Naturalness, Intrinsicality, and Duplication

Theodore Sider

Ph.D. thesis, University of Massachusetts.
Committee: Phillip Bricker (chair), Fred
Feldman, Edmund Gettier, Angelika Kratzer

May 1993

Acknowledgments

I would like to thank various people for their verbal and written comments on earlier stages of this dissertation. Of course, none of these people are responsible for the views I express. I am indebted to Mark Aronszajn, Lynne Baker, David Braun, Earl Conee, David Cowles, Max Cresswell, David Denby, Fred Feldman, Richard Feldman, Edmund Gettier, Larry Holm, Ned Markosian, Deborah Modrak, and R. Cranston Paull.

I am especially indebted to my dissertation director, Phillip Bricker. He consistently gave me excellent comments and helped me to see more mistakes than I care to remember. For whatever there is of value in this dissertation, he deserves much of the credit.

Abstract

This dissertation explores the concepts of naturalness, intrinsicality, and duplication. An intrinsic property is had by an object purely in virtue of the way that object is considered in itself. Duplicate objects are exactly similar, considered as they are in themselves. The perfectly natural properties are the most fundamental properties of the world, upon which the nature of the world depends.

In this dissertation I develop a theory of intrinsicality, naturalness, and duplication and explore their philosophical applications. Chapter 1 introduces the notions, gives a preliminary survey of some proposed conceptual connections between the notions, and sketches some of their proposed applications. Chapter 2 gives my background assumptions and introduces notational conventions.

In chapter 3 I present a theory of naturalness. Although I take 'natural' as a primitive, I clarify this notion by distinguishing and explicating various conceptions of naturalness.

In chapter 4 I give a theory of various notions related to naturalness, especially intrinsicality and duplication. I show that 'intrinsic' and 'duplicate' are interdefinable, and then give analyses of these and other notions in terms of naturalness.

If, as I think likely, naturalness cannot be analyzed, then what is the proper response? David Lewis suggests: accept naturalness as a primitive. I am sympathetic to this proposal, but not to the form Lewis gives it: chapter 5 contains an argument against Lewis's theory of naturalness.

In chapter 6 I reject the idea that naturalness is analyzable in terms of *immanent universals*. I focus on the work of D. M. Armstrong. I also criticize Armstrong's arguments against transcendent universals.

In chapter 7 I address criticisms of David Lewis's definition of 'intrinsic' offered by Mike Dunn.

In chapter 8 I discuss the possibility of analyzing our three notions. I discuss defining 'natural' in terms of supervenience and other concepts, and then criticize attempts by Jaegwon Kim and Michael Slote to analyze intrinsicality in terms of "quasi-logical" concepts.

Finally, in chapter 9 I present a new application for the notion of naturalness: the statement of "metrical realism" in the philosophy of space and time.

Contents

List of Figures

Chapter 1

Introducing Three Notions

1.1 Intrinsicality and Duplication

Today my brother got a haircut. Yesterday Mike had the property *having long hair*; today he does not. By losing this property my brother changed.

When my only brother got his much-needed haircut, I too lost a property: the property *having a long-haired brother*. But when the haircut occurred, I did *not* change. Mike's haircut changed Mike, not me, despite the fact that we each thereby lost properties we formerly had (and gained new ones as well). This is because the property Mike lost, *having long hair*, is an **intrinsic** property, whereas the property I lost, *having a long-haired brother*, is not.

Consider the property *having long hair.* When someone has this property, he has it in virtue of the way he is in himself. In contrast, when a person has the property *having a long-haired brother*, he does not have this purely in virtue of the way he is himself; he has this property partly in virtue of the way someone else is (namely, his brother) and partly in virtue of the relations he bears to that someone else (namely, the *being the only brother of* relation). When an object x has an intrinsic property, then, this is in virtue of the way x is, and not a matter of the ways other objects are, nor of the way x stands with respect to other objects. Intrinsic properties are "non-relational" and "purely qualitative". If an object has an intrinsic property, then so must any "perfect duplicate" of that object. In contrast, **extrinsic** (non-intrinsic) properties may differ between perfect duplicates. Marble A may be a perfect duplicate of marble B despite the fact that A has, while B lacks, the property *being ten feet from Ted.* Shapes, sizes, masses—these are intrinsic properties.

Locations, speeds, ownerships (e.g. *being owned by Ted*)—these are extrinsic properties.

I hope that the preceding paragraph has helped you grasp the notion of an intrinsic property, if you were not already familiar with that notion. But I suspect that you distrust the value of what I have said as an *analysis* of the notion, on grounds of circularity. You may ask me to explain the locutions I used in that paragraph: 'duplicate', and 'had in virtue of the way an object is in itself', for example. What is the "way an object is in itself", if not the conjunction of its intrinsic properties? What is it for two objects to be "duplicates", if not for them to have exactly the same intrinsic properties?

Let us focus a little more closely on the notion of **duplication**, and its relation to intrinsicality. Objects that are perfect duplicates are exactly similar, down to the last detail of the smallest part. It is likely that there are no *actual* pairs of *macroscopic* duplicates. Even two marbles made by the best manufacturing techniques are bound to differ slightly—a stray atom here or there. The differences may not be perceptible, but if they are there then the marbles cannot be duplicates, for duplicates may not have *any* (intrinsic) differences whatsoever. Of course, *no* pair of objects, even a pair of perfect duplicates, will share *all* properties. For consider any two objects *a* and *b*—only *a* will have the property *being a*. Although it is likely that there are no two actual macroscopic duplicates, it may be that there are pairs of actual *micro*scopic duplicates: perhaps pairs of duplicate electrons. And surely every object, microscopic or macroscopic, has a duplicate in some other possible world.[1]

One reason for the importance of the concept of duplication is that 'intrinsic' may be defined in terms of 'duplicate', as David Lewis has proposed:[2]

(D1) Property P is *intrinsic* iff for any possible individuals x and y, if x and y are duplicates then x has P iff y has P

In fact, the direction of definition can be reversed: 'duplicate' is definable in terms of 'intrinsic':

(D2) Possible individuals x and y are *duplicates* iff for any intrinsic property P, x has P iff y has P

[1] I ignore sets, numbers, and the like. See section 2.1.

[2] See Lewis (1986*c*, p. 62), and Lewis (1983*a*, p. 197).

We have a circle of interdefinability between 'intrinsic' and 'duplicate'.[3] Some philosophers claim to be able to break into this circle and define 'intrinsic' merely in terms of "quasi-logical" (e.g. modal, part/whole) and spatiotemporal concepts. I discuss this claim in chapter 8; my verdict is negative. Some might claim that intrinsicality is relative to human interests or cognitive capabilities. A related proposal would be that certain properties are singled out by context as those we "count as intrinsic" for the purposes of conversation. But on these views there would be no "objective" intrinsic/extrinsic distinction; there would merely be our conventions. I reject this option; I will mention below the reason for this rejection. What, then, is the remaining option? To take the distinction (or a related distinction—see below) as a primitive distinction.

Some, I know, will find this ludicrous. When I discuss intrinsicality and duplication they will refuse to understand my words; for them, my arguments will be impossible to evaluate, my claims ultimately groundless. Until I can define my terms they will reject them. But I say that the notions of intrinsicality and duplication are pre-analytically understood, and are therefore as fit a foundation for philosophical theorizing as we can reasonably demand. Go back and re-read the example of my long-haired brother. Can you honestly claim not to see the difference between *having long hair* and *having a long-haired brother*?

