intuition that I aim to capture with the argument I present in section IV. But let’s first examine some replies on behalf of the perdurantist.

5.5 Two Perdurantist Replies

I’ve found that perdurantists who endorse the *Simpliciter* Argument tend to raise two responses to this criticism. The first response is to insist that there is a sense in which some of the canonical ephemera – particularly *shape* and *size* – are had *simpliciter*. The second response is to change the desiderata for a satisfactory account of ephemera, and amend the *Simpliciter* Argument in light of this.

I think neither response is successful. Moreover, I’m not sure that any reply along these lines could help the *Simpliciter* Argument. I’ll address the individual responses first, and then come back to the deeper worry.

5.5.1 Senses of “Shape”

The perdurantist who offers the first sort of response might claim that, in the context of

ephemera, the relevant notion of *shape* isn't spatial (or three-dimensional) shape, since spatial shape is not had *simpliciter*. Rather, the relevant notion of *shape* is four-dimensional shape. Four-dimensional shape is had *simpliciter*, says our perdurantist; it does not involve relations to reference frames, times, or anything else.⁸⁰ So the response goes: any theory of persistence must capture the intuition that an object has its four-dimensional shape *simpliciter*. Since the perdurantist can capture this intuition while the endurantist cannot, perdurantism is to be preferred.

This reply fails for several reasons. First, our pretheoretic notion of shape is spatial, not spatiotemporal, so the intuitions about having *simpliciter* that the perdurantist invokes apply to three-dimensional shape, not four-dimensional shape. Now, the perdurantist might be tempted to reply that our intuitions about having *simpliciter* do apply to four-dimensional shape. After all, the notion of four-dimensional shape isn't so much harder to grasp than that of three-dimensional shape. Most of us understand the relationship between two- and three-dimensional shape, and four-dimensional shape can be described analogously: simply increase the number of dimensions by one, and call the newest addition "time." Even students in an introductory philosophy class unfamiliar with the notion of four-dimensional shape have little trouble grasping the concept of a spacetime "worm" or an object's "worldline."

But this line of thought is plausible only if one conflates two different notions of "four-dimensional shape." Call the first notion "Newtonian" and the second "Relativistic." Think of four-dimensional Newtonian shape as the shape of an object as it sweeps through time. Just as a two-dimensional coordinate system (x, y) can be expanded into a three-dimensional coordinate system (x, y, z) by adding on a z axis, so a

⁸⁰ Ted Sider has raised this response on behalf of the perdurantist in conversation.

three-dimensional coordinate system can be expanded into a four-dimensional Newtonian coordinate system (x , y , z , and t) by adding on a *time* axis. The four-dimensional Relativistic shape of an object, on the other hand, is just the information encoded by the spatiotemporal intervals between each of the object's parts.

I'll grant that four-dimensional Newtonian shape is an intuitively accessible notion of shape. But Newtonian shape isn't had *simpliciter*. An object's Newtonian shape depends on the spatial distances between its parts, and these facts are frame-dependent. On the other hand, Relativistic shape *is* had *simpliciter*. An object's Relativistic shape depends on the spatiotemporal intervals between each of its parts, and these facts are not frame-dependent. But we don't have an intuitive understanding of these spatiotemporal intervals, and *a fortiori*, we don't have an intuitive understanding of four-dimensional Relativistic shape. (Those who believe our ordinary intuitions about distance apply to these relativistic spatiotemporal intervals should reconsider: unlike ordinary spatial distances, spatiotemporal distances can be both real and imaginary, and unlike ordinary spatial distances, there can be an infinite number of distinct locations whose spatiotemporal distance is zero.)

The perdurantist's initial complaint against the endurantist was that she fails to respect the observation that certain properties of everyday acquaintance, such as shape, are had *simpliciter*. This line of response attempts to extend the complaint to four-dimensional shape, but it does not succeed. If by "shape" the perdurantist means "four-dimensional Newtonian shape," then shape is not had *simpliciter*. And if by "shape" she means "four-dimensional Relativistic shape," then shape is not a property of everyday acquaintance.

There's another reason that shifting to four-dimensional Relativistic shape does not help the perdurantist. The endurantist was initially faulted for her account of the properties involved when an object undergoes change. But four-dimensional Relativistic shape is not a property involved in matters of change: no object can have one four-dimensional shape at time t_1 and another at time t_2 . Since Relativistic shape is not a temporal property, there is nothing to stop the endurantist from agreeing that objects have their four-dimensional Relativistic shapes *simpliciter*.

In sum, both the perdurantist and the endurantist must admit that no object has its spatial extension *simpliciter*, and both may say that objects have their frame-invariant shapes *simpliciter*. The strategy of shifting the properties in question to frame-invariant ones does not make the *Simpliciter* Argument any more effective.

5.5.2 Time and Having *Simpliciter*

The perdurantist who endorses the second sort of response does not attempt to contrive a sense in which ephemera like *mass*, *color*, and *shape* are had *simpliciter*. Instead, this perdurantist objects to the way in which these properties are not had *simpliciter* on the endurantist account. That ephemera should turn out to be relational is not itself objectionable; what's objectionable is that these relations should involve *time*. The thought is this: whatever properties like *mass*, *color*, and *shape* turn out to be, they are not inherently temporal notions. Since the endurantist claims that they *are* inherently temporal, endurantism is untenable.

It's unclear why time should be a particularly egregious relatum. In any case, many of these properties *are* inherently temporal. Given relativity, any property which

varies depending on one's reference frame has a critical temporal component. This is especially easy to see with properties like *spatial shape* and *size*. Maudlin (2002)

describes this well:

Take a car and a tunnel which, when at rest relative to one another, are exactly the same length. Now get in the car and drive it through the tunnel. According to the tunnel, the car is moving and therefore suffers a contraction: the car should fit entirely inside the tunnel. On the other hand, according to the car the tunnel is moving, so the car should now be longer than the tunnel. No matter how objects shrink or grow, how can it both be the case that the car is longer than the tunnel and that it is shorter than the tunnel?...

The story of the car and the tunnel illustrates how seemingly nontemporal notions may be infected by hidden temporal aspects. The question of whether the car or the tunnel is longer appears to be a question purely about the spatial qualities of the two objects... But once we see that the car being longer than the tunnel is a matter of the front and back ends of the car being outside the tunnel *at the same instant*, it becomes clear how the relativity of simultaneity resolves the tension between the two judgments. Since the different observers disagree on which sets of events constitute an instant, they may disagree on whether there is an instant when every part of the car is within the tunnel. (Maudlin 2002, 53-55)

Our everyday notions do not always recognize the role time plays in properties like *shape* or *distance*. But the intuition that such properties are divorced from time is mistaken, and it should not count as a mark against the endurantist that she cannot capture it.

5.6 The Fundamental *Simpliciter* Argument

And yet the suspicion that there is something wrong with the endurantist's ephemera persists. Why?

For simplicity, assume that the endurantist's picture characterizes ephemera as relations between objects and times (as opposed to the adverbial variant of endurantism, which relativizes the instantiation relation⁸¹). There is nothing objectionable about these relations *per se*. After all, the perdurantist accepts the very same relations. I am not happy *simpliciter*, says the perdurantist; I bear the *happy-at* relation to every time at

⁸¹ The endurantist has a bit more trouble if her "temporary intrinsics" are relativized to times using a three-place instantiation relation. Unlike relations to times, most perdurantists do not appear to accept a three-place instantiation relation in addition to a two-place instantiation relation.

which I have a happy temporal part.⁸² So the perdurantist worry cannot simply be that the endurantist posits relations involving objects and times.

Lewis states that the endurantist's relations to times (and the relational properties built out of them) are acceptable so long as they are "not alleged to be fundamental properties of the sort that might figure in a minimal basis on which all else supervenes." (Lewis 2002, 4) Perhaps this is the source of the worry. What's wrong with the endurantist account is not so much that it posits these relations to times, but that it posits them as *fundamental*. The fundamental relations instantiated at our world do not include those required by the endurantist account of persistence.⁸³

We can capture this perdurantist objection to endurantism with a third version of the Argument from Temporary Intrinsic: the Fundamental Argument.

P1: There are no fundamental relations involving objects and times instantiated at our world.

P2: If endurantism is true, then fundamental relations involving objects and times *are* instantiated at our world.

Conclusion: Endurantism is false.

While the previous versions of the Temporary Intrinsic Argument invoked intuitions about the properties and relations involved in ordinary matters of change, this version makes a claim about the underlying properties and relations. For the endurantist, relations involving objects and times are fundamental; for the perdurantist, they are not.

⁸² Lewis writes, "I cannot object to these relations and relational properties...I accept similar relations and relational properties myself" (Lewis 2002, 4).

⁸³ The one exception may be the *occupation* relation that holds between objects and the spacetime regions they occupy. The exception does not arise for the perdurantist who identifies objects with spacetime regions.

This version of the Temporary Intrinsic Argument avoids the problems that beset the other two, and I believe it gets at the core worry behind the original Temporary Intrinsic Argument.

5.7 Fundamental Relations involving Objects and Times

(P1) is a claim about the nature of our fundamental ontology. But whether there are fundamental relations involving objects and times is partly an empirical matter. As it turns out, the gauge theories of standard particle physics – theories that describe the *gauge properties* of fundamental particles – require irreducible relations of the sort (P1) rules out.

5.7.1 Cannons, arrows, and quarks

We can bring out the relevant features of gauge theories by looking at how to compare gauge properties. Fortunately, we don't need to get bogged down in the details of gauge theory. Mathematically, comparing gauge properties is directly analogous to comparing the directions of vectors. So let's turn away from gauge theory for a moment, and look at how to compare the directions of vectors.⁸⁴

Suppose you and a friend are visiting the Castle of Carcassonne in France. According to your guidebook, the cannon in the Northern tower points in the same direction as the cannon in the Southern tower. But your friend is skeptical; she thinks the cannons are skewed with respect to one another. How might you resolve the matter?

Unfortunately, you've lost your compass. But you notice that there is a stack of tapered boards lying nearby, where construction workers are renovating the rampart that


⁸⁴ My discussion of chromodynamics draws heavily on Maudlin (2007).

connects the two towers. Here's a way you can determine whether the cannons point in the same direction. Go to the cannon in the Northern tower and lay down one of the boards parallel to it, one step to the south, with the tapered end pointing in the same direction as the cannon. Have your friend check to make sure you haven't twisted or turned the orientation of the board relative to the cannon. When she's satisfied, take another step towards the Southern tower, and put down a second board parallel to the first. Once your friend has checked that they're parallel, lay down a third board, and so on, until you reach the cannon in the Southern tower. Eventually, you have a straight path of boards running from one cannon to the other. Since each cannon is parallel to the board next to it, and the boards are parallel to each other, you deduce that the two cannons point in the same direction.

This situation is directly analogous to comparing the direction of vectors. To compare the directions of two vectors, we need to shift, or "parallel transport," one of the vectors over to the other. If we replace the cannons with vectors and the rampart with a two-dimensional plane, that's essentially what we've done in the case above.


In the example of the Castle of Carcassonne, this was relatively straightforward. The rampart connecting the two cannons was approximately flat, so we were able to give an unequivocal answer as to whether the cannons point in the same direction simply by laying out boards in the way described. But things are not so straightforward when the surface involved isn't flat.

Consider how you might compare the directions of a pair of vectors lying on the surface of a sphere, as shown in the following diagram:



One of the vectors is located at the North Pole, and the other at the South Pole. Do these two vectors point in the same direction?

If we parallel transport the vector at the North Pole to the South Pole along path α , as shown in the diagram below, then the two vectors will point in the same direction. But if we parallel transport the vector at the North Pole to the South Pole along path β , then the two vectors will point in *opposite* directions. On curved surfaces, whether the direction of one vector is the same as the direction of another depends on the path along which they are parallel transported. Two vectors on a sphere don't point in the same direction *simpliciter* – they point in the same direction *relative to a path*.



So we can see that assessing the directions of vectors in a space depends on the path through that space used to compare them. In the Carcassonne example, the space is four-dimensional spacetime, so the paths involved are *spatiotemporal* paths. Although we compared the directions of the two cannons at the same time, we could have made any number of other comparisons, such as cross-time comparisons (e.g. compare the direction of a cannon to itself, five minutes later) or comparisons across time and space (e.g. compare the direction of one cannon to the direction of the other five minutes later).

Now consider the gauge theory of quantum chromodynamics, according to which quark particles have one of three different “color” properties. The comparison of the *quark colors* of quarks is directly analogous to the comparison of the directions of vectors: it depends on the path along which they are compared. Two quarks don’t have the same color *simpliciter* – they have the same color *relative to a path*.

What are the fundamental properties and relations that ground these facts about *quark colors*? One natural account appeals to fundamental relations between pairs of quarks and a spatiotemporal path connecting them. Another appeals to fundamental

relational properties, such as *having the same color as quark q along path p*. Or perhaps the the best account of quark color will invoke fundamental relations between pairs of quarks and *every* spatiotemporal path connecting them. But whatever account of gauge properties we ultimately settle on, the underlying gauge facts appeals to relations involving objects and spatiotemporal paths.

Recall (P1) of the the Fundamental Argument: there are no fundamental relations involving objects and times instantiated at our world. But we've seen that the gauge theories of particle physics require fundamental relations between objects and paths in spacetime. *A fortiori*, the gauge theories require fundamental relations involving objects and time. And the fact that this is so has nothing to do with the metaphysics of persistence. Since (P1) is false, the Fundamental Argument is unsound.

5.7.2 A Perdurantist Response

The perdurantist might respond by protesting that the endurantist's relations to times are different from the relations required by gauge theories. The endurantist's relations hold between objects and times, while the relations underlying gauge properties hold between objects and spatiotemporal paths. Says this perdurantist: True, relations to spatiotemporal paths involve time, but they do not involve time in the same direct way as the endurantist's relations involve time. All the perdurantist has to do is slightly modify the Fundamental Argument:

P1*: There are no fundamental relations that hold *between* objects and times instantiated at our world.