Of course, it would be rash to suppose that understanding of the *word* 'intrinsic' in precisely the sense that I intend is a part of everyday wisdom, for it is not. Many have never thought of the concept of intrinsicality, and it often requires a little "coaching" to introduce the notion to someone who is not familiar with it. Moreover, some philosophers may use 'intrinsic' to mean something different from what I mean by that word, so I should not claim that intrinsicality is commonly known by that name. It may be that 'intrinsic' has several meanings.[4] But I do claim that there is a notion that is properly called 'intrinsicality' that is easily fixed on, given a few phrases (e.g. "purely qualitative", "non- relational", etc.) and suggestive examples.

In particular, I regard definition (D1) as doing the lion's share of the work in picking out the intended sense of 'intrinsic', for the notion of duplication is, I think, clear and unambiguous. Consider the claim of Mike Dunn, for

---

[3]In section 4.1 I show that 'intrinsic' and 'duplicate' are indeed interdefinable via (D1) and (D2).

[4]See Moore (1922), esp. p. 262.

example, that the property *being a* is intrinsic, for it "depends on *a* and on no other thing" (1990, p. 186). Dunn, I think, has simply got a different notion of intrinsicality in mind, and we can clear things up by stipulating that the notion of intrinsicality we are here concerned with may be defined in terms of duplication via (D1). Additionally, I exclude this property because it is not *purely qualitative*.[5]

To forestall certain objections to my claim that the notions of intrinsicality and duplication are well-understood, consider an example. Let us grant ourselves talk of the colors of visual "sensations" or "qualia". Thus, red things cause me to experience "red" qualia when I look at them in the appropriate light. Some philosophers (notably Locke) thought that color properties are somehow not as real as other properties of physical objects—they are "secondary qualities". One could interpret this as involving a claim that color properties are not intrinsic. For suppose, simplistically, that Locke's view is that x is red iff x causes "red" qualia in people. According to this view, *redness* is clearly extrinsic. For let x be any actual red object; we can imagine a world containing a perfect duplicate of x that causes no "red" qualia (perhaps that world contains no people, or perhaps the people there have "inverted spectra"). Since x and a perfect duplicate of x differ over *redness*, that property turns out extrinsic on Locke's view.

Another theory according to which color properties are extrinsic says that the color of an object is determined by the wavelength of the light it scatters. Since the wavelength of light scattered by an object is partly a matter of the laws of nature, colors on this view are extrinsic as well. For consider some actual red ball B. We can imagine a possible world with different laws governing the interaction of light with material bodies, a result of which is that a certain duplicate of B scatters light with the same wavelength as that scattered by blue objects in the actual world. The theory we are considering has the result that this duplicate of B has a different color from B: blue.

Others claim that color properties are intrinsic. One such theory holds that the color of an object is determined by the physical and microphysical properties of its surface. Since the latter properties are intrinsic, this theory entails that color properties are likewise intrinsic.

One might think that the fact that it is controversial whether color properties are intrinsic falsifies my claim that the intrinsic/extrinsic distinction is pre-analytically understood. Not so. All agree on the intrinsic/extrinsic

[5]See section 4.2.2 for a definition and explanation of 'qualitative'.

status of the relevant properties here. *Having surface structure S* is intrinsic. *Reflecting light of wavelength λ* is extrinsic. *Causing "red" qualia* is extrinsic. No disagreement there. The disagreement is over what property is expressed by the word 'red', and not over the intrinsicality of any one property. Similarly, I might know each person in the room, and still not able to answer when you ask: who is the richest person in the room?

Similar considerations arise for the case of words referring to "artifacts". Consider a chair made of homogeneous matter, and also a slab of solid homogeneous matter of the same kind. If the slab is large enough, then, given certain assumptions, there is a duplicate of the chair inside the slab.[6] Indeed, there will be many. Let us call these objects inside the slab "pseudo-chairs". Are pseudo-chairs *chairs*? I do not know. Since I do not know, I do not know whether *being a chair* is intrinsic. If pseudo-chairs are not chairs then the chair has non-chairs for duplicates, and hence *being a chair* is extrinsic. On the other hand, if pseudo-chairs are chairs, then perhaps *being a chair* is intrinsic.[7] But again, my uncertainty is not uncertainty over the extension of 'intrinsic'. I am uncertain because I do not know precisely which property is expressed by 'chair'.

Let us consider another example to clarify the concepts of intrinsicality and duplication. Consider a ball made of sponge into which I am pressing my finger. That ball has a shape that we might call "S": somewhat spherical, but deformed at the place where my finger presses in.

Someone might say: "*Having shape S* is not an intrinsic property. Surely everything is a duplicate of itself. But before I pressed my finger into the

[6]The assumptions are two. The first is the assumption that for any objects whatsoever, there is an object that is their mereological sum. Peter van Inwagen, for example, would deny this. See van Inwagen (1990), esp. pp. 74–80. The second is that the chair-shaped objects inside the slab are not disqualified from being duplicates of the original chair by the fact that they are surrounded by matter. The issue here involves "local properties". Under the rubric of "local properties" go not only properties had by an object in virtue of what that object is like, but also properties had by an object in virtue of what goes on in its *infinitesimal neighborhood*. Must all local properties be intrinsic? If all are, then a chair cannot have a duplicate that is surrounded by matter. But I assume that not all local properties are intrinsic—thus, chair-like entities inside the slab are duplicates of the original chair.

[7]Other cases, such as one involving a duplicate of the chair that happens to materialize in outer space, might be relevant as well. I am inclined to think that pseudo-chairs are not chairs, and hence that chairhood is extrinsic. But there may be no determinate fact of the matter.

ball, it had a different shape. Hence, *having shape S* can differ between duplicates—namely, between the ball and itself—and so is extrinsic." This would clearly be a mistake. To fix it, let us speak of the various "temporal parts", or "time slices" of the ball. Let s_1 be a time slice of the ball when I am pressing on it, and let s_2 be a slice of the ball when I am not pressing on it. s_1 has shape S, whereas s_2 does not. But s_1 and s_2 are not duplicates. Certainly, everything is a duplicate of itself, but this does not imply that s_1 is a duplicate of s_2, for s_1 and s_2 are distinct.

Someone might say: "The property of *being pushed on* is intrinsic. Any duplicate of the ball would have to have something pushing on the spot of the deformation to give it shape S." This too would be a mistake. To see that this property is extrinsic, consider a possible world w containing only a ball with shape S. Since the ball is alone in w, nothing is pushing on it, but we may stipulate that this ball is a duplicate of the actual ball. It may be objected that the ball would not have shape S, for the forces between the molecules constituting the ball would force it into a circular shape in the absence of a force deforming it. But this objection assumes that the causal laws of w are the same as those of our world, and there is no reason to suppose this. It surely is metaphysically *possible* that a ball have shape S without any force being exerted on it at all. This example reminds us that in definitions (D1) and (D2), we are told to consider *all* possible objects, and not just objects in worlds with the same causal laws as the actual world.[8]

In intuitively conveying the concept of intrinsicality, I have relied heavily on David Lewis's words in Lewis (1986*c*, pp. 61 ff.), and in Lewis (1983*a*, p. 197). (For a radically different conception of intrinsicality, see Shoemaker (1980).) G. E. Moore has an account similar to that of Lewis. Moore characterizes the notion of an "intrinsic kind of value" (1922, p. 260).

> To say that a kind of value is "intrinsic" means merely that the question whether a thing possesses it, and in what degree it possesses it, depends solely on the intrinsic nature of the thing in question.

He then says that x's having P "depends solely on the intrinsic nature" of x iff in any circumstances, an object that is either i) identical to x, or ii) "exactly

[8] We needn't have considered a world with different laws. Suppose I squeeze the ball even harder until it reaches an even more deformed shape S'. Now I let it go. While returning back to sphericality, we may suppose that the ball momentarily assumes shape S. Its stage at that moment has shape S, and may be stipulated to be a duplicate of the original stage of the ball, but it has nothing pressing on it.

like x" must of necessity have P iff x has P (1922, pp. 260–261). The first clause, that any object identical to x must have P iff x has P, may seem odd. But it is clear that for Moore, any object not exactly like x would necessarily be distinct from x (1922, p. 261). Thus, it seems not unfair to attribute to Moore the view that a "value" is intrinsic just when it never differs between objects that are exactly alike—i.e. are duplicates. His notion of an intrinsic kind of value, then, seems near to Lewis's notion of an intrinsic property.

It must be cautioned that Moore uses the term 'intrinsic *property*' in such a way that not all properties had solely in virtue of the intrinsic natures of their instances are intrinsic properties. Moore claims that goodness is not an intrinsic property despite the fact that it depends on the intrinsic natures of the things that have it (1922, p. 273). All intrinsic properties are properties that depend on the intrinsic natures of their instances, according to Moore, but not vice versa. Moore has an interesting struggle with the question of the difference between these two notions, and admits that he cannot draw the distinction. He says:

> I can only vaguely express the kind of difference I feel there to be by saying that intrinsic properties seem to *describe* the intrinsic nature of what possesses them in a sense in which predicates of value never do. If you could enumerate *all* the intrinsic properties a given thing possessed, you would have given a *complete* description of it, and would not need to mention any predicates of value it possessed; whereas no description of a given thing could be *complete* which omitted any intrinsic property (1922, p. 274).
>
> You cannot say that an intrinsic property is a property such that, of one thing possesses it and another does not, the intrinsic nature of the two things *must* be different. For this is the very thing which we are maintaining to be true of predicates of intrinsic value, while at the same time we say that they are *not* intrinsic properties (1922, p. 275).

It may be that by 'intrinsic property' Moore meant to express something like what I would express by 'perfectly natural property'. However, I would say that to completely describe an object, one must do more than say what perfectly natural properties it has—one must say what perfectly natural properties are had by its parts and in what perfectly natural relations those parts stand. (In fact, even this is not enough—see section 4.2.1).

1.2 Naturalness

I believe that it would be permissible to take one of our two notions as a primitive. But a similar option, the one I will take in this dissertation, is to take the related notion of a **natural** property as a primitive, and to define 'intrinsic' and 'duplicate' in terms of 'natural'. This is the suggestion of David Lewis. I share Lewis's enthusiasm for the importance of naturalness. Indeed, this dissertation is foremostly devoted to extending Lewis's project of introducing naturalness into ontology as a primitive notion and harvesting the theoretical benefits.

In chapter 3 I will devote significant effort to explaining the notion of naturalness; for now, the following important quotation from Lewis will serve to introduce the notion (1986*c*, p. 60).

> Sharing of [the perfectly natural properties] makes for qualitative similarity, they carve at the joints, they are intrinsic, they are highly specific, the sets of their instances are *ipso facto* not entirely miscellaneous, there are only just enough of them to characterise things completely and without redundancy.
>
> Physics has its short list of 'fundamental physical properties': the charges and masses of particles, also their so- called 'spins' and 'colours' and 'flavours', and maybe a few more that have yet to be discovered...What physics has undertaken...is an inventory of the [perfectly natural properties] of this-worldly things.

Naturalness is best thought of as coming in degrees, the most natural properties being the **perfectly natural** properties. The perfectly natural properties are an elite minority of the intrinsic properties—the most fundamental properties. They are nature's most basic "building blocks". Naturalness applies to relations as well: the perfectly natural relations are the most fundamental relations. Roughly speaking,[9] the entire qualitative character of our world is fixed once we fix the distribution of the perfectly natural properties and relations over its objects. If present-day physics is on the right track, then the properties of *charge*, *spin*, *mass*, the quark flavors, colors, etc. are among this world's perfectly natural properties, and the spatiotemporal relations are perfectly natural relations.

[9] See sections 3.2.1 and 3.2.2.

It should be highlighted that the notions of intrinsicality, naturalness, and duplication are not thought of as being tied to the perceptual or scientific capabilities of humans. Intrinsic properties need not be detectable, even in principle. Two objects might appear to all of our best efforts, scientific or otherwise, to be perfect duplicates; they may yet fail to be duplicates because they differ with respect to some intrinsic property that utterly evades our capabilities. If this strange property is, in some sense, a "non-physical" property, then the hope of the materialist that the world be completely describable in physical terms would be dashed, but this would not disqualify the property from being intrinsic. The same remark applies to natural properties; these too may transcend our ability to detect them; these too need not be "physical".

In chapter 3 I say more about just what the notion of naturalness comes to, but for now let us press on to the definition of 'intrinsic' and 'duplicate' in terms of 'natural'. We begin first with the definition of 'duplicate'. In fact, this definition is not irrelevant to characterizing the primitive notion of naturalness, for one way to elucidate a primitive notion is to define other notions in terms of it, notions of which we have a prior grasp.

According to Lewis, we may define 'duplicate' as follows:[10]

(D3) Possible objects *x* and *y* are *duplicates* iff there is a one-to- one correspondence between *x*'s parts and *y*'s parts such that corresponding parts stand in the same perfectly natural relations and have the same perfectly natural properties

'Intrinsic' may then be defined by (D1).

I do not claim that the notion of naturalness is as near to commonsense as that of intrinsicality or duplication. Still, I think it is reasonable to adopt it as a primitive. First, I do think that naturalness is reasonably commonsensical. We do have a notion of the fundamental properties after which physics seeks; I hope to convince you of this in chapter 3. Secondly, naturalness seems to have vast philosophical utility; I take it that this is a good reason for adopting the notion. For example, it enables analysis of intrinsicality and duplication. The notion bears other fruit as well; various applications of naturalness are discussed by Lewis in "New Work for a Theory of Uni-

[10] Lewis (1986*c*, p. 61). In section 4.2.1 I argue that (D3) must be revised.

versals", and by Anthony Quinton in "Properties and Classes".[11] In chapter 9 I discuss yet another application: in the philosophy of space and time.

Now it can be seen why I do not think the notions of this dissertation are "subjective". If naturalness, for example, were subjective then so too would be all the distinctions drawn in terms of naturalness, and this is something I would not want to admit. For example, I use naturalness in chapter 9 to characterize the facts about the geometrical structure of physical space. If naturalness were subjective, then so too would be the geometry of physical space.