P2*: If endurantism is true, then fundamental relations that hold between objects and times *are* instantiated at our world.

Conclusion*: Endurantism is false.

But when endurantists speak of time-indexed properties, they speak as if the world were not relativistic. Since relativity holds at our world, the endurantist must provide an account of her “time-indexed” or “time-relative” properties that is compatible with relativity. Sider (2001) shows how the endurantist picture can be coherently implemented within a relativistic framework. Assume a four-dimensional spacetime manifold, and assume that objects occupy regions of spacetime. Given this, the endurantist can relativize the instantiation of properties to points of spacetime (rather than time). For instance, suppose object *o* occupies a region of spacetime comprised of regions *r* and *s*. According to the endurantist, the object instantiates property *P* at spacetime region *r* and property *Q* at spacetime region *s*. (In contrast, the perdurantist will say that the part of the object that occupies *r* instantiates *P*, and the part that occupies *s* instantiates *Q*.)

The endurantist’s relations to times can be derived from these relations to spacetime regions. Let’s define a time *t* as an equivalence class of simultaneous points, relative to a frame of reference. Now we can say that an object *o* instantiates *P* at time *t* *iff* *o* instantiates *P* at the spacetime region in *t* occupied by *o*.⁸⁵ A relativistically acceptable account of endurantism, then, requires fundamental relations between objects and regions of spacetime.

⁸⁵ See Sider (2001, 81-82).

How are such relations different from those required by gauge theories? The underlying gauge facts involve relations between objects and spatiotemporal paths, and a spatiotemporal path is a spatiotemporal region. Gauge theories require fundamental relations between objects and spacetime regions, and endurantism requires fundamental relations between objects and spacetime regions. If one rejects the endurantist's relations one must also reject the relations invoked by gauge theories.

5.8 Two Morals

5.8.1 Endurantism vs. Perdurantism

In the debate over the nature of persistence, our intuitions about intrinsic change are often used as the ammunition. The original Argument from Temporary Intrinsic and Lewis's Having *Simpliciter* version are attempts at translating these intuitions into cogent arguments. But these arguments have limited force against endurantism, since neither perdurantism nor endurantism can accommodate our pretheoretic beliefs about ephemera. A more promising way to elucidate the intuitive objection behind the Temporary Intrinsic Argument is to shift the context of the argument to the fundamental properties. Yet this argument fails as well. It seems that no argument presented in the spirit of the original gives us a compelling reason to reject endurantism.

These aren't grounds to reject perdurantism. But it does suggest that if we want decisive reasons to prefer perdurantism to endurantism, we should look for them elsewhere.

5.8.2 Humean supervenience

In Chapter (3), we saw several reasons to reject the claim that Humean supervenience is just Locality. Gauge relations give us yet another reason to prefer a formulation of Humean supervenience along the lines of (HS₃) or (HS₄). If the fundamental relations at our world include gauge relations, then Locality is false. Since Humean supervenience entails Locality, then Humean supervenience is also false. Nonetheless, it seems that we should be able to adapt Humean supervenience in order to accommodate such advances in physics. Recall Lewis's defense of Humean supervenience when faced with evidence that fundamental physics is not local: "If I defend the philosophical tenability of Humean Supervenience, that defence can doubtless be adapted to whatever better supervenience thesis may emerge from better physics." (Lewis [1994a] 1999a, 226)

So if a better physics proposes a non-local theory, we should be able to adapt Humean supervenience to that theory. But to do so, we will need to reject Locality. If Humean supervenience is just Locality, then there is nothing left to adapt. So we have even more reason to reject the characterization of Humean supervenience in terms of Locality, and instead adopt one in terms of qualitative_O or qualitative_H.

Chapter 6:

Armstrong on Quantities and Resemblance

6.1 Introduction

A theory of universals takes at face value the idea that things share properties. Such a theory holds that universals can be instantiated by numerically distinct objects. One of the natural applications of this theory is to explain how two things resemble one another, and thus to offer an answer to the so-called Problem of Resemblance: two things intrinsically resemble one another if and only if they share some of their universals.⁸⁶

David Armstrong claims that universals provide the only tenable account of resemblance, because they provide the only *reductive* account. (Armstrong 1978 and 1989a) But whether universals provide an attractive analysis of resemblance hinges on a crucial question: can a theory of universals account for resemblance relations among properties as well as resemblance relations among objects? Armstrong believes so. He offers an account according to which the more parts two properties share, the more similar they are. (Armstrong 1988 and 1989a, 101-105)

This strategy is fatally flawed. As a result, I argue, a theory of universals cannot count an analysis of resemblance among its virtues. Since one of its alleged strengths is an elegant and reductive analysis of resemblance, the failure to produce such an account is a mark against the theory. (I will not be weighing other costs and benefits here.)

⁸⁶ In this chapter I am interested only in intrinsic resemblance, not extrinsic resemblance. For instance, I do not address cases where there is some sense in which two things resemble (perhaps each has the property of *being five feet from a poodle*), but where this resemblance does not arise from the intrinsic properties of each object alone.

In this chapter, I will look at how Armstrong's theory deals with quantitative properties, particularly those of classical mechanics. I do this for three reasons. First, Armstrong himself claims that universals are in a unique position to accommodate quantitative properties. (Armstrong 1989*a*, 101) Second, a world where the laws of classical mechanics hold is metaphysically possible, and Armstrong should be able to account for such a world. Third, if Armstrong's theory cannot accommodate the properties of classical mechanics, there is little hope it will be able to accommodate the quantitative properties of more sophisticated physical theories.

6.2 Armstrong's Picture: Resemblance as Partial Identity

On Armstrong's picture, universals are sparse; they carve nature at the joints. The paradigmatic universals are the fundamental quantities expressed by predicates in an ideal physics. (Armstrong 1988, 87) Gruesome predicates have no correlates in the world of universals. Armstrong is especially conservative with his ontology of higher-order universals, or universals instantiated by universals. (The relation of nomic necessitation is one of the few higher-order relations he allows.) Although positing higher-order universals may seem a natural way to account for property resemblance, Armstrong has several reasons to reject this approach.⁸⁷ As a result, he instead proposes a different strategy to account for resemblances among properties.

Armstrong claims that universals can be constituents of other universals, just as objects can be parts of other objects. Universals made up of constituents are *structural*

⁸⁷ One reason is the desire for ontological parsimony. Another is the fact that the natural candidates for such higher-order universals are instantiated necessarily by first-order universals, which is at odds with Armstrong's combinatorial view of possibility. See Armstrong (1978, 105-108), (1983), and (1989*b*) for more discussion.

universals, while universals with no constituents are *simple universals*. On his account, the structure of universals mirrors the structure of the objects that instantiate them. Any object that instantiates a structural universal must have proper parts which instantiate that universal's constituents. Consider a structural universal F with constituents F_1 through F_n . If object a instantiates F , then a must have numerically distinct proper parts a_1 through a_n which instantiate F_1 through F_n , respectively.⁸⁸

Two structural universals resemble one another to the extent to which they share constituents. If two universals do not share any constituents, they do not resemble one another in any respect. If they share at least one constituent, they resemble one another at least somewhat. The more constituents they share, the more similar they are: all resemblance is reduced to partial or whole identity.