Although Lewis has been instrumental in bringing the notion of naturalness to the attention of contemporary writers, others have discussed related notions. Demopoulos and Friedman (1985, pp. 635–7) discuss a notion of "foundedness" that Carnap introduced in Carnap (1929). Foundedness applies to the extensions of relations that are "experienceable and 'natural"' (quotation from Demopoulos and Friedman p. 636). The tie between foundedness and experienceability is a point of difference between foundedness and Lewis's naturalness. Quinton (1958) introduces a distinction between natural and unnatural classes. Quinton's naturalness comes in degrees (p. 47). However, Quinton links naturalness with human capabilities—he says of natural classes that "people who are introduced to a few of their members can go on to pick out others without hesitation…" (p. 47). George Bealer has a related distinction: among the properties and relations are to be found an elite minority of "qualities" and "connections". However, Bealer's distinction does not admit of degrees and is tied to human perception. See Bealer (1982, pp. 177–83). Quine ties the notion of a natural kind to similarity, subjunctive conditionals, and causation in Quine (1969). He takes a characteristically dim view of these notions. Also, D.M. Armstrong's theory of sparse universals, as presented in Armstrong (1978*a*,*b*), has commanded a good deal of attention, and many of the philosophical jobs for naturalness are likewise jobs for sparse universals. See chapter 6 of this dissertation on sparse universals.

[11] See also Bricker (1991, 1992), although the notion of naturalness Bricker employs is an extension of the notion I consider since it applies to "mathematical" objects.

1.3 The Structure of the Dissertation

In this dissertation I will explore various issues surrounding our three notions. I hope that this chapter has at least convinced you of the existence and potential importance of the notions. Chapter 2 makes explicit my background metaphysical assumptions and introduces the language I will be using throughout the dissertation to talk about entities and their properties.

In chapter 3 I take up the task of presenting a theory of natural properties and relations. I do not define 'natural'—rather, I distinguish and elaborate different conceptions of naturalness, and select one on which to focus in this dissertation.

In chapter 4 I produce a theory of various notions related to naturalness: the notions of intrinsic properties, duplicates, internal and external relations, intrinsic profiles, and qualitative properties and relations. I follow David Lewis in analyzing these notions in terms of naturalness. However, because of an issue raised in chapter 3, Lewis's account of the relation between naturalness and duplication is too simplistic. One task of chapter 4 is to make the necessary complications. Further tasks include investigating the relation between duplication and intrinsicality, and stating various consequences of the definitions that I offer.

If, as I think likely, naturalness cannot be analyzed, then what is the proper response? David Lewis expresses sympathy with the project of including naturalness in our ontology as an unanalyzed concept. I call this approach **primitive naturalism**, and while I am sympathetic to primitive naturalism, I do not believe it can be incorporated into Lewis's ontological system. In particular, Lewis's "class nominalism" is the culprit; chapter 5 develops this objection.

The question of whether we can analyze a notion is vague until we have specified what concepts may be employed in the analysis. In chapter 8 I will consider some proposed analyses of our concepts using only "quasi- logical" notions such as modal and mereological concepts, and also spatiotemporal concepts. But an analysis of naturalness may be possible in terms of a *sparse theory of universals*.

The distinctive feature of a *sparse* theory of universals is that a universal is not admitted for every meaningful predicate or every class of possible objects; only a select few classes of objects and predicates correspond to universals. Intuitively, a predicate expresses a sparse universal only if whenever it applies to two objects, those objects must thereby be *genuinely similar*. 'Having unit positive charge' may express a universal, but 'not having unit positive charge' would not; neither would 'having unit positive charge or having unit negative charge'. Certainly, 'having a friend named "George"' would not express a universal.

The concept of genuine similarity just alluded to perhaps requires comment. Genuine similarity is *not* just a matter of shared properties. Two objects can share the property *being owned by Ted* and still not be truly similar in any respect. Neither is genuine similarity a mere matter of shared *intrinsic* properties, for completely dissimilar objects may share the intrinsic property *not being spherical*.[12] Now it must be granted that there are everyday uses of the word 'similar' that do not express this concept that I have called "genuine similarity". The sentence "John and Jane are similar in that each has a doctor for a father" seems to be capable of expressing a truth, given an appropriate context. The meaning of the word 'similar' is context-dependent. Still, I think there is an intuitive notion of *genuine* similarity that is appropriately "objective", and it is this sort of similarity that concerns us here. Consider two perfect duplicate black marbles, x and y, and a pink elephant z. I find it compelling to say that it is an objective fact that x and y resemble each other perfectly, but neither resembles z.

If there are universals and they are properly sparse, then, it is claimed, they may afford definition of 'perfectly natural':[13]

(D4) Property P is *perfectly natural* iff there is some universal U such that the set of P's possible instances is identical to the set of U's possible

[12]In section 4.2.3, principle A3, which states that the set of intrinsic properties is closed under the Boolean operations, is shown to follow from the definition of 'intrinsic'.

[13]A variant on the universals definition of 'perfectly natural' should be mentioned. Some theorists replace universals, which are supposed to identically recur as non-spatiotemporal parts of their instances, with "particularized properties" called **tropes**. See Stout (1923); Williams (1952); Campbell (1981). Suppose the universals theorist explains the similarity of two unit positively charged protons by postulating a shared universal of unit positive charge. The tropes theorist would say rather that there is a distinct trope of unit positive charge for each proton. Each trope of unit positive charge is exactly similar to every other—indeed, the similarity relation among tropes is taken as a primitive equivalence relation. The tropes theorist offers an analog of (D4):

(D5) Property P is *perfectly natural* iff there is a family of similar possible tropes S such that the set of P's possible instances is identical to the set of possible objects that instantiate members of S

I will not discuss tropes further in this dissertation.

instances

Chapter 6 critically discusses this project of analyzing naturalness in terms of sparse universals, focusing primarily on the work of D.M. Armstrong. In that chapter I also defend the theory of "abundant transcendent" universals against attacks from Armstrong, for that account of universals seems well-suited to my purposes.

In chapter 7 I will address Michael Dunn's criticisms of Lewis's definition of 'intrinsic'. I conclude that the criticisms are largely based on misunderstanding.

In chapter 8 I ask questions of analysis. First, I ask whether naturalness can be analyzed purely in "quasi-logical" terms, then whether it can be analyzed in terms of the notion of "qualitative" properties and relations. Then I consider the question of whether intrinsicality can be analyzed in "quasi-logical" terms. To each question I answer "no", but in each case my evidence is rather meager: the failure of a few proposed attempts. In the second half of chapter 8 I consider proposals by Jaegwon Kim and Michael Slote.

In chapter 9 I will present a new application for the notion of naturalness: the statement of "metrical realism" in the philosophy of space and time. We realists who believe in facts about the metric structure of space, time, and spacetime have a challenge: in what do these facts consist, and how does our language express them? Naturalness provides a solution.

Chapter 2

Preliminaries

2.1 Metaphysical Assumptions

I want to lay out some metaphysical assumptions that I will make in this dissertation. Or rather, they will *appear* to be my metaphysical assumptions. On many I have no opinion, and most have no effect on the arguments in the dissertation. Even though I do not commit myself to these assumptions, there is still value in "taking a stand". The alternative would be a dissertation cluttered by endless qualification, restatement, etc. Better to make some (somewhat arbitrary) assumptions. As will be clear, paraphrase of what I say will usually be possible for those whose favorite metaphysical doctrines are slighted. I will note some possible paraphrases in this chapter, and then ignore them for the rest of the dissertation.