Armstrong claims that every quantitative property is a structural universal. Think of quantitative universals as Russian nesting dolls. Within each doll there is a smaller doll, and a smaller one, *ad infinitum*. The largest doll "contains" all the other dolls: it shares many nested dolls with the second-largest doll, slightly fewer nested dolls with the third-largest doll, and so on. In terms of the quantity of shared dolls, the largest is more similar to the second-largest than to the third-largest. Likewise for quantitative universals – every quantitative property has an infinite number of "nested" constituent universals. Intuitively, the more constituents two quantitative universals share, the more similar they are. For example, the *five-grams-mass* universal shares many constituents

⁸⁸ If one axiomatizes the part-whole relation using a mereology that abandons the axiom of unique fusion, then Armstrong's *constituency relation* can be identified with the mereological relation of part to whole. On the other hand, if the part-whole relation is constitutively tied to the unique fusion axiom of classical mereology, then Armstrong must accept the *constituency relation* as a primitive in his ontology, albeit one that closely approximates the notion of traditional parthood. See Lewis (1986 ??) for discussion.

with the *four-grams-mass* universal, hence the property of *five-grams-mass* closely resembles the property of *four-grams-mass*.

Armstrong's scattered informal remarks can be usefully formalized by the following two principles. The first principle provides an intuitively plausible way to determine the constituents of a structural universal. Call it the *constituency principle*: a universal x is a constituent of universal y *iff* every object in every possible world that instantiates y has some proper part that instantiates x . This principle links the structure of universals to the structure of objects – universals have constituents when the objects that instantiate them have parts. The second principle provides an intuitively plausible connection between the constituency relation and the resemblance relation. Let “ $x < y$ ” mean “ y has all of the constituents of x but x does not have all of the constituents of y ,” where x and y are universals. Call this the *resemblance principle*: a is more similar to b than to c , and c is more similar to b than to a , *iff* $a < b < c$. This principle links resemblance to constituency – two properties are similar when they share constituents. Note that the resemblance principle applies only in cases where a , b , and c share at least one constituent. If they have no constituents in common, they are utterly dissimilar and so cannot be compared along any axis of similarity.⁸⁹

The *constituency principle* does a tremendous amount of work for Armstrong. First, it provides an algorithm for determining the constituents of structural universals. It explains, for example, why a charge universal is never a constituent of a mass universal – because not every massive object has a charged proper part. It also explains why a mass universal never has constituents of greater mass – because no object with mass x ever has a proper part with a mass greater than x . Second, the *constituency principle* grounds the

⁸⁹ For supporting texts, see especially Armstrong (1978, 116-131), (1988, 312-316), and (1989a, 106-107).

structure of quantitative universals in the structure of objects and their proper parts. In so doing, it provides Armstrong with a justification for constructing quantitative universals as he does. Without this principle, we must brutally posit what the constituents of quantitative universals are. Tying resemblance to constituents that are themselves ungrounded is not an improvement over positing primitive resemblance.

Thus, Armstrong analyzes resemblance in two steps. First, the *constituency principle* grounds the ontology of universals in the ontology of objects; second, the *resemblance principle* uses the ontology of universals to ground resemblance relations.

In the following sections, I present several problems with this analysis. I conclude that Armstrong's account is not a plausible theory of quantitative properties, and thus is not an improvement over accounts that posit primitive resemblance.

6.3 The Metric Function

Armstrong uses *mass* as a paradigm example to illustrate his account of quantities, so I will focus on it in the next two sections. However, my criticisms in these sections apply to any quantitative property. For now, I shall assume that mass is spread out over regions, i.e. that only occupants of regions with finite volume have mass. On Armstrong's picture, every mass universal has smaller mass universals as constituents. For instance, the *five-grams-mass* universal has the universals *four-grams-mass*, *three-grams-mass*, etc., as constituents.

Take three mass universals: *two-grams-mass*, *three-grams-mass*, and *two-thousand-grams-mass*. According to the *constituency principle*, the following is true: *two-grams-mass* shares all of its constituents with *three-grams-mass*, and *three-grams-*

mass shares all of its constituents with *two-thousand-grams-mass*, but *two-thousand-grams-mass* does not share all of its constituents with *three-grams-mass*, and *three-grams-mass* does not share all of its constituents with *two-grams-mass*. Apply the *resemblance principle*, and *two-thousand-grams-mass* is more similar to *three-grams-mass* than to *two-grams-mass*, and *two-grams-mass* is more similar to *three-grams-mass* than to *two-thousand-grams-mass*. This provides Armstrong with an intuitive resemblance ordering: *two-grams-mass* is “smaller” than *three-grams-mass*, which is “smaller” than *two-thousand-grams-mass* – therefore *two-grams-mass* is more similar to *three-grams-mass* than to *two-thousand-grams-mass*.

But this structure alone does not entail the resemblance relations that Armstrong claims it does.⁹⁰ So far, all we have is an ordering of the mass universals. Consider three mass universals, a , b , and c , such that $a < b < c$. This tells us that a is closer to b than to c (and that c is closer to b than to a), but not whether b is closer to a or to c . We know the ordering of a , b , and c ; what we do not know are the distances between these universals.

Although the resemblance principle provides nothing more than a resemblance ordering, there are natural ways to extend the principle to capture the distances between mass universals. Here is one way. Let “ $x <_d y$ ” mean “ y has all of the constituents of x and x does not have d constituents of y (where d is a number).” Now we can apply the following principle: If $a <_m b <_n c$, then b is more similar to a than to c if and only if m is less than n . Applied to universals a , b , and c , this extension appears to tell us how many constituents “away” b is from both a and c , and thus which universal it resembles more.

⁹⁰ See Armstrong (1978, 116-131).

Unfortunately, this resemblance principle does no better than the original. This is easy to see if we assume that the mass quantities are isomorphic to the real numbers; i.e. that there are an uncountably infinite number of mass quantities. Every mass universal has an infinite number of constituents, so any two mass universals will have an infinite number of constituents in common as well as an infinite number *not* in common. So this resemblance principle will not yield the result that *three-grams-mass* and *two-grams-mass* are more similar *three-grams-mass* and *two-thousand-grams-mass*.⁹¹

What we need to know is how “far apart” the mass quantities are. We can do this by using a mathematical metric. A metric is a function $d(x, y)$ defined on a set which assigns a non-negative value to each pair of elements. For all x, y , and z in the set, the metric satisfies the following properties: (1) the value it assigns to (x, y) equals the value it assigns to (y, x) , (2) it assigns 0 to (x, y) if and only if $x = y$, and (3) $d(x, z) \leq d(x, y) + d(y, z)$. (Abbott 2001, 222)

Intuitively, a metric tells us how close any two values are to one another. For instance, a mass metric function might tell us that the distance between *two-grams-mass* and *ten-grams-mass* is eight grams, and the distance between *ten-grams-mass* and *twelve-grams-mass* is two grams. The mass metric function captures both the ordering of mass quantities as well as the distances between them. Thus it gives us everything we need to

⁹¹ The same problem arises if there are countably infinite numbers of constituents. In fact, Armstrong’s account cannot guarantee the correct resemblance relations even if universals had finite numbers of constituents. Suppose there were only three mass universals: a , b , and c , and that a is simple, a is the only constituent of b , and a and b are the only constituents of c . According to Armstrong’s resemblance principle, b is equally similar to a as to c . Now suppose that a is 1 gram, b is 10 grams, and c is 11 grams. It is clear that the resemblance principle gives us the *wrong* answer: *10 grams* is not equally similar to *11 grams* and to *1 gram*. This shows us that constituency facts do not ground resemblance; rather, it is the facts about the number of grams (i.e. the facts about the metric structure of the property) that play that role. When constituency facts and metric facts come apart, Armstrong’s account delivers the wrong similarity judgments. And even if, fortuitously, they do not come apart, Armstrong’s account is fundamentally mistaken about which facts ultimately ground resemblance.

account for resemblance among mass quantities: any two quantitative universals are similar to the extent to which the distance (given to us by the metric function) between them is minimal. So, *ten-grams-mass* is more similar to *two-grams-mass* than to *one-hundred-grams mass* because the distance between *ten-grams-mass* and *two-grams-mass* is less than the distance between *ten-grams-mass* and *one-hundred-grams-mass*.

Let's step back and assess Armstrong's position. Armstrong's *constituency relation* provides an ordering of quantities, but not the distances between them. In order to capture these distances, Armstrong must enrich his account by adding structure that encodes metric information. This metric information alone provides both the ordering and the distances between quantities. Thus it provides Armstrong with everything he needs, rendering the *constituency relation* superfluous.

So where does this leave Armstrong with respect to other accounts of quantity? The boon of a theory of universals, according to Armstrong, is its ability to capture resemblance facts without the extra structure required by rival theories. We've seen, however, that both he and his rivals must incorporate metric structure in order to capture resemblance facts. So far, Armstrong's account has no advantage over any competing accounts.⁹²

6.4 Point Particles

Grant Armstrong a metric function over the mass quantities. His account runs into further problems when we abandon the assumption that mass is only instantiated by

⁹² One might follow Hartry Field (1980) in grounding metrical assignments on a fundamental level of relations of congruence and betweenness. But even then the relevant judgments of resemblance that Armstrong is interested in would be captured by the pattern of distribution of these relations and not by the facts of constituency.

objects spread out over regions. Recall the *constituency principle*, which states that a universal x is a constituent of y iff every object that instantiates y has a proper part that instantiates x . So no object without proper parts can instantiate a structural universal. Now take a single electron. An electron is point-sized – it has no radius and no apparent proper parts – yet it has a finite mass. How can this be?

Armstrong suggests that point particles like electrons *do* have proper parts – an infinite number of them. At every spacetime point where an electron is located, Armstrong claims there are an infinite number of particles. (Armstrong 1988, 315). Each of these particles is a proper part of the single electron, and each has a part of the electron's mass. The mass property instantiated by the point-sized particle now satisfies the *constituency principle*, and it appears the problem is solved.

But this solution is unsatisfactory for several reasons. First, on this account it is not metaphysically possible that just *one* object which instantiates a quantitative property like mass can occupy a point at a time. This is highly counterintuitive. Not only is it possible that one object *could* occupy a spacetime region, it seems very plausible that this is *actually* the case.

Second, there is an important disanalogy between Armstrong's treatment of quantitative properties instantiated by point particles and those instantiated by objects spread out over a region. Armstrong claims that any point particle instantiating a quantitative property has the same pattern of parts that it would have if it were spread out. If this is true, it seems we should be able to isolate the parts of point particles just as we can isolate the parts of spread out objects. But it seems we cannot. For example, on Armstrong's account particles with half of the mass of the electron must exist, but to date

physics has found no evidence of such particles. Why is this? Is it because there are heretofore unknown laws of physics preventing them from existing independently? The burden is on Armstrong to explain why these particles elude us.

Finally, we have no independent justification for Armstrong's assumption that every point particle has an uncountably infinite number of parts. Recall that Armstrong grounds the construction of universals in the mereology of objects. This tactic seems plausible in cases where we already have good reason to believe in these mereological relations, prior to the adoption of some particular metaphysical theory of properties. So, for example, it is plausible to postulate proper parts of a massy object when that object is extended over a region. This is not the case for point particles. The only reason we have for assuming the parts of point particles exist is that Armstrong's account *requires* their existence. Similar concerns apply, *mutatis mutandis*, to cases in which spacetime points instantiate fundamental field values. Just as Armstrong required point particles to have an infinite number of further point particles as parts, here he requires every spacetime point to have an infinite number of further spacetime points as parts. A theory that posits infinitely many spacetime points at each spacetime location just to salvage a rather abstruse principle about universals is just the sort of theory that gives metaphysics a bad name.

6.5 Positive and Negative Quantities

Armstrong seems to assume that his theory generalizes from the paradigm of *mass* to all fundamental quantitative properties. But this isn't the case. In this section, I will look at the fundamental property of *charge*. As with mass, Armstrong needs a metric function to

capture resemblance between charge properties. And, as with mass, problems arise when we consider finite charge values instantiated at points. So let's grant Armstrong a similarity metric over charge properties, and let's assume that *charge* is only instantiated by objects spread out over regions of spacetime.

Although Armstrong provides a general sketch for the construction of quantitative universals, he never explicitly discusses *charge*. It is safe to assume, though, that Armstrong believes *charge* universals are constructed in the same way as any other fundamental quantitative property. *Charge* properties – like *mass* – are structural universals whose constituents are other charge universals of smaller magnitudes. For example, a charge universal of +2 coulombs has constituents with values between 2 and 0 coulombs, and a charge universal of -2 coulombs has constituents with values between -2 and 0 coulombs.

An immediate problem emerges with this construction: positive and negative charge properties share no constituent universals. By the *resemblance principle*, positive and negative charges do not resemble each other at all; they have as much in common with one another as each has with *mass*. But surely this isn't right. Positive charge has something very important in common with negative charge – they are both charge!

How can Armstrong account for the apparent resemblance between positive and negative charge? Here is one option. Even though positive and negative charge do not share constituent universals, perhaps their similarity lies in the role they play in the laws of nature. So any resemblance between positive and negative charge consists solely in how they are treated by the laws. But this is resemblance in virtue of causal role, not in virtue of shared universals. Such a solution undercuts the motivation behind Armstrong's

theory of structural universals, since the point of the project is to reduce resemblance to partial identity.

Here is another option. Suppose we assume that properties of positive and negative charge *do* share constituent universals. In that case, every charge universal has constituents in common with every other charge universal – both positive *and* negative. Resemblance between positive and negative charge naturally follows.

While this option is attractive in some regards, it fails on several counts. First, it leads to odd constraints on metaphysical possibility. Recall that the *constituency principle* states an object instantiating a universal must have parts instantiating the constituents of the universal. On this version of Armstrong's account, then, any object that instantiates positive (or negative) charge must have positively *and* negatively charged parts. This rules out common cases in classical mechanics according to which there are such things as spheres of uniform positive charge. On this account such a sphere must have a part that instantiates negative charge – but as described, the sphere has no parts that instantiate negative charge. It turns out that many common classical mechanical cases are metaphysically impossible.