The assumptions are, for the most part, those in David Lewis's *On the Plurality of Worlds*. There are two reasons for this. The first is that I have found this framework convenient. The second is autobiographical: my thinking about the topics in this dissertation has been influenced at every turn by what Lewis has said.

First, there are possible worlds. I presume the existence of the usual plenitude of possible worlds: one for every way the world could possibly have been.

Most believe that possible worlds talk is in need of paraphrase. One way is to find suitable entities to play the role of possible worlds. For example, if our ontology contains propositions then we might choose to call maximal consistent propositions "possible worlds".[1] In contrast, David Lewis notoriously takes possible worlds talk at face value: as making reference to full blooded "concrete" entities. I do not defend Lewis's modal realism, but I will talk *as if* Lewis is right in his claim that other possible worlds are "of a kind with" the actual concrete world. (For example, I will treat non-actual possibilia as being parts of the worlds they inhabit.)

I need a blanket term for "concrete particular". I choose 'object', reserving the broader terms 'thing' and 'item' to apply to anything whatsoever. The actual objects include tables, people, electrons, etc. If there are points and regions of space, electromagnetic fields, etc., then these too are objects. If there are gods, souls, etc., then these are also objects. However, sets are not objects, nor are properties, propositions, numbers, etc.; nor are universals (see below).

I accept talk of non-actual objects. Excepting Lewis and a few others, the usual wisdom is that talk of other-worldly objects is in need of paraphrase. One strategy identifies these objects with certain properties: their essences.[2] I make no commitment either way on this issue.

I assume the usual characterizations of modal notions in terms of possible worlds and objects. A statement is *necessarily* true if true at all possible worlds, *possibly* true if true at some possible world, *impossible* if true at no world. Similarly, a property is impossible if no object has it, possible if some object has it, and necessary if every object has it.[3] (By 'object' here I do not mean merely 'actual object'; I mean to be quantifying over all objects.) Two properties are incompatible if no object has both; one property *entails* another if every object with the first has the second; two properties are *necessarily coextensive* if each entails the other.

I assume the existence of properties. Moreover, I assume that these are properties construed "abundantly"[4], and so I make no assumption that objects that share a property need be genuinely similar. I make a strong as-

[1]See Plantinga (1976, 1974, pp. 44–5).

[2]See Plantinga (1974, pp. 70–87) on the nature of essences, and Plantinga (1976, pp. 268–72) for an account of the reduction of possibilia to essences. Actually, Plantinga reduces *possible* objects to essences. See below for the distinction between objects and possible objects.

[3]Notice that on this terminology, a possible property might be such that it is not possible that it is instantiated, for it might be had only by *impossible* objects (fusions of objects from distinct worlds). See below for the distinction between possible and impossible objects.

[4]This is Lewis's terminology—see Lewis (1986*c*, p. 59 ff.)

sumption about the abundance of properties:

Abundance for any class S of things (actual or otherwise) there is a property P had by all and only the members of S

P may be extremely unnatural, "disjunctive", or "gerrymandered", but it will still count as a property.

I will also assume that properties are individuated by necessary coextension:

Individuation if properties P and Q are such that for any object u (actual or otherwise), u has P iff u has Q, then P=Q.

This is convenient, for it means that we may specify a property by specifying which things have it. Again, paraphrase is available for dissenters. Where I speak of *the* property had by exactly the things in set S, I can be interpreted as referring to the conjunction of all such properties.

I make analogous assumptions for relations. Relations are individuated by necessary coextension, and for any $n \geq 2$, and any class S of n-tuples of things, there is a relation R such that S is the set of n-tuples that stand in R. Moreover, I assume that for every relation R, there is an integer $n \geq 2$, and a class of n-tuples that is the class of all and only the 'tuples that instantiate R.[5]

The two assumptions, Abundance and Individuation, guarantee a one-one correspondence between the properties and the classes of possible individuals. This means that the finitary and infinitary Boolean operations are well-defined. For example, the conjunction of the properties in a (possibly infinite) set S is the property that corresponds to the set of possibilia that is the intersection of the sets of possibilia that correspond to members of S. I will make free use of facts about these operations on properties. *Mutatis mutandis* for relations.[6]

[5]With this last assumption I eliminate infinite-place relations and multi-grade relations. This I do for simplicity.

[6]A succinct summary of the formal properties of the structure consisting of the set of properties together with the operations of finitary conjunction and disjunction, and negation is to say that it is a **complete atomic Boolean Algebra**. For this structure is isomorphic to the Boolean algebra of the class of subclasses of the class of possible objects, and so by Theorem 5 from Halmos (1963) is a complete atomic Boolean algebra. See Halmos pp. 3–9 on Boolean Algebras, and p. 25 for material on infinitary conjunction and disjunction in Boolean Algebras.

A terminological note regarding 'universal' and 'property'. Sometimes these terms are used interchangeably. I reserve the word 'property' for the abundant properties—the properties that obey the abundance assumption. When I discuss 'universals', I leave it open whether or not they obey this abundance assumption. Some, like D. M. Armstrong, construe universals "sparsely". For example, on Armstrong's view, distinct universals do not have disjunctions (1978*b*, p. 19 ff.). When I discuss sparse universals in chapter 6, 'universal' must be taken to refer to a different sort of entity than the abundant properties.

Note, however, that in chapter 6 I also discuss universals construed abundantly, so that they do obey the abundance assumption. Then, I may simply be taken to be discussing the abundant properties. I keep the term 'universal' for that discussion since Armstrong uses it in his criticisms of that view.

At times I will need to make certain "abundance" assumptions for possible worlds and possible objects. I quoted the customary slogan above: there is a possible world "for every way the world might have been". But what are the different ways the world "might have been"? "Recombination principles" are principles that help to specify what possibilities there are. They take the form "*if* such and such is possible, *then* such and such must also be possible". They tell us when we can "recombine" possibilities to obtain more possibilities. David Lewis has an excellent discussion of principles of recombination in Lewis (1986*c*, section 1.8). I do not commit myself to any general principle of recombination; rather, when the need for recombination arises, I will help myself to whatever particular principles seem useful and plausible.

For each possible world, there are some objects, properties, and relations that exist *at* that possible world. The objects that exist at a world are exactly the *parts* of that world. I discuss *existence at* as applied to properties more fully in section 3.2.1.

In the previous paragraph and throughout the dissertation I use the notions of mereology: the theory of the part/whole relation.[7] In particular, whenever there are some objects, I accept the existence of the mereological fusion, or sum, of those objects—the smallest thing that contains those objects as parts. I assume that the mereological sum of some objects is itself an object.

---

[7] See Leonard and Goodman (1940) on mereology. Also see Lewis (1991, p. 72 note 5) for more references on mereology.

When I accept arbitrary mereological sums I introduce the possibility of an object that does not exist at any one world. For if there is no world at which both x and y exist, then their fusion exists at no world, for it is a part of no world. If an object exists at some world, I call it a *possible* object, otherwise it is *impossible*.[8]

In this dissertation I will be concerned only with properties and relations of possible objects, since I am not interested in properties and relations of impossible objects, nor in properties and relations of sets, properties, universals, etc.[9] When I quantify over properties and relations, I mean to quantify only over properties and relations of possible objects. Moreover, my object-quantifiers will always range over *possible* objects only; after this chapter I will sometimes drop the qualifier 'possible'. I will sometimes say 'possibilia' in place of 'possible objects'.

I will also assume the existence of times (both intervals and instants) and places (both points and regions). Again, I invite disbelievers to paraphrase.

I will assume the concepts of set theory, along with the assumptions about those concepts that are captured in the Zermelo Frankel system.[10] I will make free use of set theoretic notation, concepts, and principles. For example, "$A \cap B$" denotes the intersection of A and B; "$A \subset B$" means that A is a proper subset of B. A function is a set f of ordered pairs such that no two pairs have the same first member. If $\langle x, y\rangle \in f$, we say $f(x) = y$, or that y is the "value of f for argument x". The set of arguments of f is its *domain*; the set of values for f is its *range*; f is said to be a function *from* its domain *onto* its range. If we never have $f(x) = f(y)$ unless $x = y$, then f is *one-one*. "$f \circ g$" denotes the *composition* of f with g: the function h with the same domain D as f such that for any $x \in D, h(x) = g(f(x))$. Where A is the domain of function f and $B \subseteq A$, the "restriction of f to B" is the function h with domain B such that for any $x \in B, h(x) = f(x)$.

So much for the entities I will accept and the concepts I will employ.

[8]An object that is the fusion of several possible objects, but does not itself exist at any world is one example of an impossible object. Though this will not concern us, it is natural to assume that there are no other kinds of impossible objects. That is, I assume that every object is the fusion of some possible objects. Equivalently, the fusion of all the objects is identical to the fusion of all the possible worlds.

[9]Exceptions: I am, of course, interested in the *naturalness* of properties and relations, and I will be interested, for example, in the property *making for similarity* which is had by properties (see section 3.3.1). But I will not be interested in these when I quantify over properties and relations.

[10]See, for example, Mendelson (1987, pp. 222–4).

Next I want to discuss two metaphysical issues. Above, I said things like "object u has property P", rather than saying "u has P *at time* t" or "u has P at t *at world* w". Throughout this dissertation I will always speak of properties (and relations) of objects simpliciter, rather than relative to world or time. This involves taking a (nominal) stand on two metaphysical issues, one involving time, the other involving modality.

The temporal issue is that of identity through time. My assumption will be that objects do *not* endure identically through time; rather, an object "persists" through an interval of time by having instantaneous temporal parts in all the instants that comprise the interval.[11] This assumption allows me to re-interpret claims that are apparently about properties had *at* times by continuants as claims about properties had simpliciter by time-slices.

The modal issue is that of transworld identity. My assumption will be that each possible object *exists at* exactly one world. Each possible object is confined to its own world, and hence has its properties absolutely, not relative to one world or another.

It is standard to use "counterpart theory" to analyze *de re* modal predication, if possible objects are thought of as worldbound (Lewis, 1968). On such a view, an object possibly has property P just in case that object has some *counterpart* in some possible world with property P. However, *de re* modality will seldom be mentioned and so counterpart theory need not concern us much.

Those who disagree with these assumptions should be able to easily paraphrase my remarks. They may construe my talk of properties had *simpliciter* by time-slices as talk of properties had *relative to time and world* by continuants. For example, suppose Fred sits at time t in the actual world. Let F be the time-slice of Fred at time t. I would say that F has the property *sitting*. But one who disagrees with my assumptions could say rather that Fred has the property *sitting* **in** the actual world, **at** time t.

The doctrine of temporally bound individuals raises an issue that requires clarification. Most philosophers who accept the doctrine of temporal parts believe that references to ordinary objects are references to fusions of many temporal parts: "four dimensional space-time worms".[12] 'Ted Sider', on this view, refers to the fusion of all my instantaneous time-slices throughout my life. Consider a property I currently have: the property of *sitting*. According

[11] See Lewis (1986*c*, p. 202) for the "endurance"/"persistence" terminology.

[12] Lewis is a representative example. See Lewis (1986*c*, pp. 202–4).

to temporal parts theorists, what makes the sentence 'Ted is sitting now' true is that the current temporal part of the referent of 'Ted', that is, my current temporal part, has the property of *sitting*. Remember that the referent of 'Ted' is, on this view, a 4D space- time worm that is not itself in any particular posture *simpliciter*; rather, it has temporal parts in various postures. The point I want to stress here is that many properties we attribute in common speech, like the property of *sitting*, are had by instantaneous time-slices, and not by the space-time worms we commonly refer to.

Of course, temporal continuants stand in relations corresponding to properties attributed in commonsense speech. I stand in the relation *sitting at* to the time that is now current—I do this in virtue of my current time-slice possessing the property *sitting*. Moreover, the space- time worms themselves have properties as well as standing in relations. Consider the property of *having a temporal part that exists at time t and is sitting*, or the property of *having temporal parts that span seventy years*. But even for these trans-time fusions, the properties are had simpliciter, and not relative to any time.

There is an analogous modal point. Most properties with which we are concerned are had by worldbound objects. The property *living 70 years* is had by continuant persons bound to a world, not by transworld fusions of persons. Moreover, that property is had by those worldbound objects *simpliciter*, and not relative to this world or that. There is a corresponding relation we might define for transworld fusions of persons, *living 70 years at*, which holds between a transworld object x and a world w iff x has a part in w that is a person who lives 70 years.

Although time and modality are analogous in many ways, there is an interesting disanalogy. It is common among those who believe in temporal parts to take ordinary references to physical objects as being references to temporally extended 4D worms. But few modal realists would take ordinary references to physical objects as being references to transworld entities (and not merely because there are few modal realists to begin with!).[13]

These, then are my metaphysical assumptions. As noted, many are made for convenience only, and for the sake of definiteness. But even given this disclaimer, I realize that some will find my assumptions unacceptable. Perhaps some will find possible worlds talk, say, as being utterly incapable of being paraphrased into an ontologically acceptable language. There is no space here to argue this point, so this dissertation must be seen as beginning:

[13]See Lewis (1986*c*, pp. 218–20).

"*if* you can accept my framework, *then* ...". And even those who cannot accept my assumptions can perhaps find arguments that are analogous to mine within their own systems. In sum, I hope that my discussion of naturalness, intrinsicality, and duplication is independent of most of the metaphysical assumptions I have made.

2.2 Language

Next I want to get clear (informally) on the language I will be using to talk about possible worlds and their contents. Here and elsewhere, I will engage in some use-mention sloppiness in the interest of smooth exposition.

My primary language is a possibilist language much like that discussed by David Lewis in 'Counterpart Theory and Quantified Modal Logic'. It contains no modal operators; in their place we have quantifiers that range over possible worlds (variables: w, w', etc.) and quantifiers that range over all possible objects (variables: x, y, z, u, v, etc.). (Notice that I restrict the latter quantifiers to *possible* objects.) So the sentence "$\forall x Fx$" says that *every* possible object, *both actual and merely possible*, is F. Usually, I will use English phrases like 'for every possible object x' instead of '$\forall x$'; the latter is only an abbreviation—it does not indicate any increased formality or level of rigor.