Second, problems arise concerning the construction of charge universals. What constituents does, say, the *+3-coulombs* universal have? Following Armstrong's construction of *mass* universals, it is natural to take *charge* universals to have constituents of smaller charge magnitudes. So *+3-coulombs* has a constituent for every charge value between +3 and -3.⁹³

⁹³ This is a natural move, I believe, but notice that we must simply assume that charged objects always have these parts. Already the theory is beginning to look uncomfortably *ad hoc*.

But now we cannot deliver resemblance relations via the *resemblance principle*. Charge is an additive property: if we partition a charged object (divide it into mutually exclusive and exhaustive parts) the charge of these parts will add up to the charge of the whole. Positive and negative values will cancel each other out, so that an object with a -10-coulombs part and a +10-coulombs part has zero net charge.

Take an object of +10 coulombs. We've assumed that such an object has a part for every charge value between +10 and -10 coulombs. So let's pick out a part of the object that has -9 coulombs of charge; call this part *a* and the remainder part *b*. In order to preserve the charge value of the whole (+10 coulombs), part *b* must have +19 coulombs of charge. Now, we know that *b* likewise has a part with every charge value between +19 and -19 coulombs. So let's pick out a part of *b* that has -18 coulombs of charge; call this part *c* and the remainder part *d*. Again, in order to preserve the charge value, the part *d* must have a charge of +37 coulombs. And so on. Eventually, we find that every charged object has proper parts that instantiate arbitrarily large and arbitrarily small positive and negative charge magnitudes. For any charge magnitude *c*, every charged object will have a part that instantiates *c*.

By the *constituency principle*, every charged object has parts that instantiate every possible charge value. So every charge universal will have constituents of every possible charge value. Thus, *all* charge universals have the same constituents, and so they all have the same constituents in common. But then, what grounds resemblance? Not the *resemblance principle* (in conjunction with the similarity metric), which tells us that the more constituents two universals share, the more they resemble each other. So Armstrong is still left without an account of resemblance.

(This account runs into further trouble if Armstrong takes two universals with the very same constituents to be numerically identical. By definition, qualitatively identical *simple* universals are numerically identical, but it is unclear how Armstrong deals with qualitatively identical *structural* universals. If universals with the same constituents are numerically identical, then every *charge* universal is numerically identical to every other *charge* universal – so a *+2-coulomb* universal would be identical to a *+100-coulomb* universal, an undesirable result.)

This extension of Armstrong's account fails. Perhaps other extensions will be contrived. But it's hard to see how any account will be able to capture all of the resemblance facts in a manner true to Armstrong's initial ambitions. In any case, I will not explore this further. One of the goals of a reductive project is to purge our ontology of dubious and otherwise unjustified entities. A reduction of resemblance that generates a dubious and otherwise unjustified ontology is not a reduction worth pursuing.

6.6 Vector Quantities

A vector is a quantity characterized by a magnitude and a direction. Any account of resemblance between properties needs to account for resemblance between vector properties.

Armstrong discusses vector properties in Armstrong (1988). He begins by allowing himself an ontology of fundamental properties and relations instantiated by spacetime points and intervals, but he does not allow any fundamental *quantitative* relations. He then claims that all vector properties can be reduced to these fundamental properties and relations. He backs up this claims with two examples: *being x amount*

later than and being x amount east of. Both of these vector quantities may be reduced to a magnitude and spatiotemporal distance relation. (Armstrong 1988, 310-311)

But there are vector properties which cannot be reduced in this way. An *electric field* vector, for example, cannot be decomposed into electric field magnitudes and spatiotemporal interval properties. The pattern of electric field “magnitudes” and the spatiotemporal intervals between them leaves out crucial information given by the electric field vector. In particular, it will not tell us what direction each magnitude “points” in. Consider a polarized ray of light – the locations and values of the ray’s electric field magnitudes alone do not give the direction of polarization. So it is not the case that all vector properties can be reduced in the way Armstrong suggests.

Armstrong must enrich his fundamental ontology in some way if he wishes to account for vector properties. Here are three natural ways Armstrong might do this. First, he might be more liberal in the fundamental relations that he allows into his ontology. Second, he might decompose vectors into two fundamental properties. Third, he might simply allow fundamental vector properties into his ontology. Each of these strategies has its pitfalls, but the one they all share is a failure to account for resemblance.

Let’s see how the first strategy of allowing fundamental relations works for electric field vectors. One way to pursue this strategy is to add two new fundamental elements to one’s ontology: an *electric field magnitude* property and an *electric field points-at* relation. The *electric field magnitude* is the value associated with the electric field vector, and the *electric field points-at* relation is an asymmetric relation that holds between pairs of spacetime points that encodes the vector’s direction. A 5 volts/meter electric field vector pointing north is reduced to a 5 v/m *electric field magnitude* and the

fact that the *electric field points-at* relation holds between the spacetime point instantiating that magnitude and each spacetime point due north of it.

This version is unsatisfactory for several reasons. One of Armstrong's goals in providing a reductive account of vectors is to reduce the number of strange and unnecessary primitives in his ontology. But this version does the opposite – it increases both the number and the types of primitives required, and these primitives are arguably more exotic than the vector properties being reduced. Moreover, this account results in bizarre metaphysical possibilities. Armstrong believes in a combinatorial theory of possibility, according to which the fundamental properties and relations of our world can be recombined to produce a possible world. (Armstrong 1989*b*) If *electric field magnitudes* and *electric field points-at* relations are fundamental, then it is metaphysically possible that one could be instantiated without the other. In other words, two points could stand in the *electric field points-at* relation even though an *electric field magnitude* is not associated with either point.

A second way Armstrong might try to analyze vectors is by invoking pairs of monadic properties. For example, he might postulate a *magnitude* property that encodes the magnitude of the vector, and a *direction* property that encodes its direction. There are an infinite number of *direction* properties, one for each direction in which a vector can point. This strategy has the same demerits as the last one. It increases the number and types of primitives in Armstrong's ontology, and his combinatorialism entails that an *electric field direction* property may be instantiated at a point with no associated *electric field magnitude*, a strange consequence.

More importantly, neither of these accounts can fully capture resemblance between vector properties. Different directions resemble each other to different degrees. Consider three electric field vectors, a , b , and c , of the same magnitude, with the same origin, and lying in the same plane. Vector a points north, b points 10° in a clockwise direction from a , and c points 10° clockwise from b . It seems that a is more similar to b than to c because a and b are more similar with regard to direction.

Any account of resemblance must explain this, but an analysis of vectors in terms of either *magnitudes* and *direction* properties or *magnitudes* and *points-at* relations cannot. On Armstrong's account only quantitative fundamental universals have constituents, and neither *direction* properties nor *points-at* relations are quantitative. No *direction* universals have any constituents in common, so no *direction* universal ever resembles another *direction* universal. Likewise, no *points-at* relations have any constituents in common, so no *points-at* relation ever resembles another *points-at* relation. Thus, we cannot say that vector a resembles b more than c because a and b point in more similar directions. Neither strategy can capture all of the resemblance relations between vectors, since the sharing of constituents cannot explain any similarity between directions.