Occasionally, I will make modal claims using the language of "boxes and diamonds"; I will use locutions like "necessarily", "possibly", "entails", etc. This will usually take place when I am discussing the work of a philosopher who uses this sort of language. In such cases, my words should be interpreted in the way David Lewis outlines in his 1968 paper "Counterpart Theory and Quantified Modal Logic". For example, a *de re* modal claim like "x might have had property P" means that a counterpart of x in some possible world has property P. Moreover, when I use the language of boxes and diamonds, my quantifiers are *actualist*, not possibilist. For example, if I say "$\Box \exists x Fx$" in such a context, this means that at every possible world w, there is some possible object **in** w that is F.

I will also permit quantifiers ranging over *properties* and *relations*, and corresponding variables 'P', 'Q', 'R' etc. These range over all properties and relations that are had by nothing other than possible objects. I will often use variables that have properties or relations as values as if they were predicates. Thus, to assert that object x has property P, I will say simply "Px". When discussing universals, I will use variables like 'U' and 'V'; it will be clear from the context whether the universals quantified over are sparse or abundant.

I will also employ [a nonstandard version of] "lambda abstraction". For example, where ϕ is some formula with exactly one free variable, x, the expression $\ulcorner \lambda x \phi \urcorner$ shall be understood as [a term, not a predicate] denoting the property expressed by $\ulcorner$being an x such that $\phi\urcorner$.

Finally, I will employ the "description operator". Thus, the term '$\iota x F x$' refers to the F if there is exactly one F. Whether this is interpreted possibilist ("the one possible object that is F") or actualist ("the one possible object in such and such possible world that is F") will depend on context. I take no particular stand on the fate of a sentence containing such a term when there is not exactly one F.

Chapter 3

Conceptions of Naturalness

3.1 Introduction

In this chapter I distinguish and discuss different conceptions of naturalness.

Naturalness is to be taken as a primitive, but even without giving reductive analyses it is still possible to clarify and distinguish different notions of naturalness. I will choose one conception of naturalness as primary, and attempt to give a rationale for this choice.

I should note that I do not intend what I say about naturalness to extend to "abstract" entities such as numbers. This is not to say that naturalness does not apply to abstract entities. I simply do not know how to illuminate the application of naturalness to these items.[1]

I assume that naturalness comes in degrees (in section 3.4 I justify this assumption). The primitive I will elucidate is the relation *at least as natural as*, from which the relations *more natural than* and *equally as natural as* may be defined. I discuss these relations in detail in sections 3.2.3 and 3.3.2.

The most natural properties and relations will occupy much of our attention; these are *perfectly* natural. A fundamental assumption about perfectly natural properties that I will not question is the following:

(0) Every perfectly natural property is intrinsic[2]

[1] Phillip Bricker pointed out that my conceptions of naturalness do not seem to apply to abstract entities.

[2] A related principle, which I also affirm, is that every perfectly natural *relation* is *external*. See sections 4.2.2 and 4.2.3.

However, the "converse" of (0), the claim that every intrinsic property is perfectly natural, is implausible. The property *having unit positive charge* seems intrinsic, and if it is then so is its negation, the property *not having unit positive charge*.[3] But only the former would be a serious candidate for perfect naturalness.[4] The set of perfectly natural properties, then, is a proper subset of the set of intrinsic properties.

Which proper subset? David Lewis has the following to say (1986*c*, p.60):

> Sharing of [the perfectly natural properties] makes for qualitative similarity, they carve at the joints, they are intrinsic, they are highly specific, the sets of their instances are *ipso facto* not entirely miscellaneous, there are only just enough of them to characterise things completely and without redundancy…
>
> Physics has its short list of 'fundamental physical properties': the charges and masses of particles, also their so- called 'spins' and 'colours' and 'flavours', and maybe a few more that have yet to be discovered…What physics has undertaken…is an inventory of the [perfectly natural properties] of this-worldly things.

This passage seems to me to contain the seeds of two distinct conceptions of naturalness:

Conception 1 naturalness as fundamentalness

Conception 2 naturalness as the source of similarity

In what follows I investigate how Conceptions 1 and 2 characterize the notion of naturalness. Each Conception will be explicated. As I will show, the Conceptions are inconsistent with each other, but parts of the Conceptions can be shared. The goal will be to clarify the Conceptions of naturalness, partly in an informal intuitive fashion, and partly through articulating various principles suggested by the Conceptions.

[3] The closure of the set of intrinsic properties under the Boolean operations is discussed in sections 4.1 and 4.2.3.

[4] This is so on either Conception. The property *not having unit positive charge* is clearly not one of the most fundamental properties, nor does it make for similarity.

Section 3.2 contains my explication of Conception 1. In section 3.2.1 I clarify the conception of naturalness as fundamentalness by relating the concept of fundamentalness to the concepts of supervenience and microphysicality, and by discussing the naturalness of "combinations" of properties. In section 3.2.2 I discuss an important complication in a principle from section 3.2.1. Whereas in section 3.2.1 I focus mostly on perfect naturalness, in section 3.2.3 I discuss relative naturalness as conceived by Conception 1.

Section 3.3 deals with Conception 2. In section 3.3.1 I articulate Conception 2 of perfect naturalness, and show its relation to that of Conception 1. In section 3.3.2 I do the same for relative naturalness.

In section 3.4 I discuss the choice of a Conception, and I conclude in section 3.5.

Before I begin, I want to make a methodological point. The goal of this chapter is *clarification* of naturalness, not *analysis*. I take naturalness as primitive because I cannot analyze it. Therefore, do not look for non-circular necessary and sufficient conditions for a property's being perfectly natural, or for one property's being more natural than another. At times I will say things that are "circular" in the following sense: I will characterize naturalness by relating it to notions that I ultimately analyze in terms of naturalness. Another kind of "circularity": sometimes I will relate naturalness to naturalness itself, or to near relatives. If such "circularity" enables the reader to fix on naturalness, then it has achieved its purpose. Since my goal is clarification rather than analysis, this is fair play.[5]

3.2 Conception 1

3.2.1 Supervenience, Microphysicality, and Property "Combinations"

The first component of Conception 1 is that the perfectly natural properties and relations are the most "fundamental" properties and relations. In the actual world, these seem to be those investigated by physics—the "most basic building blocks of the universe". However, perhaps at some other worlds the most fundamental properties are, in some sense, "non-physical". I seek

[5]Compare what Sydney Shoemaker has to say about similar "circularities" in Shoemaker (1980, pp. 123–4).

clarification of 'fundamental', but I will not offer any definition. Since I cannot define 'natural', neither can I define any of its synonyms.

It will be seen that the notion of the most fundamental properties and relations is exceedingly difficult to specify. Ultimately, all I will do is discuss principles that seem to "flow from" the identification of naturalness with the intuitive notion of fundamentalness. The ideal situation would have been to come up with a precise version of the notion of fundamentalness, and then to *derive* the principles, but this was not to be.

I turn first to the relationship between supervenience and fundamentalness. Supervenience is a relation of functional dependence (Paull, 1994, chapter 7). When A supervenes on B, this means that the instantiation of the B- properties and relations in some sense functionally determines the instantiation of the A-properties and relations.

The precise concept of supervenience I employ is one of *global* supervenience.[6] For any world w, denote the set of w's objects by "$D(w)$".