A third way Armstrong might try to account for vectors is by simply admitting them as primitives into his ontology. Yet problems arise here as well. Suppose every vector universal has constituents. Since vectors are instantiated at points, this leads to the kinds of problems discussed in section 3, those that arise when finite magnitudes are instantiated at points. But unlike *mass* and *charge*, there is no story to tell about vector properties according to which they are extensive or spread out over a region. So in this

case we do not even have a spread-out template that we can use to discern the structure of vector universals.

Resemblance is even more of a problem on this account than on the previous two. By the *resemblance principle*, two vector universals resemble each other when they have constituents in common. Suppose that a vector's constituents are vectors with smaller magnitudes that all point in the same direction. Then Armstrong can account for the similarity of vector universals that point in the same direction.

But he cannot account for the similarity of vector universals that point in different directions. (Any attempt to construct a vector property with constituents that point in different directions will run into the same troubles as we saw when we tried to construct charge properties with positive and negative constituents (see section 4).) Consider two unit electric field vectors 10° apart that have the same origin and lie in the same plane. On this account, they do not resemble each *at all*, since they share no constituents. But it seems they have a great deal in common; in fact, they exactly resemble in several respects. First, they are both vectors, and two vector properties are more similar than a vector and a scalar property. Second, they both electric field vectors, and two electric field vector are more similar than an electric field vector and a magnetic field vector. Third, they have the same magnitude, and two vectors of the same magnitude pointing in different directions are more similar than two vectors of *different* magnitudes pointing in different directions. Yet on this account these two electric field vectors resemble each other no more than each resembles a *mass* universal.

I shall not speculate about how Armstrong would revise his theory of vectors when faced with the decisive problems with the account he actually states. I have no idea

which package of costs would seem more palatable and, in particular, to what extent he would be willing to accept additional fundamental relations, to what extent he would be willing to deny the authenticity of various apparent resemblance facts, and to what extent he would be willing to bite various bullets on combinatorial possibility. What is inevitable is a breakdown of the original picture that sought to explain all intrinsic resemblance by the sharing of constituents. At best, Armstrong cannot capture all the resemblances between vector properties, and at worst he cannot accommodate vector properties at all.

6.7 Resemblance in Different Respects

It is no surprise that our resemblance judgments depend significantly on context. Armstrong lauds universals for readily explaining this phenomenon when it comes to comparing similar objects: “If resemblance is a matter of different identities in different cases, it is easy to see that degrees of resemblance will be a partially subjective matter, depending upon what particular properties we happen to be interested in, in the particular context.” (Armstrong 1998, 103) But while Armstrong can accommodate context-sensitivity when it comes to objects, he cannot do the same for properties.

Consider three *charge* universals: *+5 coulombs*, *-1 coulomb*, and *-5 coulombs*.

One way to judge similarity between them is in terms of difference in charge value – the *-5 coulombs* universal is more similar to the *-1 coulomb* universal than it is to the *+5 coulombs* universal because *-1* and *-5* are only four units of charge “apart”, while *-5* and *+5* are ten units “apart”. But there is another way to judge similarity. We could say instead that the *-5 coulombs* universal is more similar to the *+5 coulombs* universal than

the *-1 coulomb* universal because -5 and $+5$ are the same magnitude of charge. Neither judgment of resemblance is prohibited; rather, the first seems right when similarity is judged in one context, the second seems right when similarity is judged in another.

The same can be said for vector quantities. Consider three unit *electric field* vectors lying on the same plane with a common origin: one points north, one points west, and one points south. Again, which are more similar? On the one hand, the vector pointing north is more similar to the one pointing west because the angle between them is smaller. On the other, the vector pointing north and the vector pointing south are more similar because they are parallel, while the vector pointing west is perpendicular.

This generates a challenge for Armstrong. We have seen how he wishes to account for the context-dependence of resemblance judgments in the case of objects – we restrict our attention to a subset of the universals in play. And we have seen that there is certainly context dependence of resemblance judgment in the case of magnitudes. For his account of resemblance to have desirable unity, he would need to explain context dependence in the same way: in terms of attention being restricted to a subset of the universals in play. But in the cases described above, there seem to be more dimensions of context-dependence to judgments of resemblance than there are families of universals present. Some kind of breakdown in the account once again seems imminent.

6.8 “The Fate of the Universals”

Why universals?

For Armstrong, a major selling point of his theory of properties is its power to analyze resemblance. Universals can account for our resemblance intuitions while

providing a reduction of resemblance, he says. No other theory of properties can do the same.

Armstrong offers a theory of universals that is intended to provide a unified account of resemblance for both objects and properties. Objects instantiate universals, and the more universals objects share the more similar they are. Universals are constructed out of constituents, and the more constituents universals share the more similar they are. The construction of universals is linked to the construction of objects: only objects with the appropriate proper parts can instantiate structural universals.

We've seen that this construction of structural universals does not deliver an adequate account of resemblance among properties. First, appeal to the "number of" shared constituents is no substitute for a similarity metric. Second, the applications to point-sized particles are bizarre and *ad hoc*. Third, the account faces special troubles from positive and negative quantities. Fourth, it faces worse troubles from vector quantities. Finally, the theory provides no basis for a unified account of context-sensitivity for resemblance.

Armstrong writes: "The fate of the Universals theory may turn on the question of the inexact resemblance of universals." (Armstrong 1988, 139)

I concur.

Bibliography

- Abbott, Stephen (2001). *Understanding Analysis*. New York: Springer-Verlag Press.
- Armstrong, D.M. (1978). *Universals and Scientific Realism, Volume II: A Theory of Universals*. Cambridge University Press.
- Armstrong, D.M. (1983). *What is a Law of Nature?* Cambridge University Press.
- Armstrong, D.M. (1988). "Are Quantities Relations? A Reply to Bigelow and Pargetter." *Philosophical Studies* 54: 305-316.
- Armstrong, D.M. (1989a). *Universals: An Opinionated Introduction*. Westview Press.
- Armstrong, D.M. (1989b). *A Combinatorial Theory of Possibility*. Cambridge University Press.
- Arntzenius, Frank (2003). *Something something*. Manuscript.
- Bigelow, John, John Collines & Robert Pargetter (1993). "The Big Bad Bug: What are the Humean's Chances?" *British Journal of the Philosophy of Science* 44: 443-462.
- Bricker, Phillip (2001). "Island Universes and the Analysis of Modality." In Pregyer & Siebelt, eds. *Reality and Humean Supervenience: Essays on the Philosophy of David Lewis*. Rowman & Littlefield.
- Bricker, Phillip (2007). "Concrete Possible Worlds." In Hawthorne, Sider, & Zimmerman, eds., *Contemporary Debates in Metaphysics*. Blackwell.
- Caplan, Ben (2003). "Why So Tense about the Copula?" *Mind* 114: 703-708.
- Carroll, John (1994). *Laws of Nature*. Cambridge University Press.
- Chalmers, David J. (2002). "Does Conceivability Entail Possibility?" In Gendler & Hawthorne, eds. *Conceivability and Possibility*. Oxford University Press.
- Dunn, J. Michael (1990). "Relevant Predication 2: Intrinsic Properties and Internal Relations." *Philosophical Studies* 60: 177-206.
- Earman, John (1986). *A Primer on Determinism*. Dordrecht: Reidel.
- Earman, John and John T. Roberts (2005). "Contact with the Nomic: A Challenge for Deniers of Humean Supervenience about Laws of Nature. Part I: Humean Supervenience." *Philosophy and Phenomenological Research* 71: 1-22.
- Ellis, Brian (2001). *Scientific Essentialism*. Cambridge University Press.
- Field, Hartry (1973). "Theory Change and the Indeterminacy of Reference." *The Journal of Philosophy* 70: 462-481.
- Field, Hartry (1980). *Science Without Numbers*. Princeton University Press.
- Fine, Kit (1994). "Essence and Modality." *Philosophical Perspectives* 8: 1-16.
- Hall, Ned (2004). "Two Concepts of Causation." In Collins, Hall, & Paul, eds., *Causation and Counterfactuals*. MIT Press.

- Haslanger, Sally (1989). "Endurance and Temporary Intrinsic." *Analysis* 49: 119-125.
- Hawthorne, John (2006). *Metaphysical Essays*. Oxford University Press.
- Hawthorne, John, Ryan Wasserman and Mark Scala (2004). "Recombination, Causal Constraints, and Humean Supervenience: An Argument for Temporal Parts?" In Zimmerman, ed., *Oxford Studies in Metaphysics, Vol. 1*. Oxford; reprinted in Hawthorne 2006: 71-84.
- Hinchliff, Mark (1996). "The Puzzle of Change." *Philosophical Perspectives* 10: 119-136.
- Jackson, Frank (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.
- Kim, Jaegwon (1982). "Psychophysical Supervenience." *Philosophical Studies* 41: 51-70.
- Kim, Jaegwon (2005). *Physicalism, or Something Near Enough*. Princeton University Press.
- Langton, Rae and David Lewis (1998). "Defining 'Intrinsic'." *Philosophy and Phenomenological Research* 58: 333-345; reprinted in Lewis 1999a: 116-132.
- Lewis, David (1973). "Causation." *Journal of Philosophy* 70: 556-567; reprinted with postscripts in Lewis 1986c: 159-213.
- Lewis, David (1983a). "Extrinsic Properties." *Philosophical Studies* 44: 197-200; reprinted in Lewis 1999a: 111-115.
- Lewis, David (1983b). "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 434-377; reprinted in Lewis 1999a: 8-55.
- Lewis, David (1984). "Putnam's Paradox." *Australasian Journal of Philosophy* 62: 221-246; reprinted in Lewis 1999a: 56-77.
- Lewis, David (1986a). "Against Structural Universals." *Australasian Journal of Philosophy* 64: 25-46; reprinted in Lewis 1999a: 78-107.
- Lewis, David (1986b). *On the Plurality of Worlds*. Blackwell.
- Lewis, David (1986c). *Philosophical Papers, Volume II*. Oxford University Press.
- Lewis, David (1988). "Rearrangement of Particles: Reply to Lowe." *Analysis* 48: 65-72; reprinted in Lewis 1999a: 187-195.
- Lewis, David (1989). "Dispositional Theories of Value." *The Proceedings of the Aristotelian Society* 63: 113-137; reprinted in Lewis 2000: 68-94.
- Lewis, David (1994a). "Humean Supervenience Debugged." *Mind* 103: 473-490; reprinted in Lewis 1999a: 224-247.
- Lewis, David (1994b). "Reduction of Mind." In Guttenplan, ed., *Companion to Philosophy of Mind*. Blackwell; reprinted in Lewis 1999a: 291-324.
- Lewis, David (1999a). *Papers in Metaphysics and Epistemology*. Cambridge University Press.

- Lewis, David (1999b). "Zimmerman and the Spinning Sphere." *Australasian Journal of Philosophy* 77: 209-212.
- Lewis, David (2000). *Papers in Ethics and Social Philosophy*. Cambridge University Press.
- Lewis, David (2002). "Tensing the Copula." *Mind* 111: 1-13.
- Loewer, Barry (1996). "Humean Supervenience." *Philosophical Topics* 24: 101-126.
- Loewer, Barry (2004). "David Lewis's Humean Theory of Objective Chance." *Philosophy of Science* 71: 1115-1125.
- Maudlin, Tim (2002). *Quantum Non-Locality and Relativity*. Oxford: Basil Blackwell.
- Maudlin, Tim (2007). *The Metaphysics Within Physics*. Oxford University Press.
- McDaniel, Kris (2007). "Brutal Simples." In Zimmerman, ed. *Oxford Studies in Metaphysics: Volume 3*. Oxford University Press.
- Merricks, Trenton (1994). "Endurance and Indiscernibility." *The Journal of Philosophy* 91: 165-184.
- Schaffer, Jonathan (2003). "Principled Chances." *British Journal of Philosophy of Science* 54: 27-41.
- Schaffer, Jonathan (2009). "On What Grounds What." In Chalmers, Manley, and Wasserman, eds. *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford University Press.
- Shoemaker, Sydney (1980). "Causality and Properties." In van Inwagen, ed. *Time and Cause: Essays Presented to Richard Taylor*. Dordrecht: Reidel; pg. 109-35.
- Sider, Theodore (1995). "Sparseness, Immanence, and Naturalness." *Nous* 29: 360-377.
- Sider, Theodore (1996). "Intrinsic Properties." *Philosophical Studies* 83: 1-27.
- Sider, Theodore (1999). "Global Supervenience and Identity Across Times and Worlds." *Philosophy and Phenomenological Research* 59: 913-927.
- Sider, Theodore (2001). *Four-Dimensionalism*. Oxford University Press.
- Skow, Bradford (2007). "Are Shapes Intrinsic?" *Philosophical Studies* 133: 111-130.
- Teller, Paul (1985). "A Poor Man's Guide to Supervenience and Determination." *Southern Journal of Philosophy* 22: 137-67; reprinted in Tooley 1999: 9-24.
- Tooley, Michael (1987). *Causation: A Realist Approach*. Clarendon Press.
- Tooley, Michael, ed. (1999). *Laws of Nature, Causation, and Supervenience*. Garland Publishing.
- Trogdon, Kelly (2009). "Monism and Intrinsicity." *Australasian Journal of Philosophy* 87: 127-148.
- Van Inwagen, Peter (1990). "Four-Dimensional Objects." *Nous* 24: 245-255.
- Wasserman, Ryan (2003). "The Argument from Temporary Intrinsic." *Australasian Journal of Philosophy* 81: 413-419.

Curriculum Vita

Maya Eddon

Education

1998-2002 A.B. in Philosophy, Princeton University
2002-2009 Ph.D, in Philosophy, Rutgers University

Employment

2007-2010 Post-Doctoral Fellow in Philosophy, University of Massachussetts,
Amherst