> Where A and B are sets of properties and relations, A **supervenes** on B iff for any possible worlds w_1 and w_2, any B-isomorphism between $D(w1)$ and $D(w2)$ is also an A-isomorphism.

We say that *property P* supervenes on set B iff $\{P\}$ supervenes on B.

We must define the notion of an A-isomorphism between sets S_1 and S_2 for a set of properties and relations A. I will use arrows overtop of variables as variables over finite sequences of possibilia. When $\vec{x}$ stands for a sequence $\langle x_1 \ldots x_n \rangle$ and R is a relation such that $R(x_1 \ldots x_n)$, then I will write: $R(\vec{x})$. When f is a one-place function defined over the members of a sequence $\vec{x}$, I will use "$f(\vec{x})$" to denote the sequence gotten from $\vec{x}$ by replacing each member of $\vec{x}$ by the value assigned to it by f. A function f is a *part-whole isomorphism* iff for any x, y in f's domain, x is a part of y iff $f(x)$ is a part of $f(y)$ (note that such a function is guaranteed to be one-one). A function f is an *A-isomorphism* iff f is a *part-whole isomorphism* such that i) for any x in the domain of f and property $P \in A$, x has P iff $f(x)$ has P, and ii) for any sequence $\vec{x}$ of objects from f's domain and relation $R \in A, R(\vec{x})$ iff $R(f(\vec{x}))$.

[6]For references on global supervenience see Paull (1994, chapter 4 section A). The sets A and B need not be closed under the Boolean operations (Paull and Sider, 1992, Appendix). Paull calls my formulation of global supervenience (minus the built-in part-whole relation) "strong" global supervenience, and distinguishes it from other formulations. See Paull (1994, chapter 4 section B).

Finally, f is an *A-isomorphism between* S_1 *and* S_2 iff f is an A-isomorphism from S_1 onto S_2.

I will appeal to a few facts about (global) supervenience:[7]

(S1) If A supervenes on B, and $B \subseteq C$, then A supervenes on C

(S2) If A supervenes on C, $B \subseteq C$, and B supervenes on $C - B$, then A supervenes on $C - B$

(S3) If property P is (finitely) "definable" in terms of the properties and relations in set S plus the part-whole relation, then $\{P\}$ supervenes on S.

(S4) Supervenience is transitive and reflexive

Where A and A' are sets of properties and relations, say that A and A' are "negation-images"—"$N(A,A')$"—iff for every property or relation P in either set, the other set contains either P or ∼P (or both). The final principle is:

(S5) If $N(A,A')$ and $N(B,B')$, then A supervenes on B iff A' supervenes on B'

Conception 1 involves a supervenience relation between the perfect natural properties and relations and certain other properties and relations. Once you fix the distribution of a world's perfectly natural properties and relations over its objects, you have fixed the distribution of a certain wider range of properties and relations over those objects. But, as we will now see, it is a tricky matter to say exactly which properties and relations those are.

For example, we should not claim that all properties instantiated at a world supervene on the set of perfectly natural properties and relations at that world, for *haecceities* do not supervene on the perfectly natural properties and relations. Consider a world w containing two objects a and b, neither with any proper parts, and no other objects. For simplicity, let us

[7](S1) and (S4) are trivial. (S3) is proved in the appendix of Sider (1991). (S2): assume the antecedent, and let f be a $C - B$ isomorphism. Since B supervenes on $C - B$, f is a B-isomorphism; thus, f is a C isomorphism. Thus, by the supervenience of A on C, f is an A-isomorphism. (S5) follows from the fact that if $N(C,C')$, then any C-isomorphism is a C'-isomorphism.

suppose that the binary spatial *distance* relations are the only perfectly natural relations at w, and that the only perfectly natural property at w is property P. Suppose further that objects a and b are 1 meter apart, and that each has property P. Let R be the set of perfectly natural relations at w, and let $N = R \cup \{P\}$. Thus, N is the set of perfectly natural properties and relations at w. Call the fusion of a and b "$a+b$". Notice that $D(w) = \{a, b, a+b\}$. Let set B contain just one property: the property of *being identical to a*. The following function f is an N- isomorphism between $D(w)$ and itself, but it is not a B-isomorphism. Hence, B does not supervene on N:

$$f: \quad \begin{array}{lcl} a & \Rightarrow & b \\ b & \Rightarrow & a \\ a+b & \Rightarrow & a+b \end{array}$$

Function f is clearly a part-whole isomorphism. Since a and b each has P, the only property in N, f preserves the properties in N. As for the relations, by the symmetry in the world the function f preserves these as well. For example, since a and b stand in the relation *being 1 meter from*, b and a stand in this relation as well. So f is an N-isomorphism. But it is clearly not a B-isomorphism, for a has the property *being identical to a*, whereas $f(a)$ (that is, b) does not.

The natural conclusion to draw from this example is that we should require only *qualitative* properties and relations to supervene on the perfectly natural properties and relations. 'Qualitative' here is not supposed to be opposed to 'quantitative'. Rather, 'qualitative' is intended in the sense of "purely descriptive", "purely general", and "non- haecceitistic". Qualitative properties and relations do not "involve" particular objects. The idea is that specifying the distribution of perfectly natural properties and relations need not fix the identities of the objects that exist there, but only the qualitative facts about the world, such as the number of objects that exist there and their qualitative properties and relations. See section 4.2.2 for information on qualitativeness.

Say that a set S is an *Q-base* for a world w iff the set of qualitative properties and relations at w supervenes on S. Conception 1 seems to involve the claim that to fix the qualitative properties and relations of a world, you need only fix that world's perfectly natural properties and relations:

(1) for any world w, the set of perfectly natural properties and relations at w is a Q-base for w.

There are complications with (1), however. One of these involves the possibility of worlds with no perfectly natural properties, but I defer discussion of this issue until the next section since there is a more pressing problem. (1) employs the notion of a property existing *at* a world. This notion requires comment.

In one sense of "exists at", a property exists at a world iff it has an instance in that world, but this is *not* the sense I intend. I will show that, thus interpreted, (1) would be equivalent to the following:

(1?) The set of *all* qualitative properties supervenes on the set of *all* perfectly natural properties and relations

For any world and any property or relation P, either P or its negation will be instantiated at that world (I assume for simplicity's sake that every world has at least one object)[8]. So the set of qualitative properties and relations *at* a given world will contain, for each qualitative property or relation whatsoever, either it or its negation. Similarly, the set of perfectly natural properties and relations *at* a world will contain, for each perfectly natural property whatsoever, either it or its negation. By (S5), (1) is equivalent to (1?).

(1?) is not what I intend. I do not deny (1?), but I want to assert something stronger. I want to assert that properties like *being a hydrogen atom*, *being within 10 feet of something cubical*, etc. supervene on the properties and relations of physics alone, the actual perfectly natural properties and relations, without the help of alien perfectly natural properties and relations.

What, then, is it for a property to exist at a world? I myself have no analysis of this notion, and if none is to be had it should be taken as primitive. This is not suspect, I say, for existence-at seems to be an extremely basic, quasi-logical notion. Other such basic notions, such as the part-whole relation, are legitimately taken as primitive. Each possible world has various items that exist *at* that world. There are the concrete objects that exist at that world. But there are also the properties and relations that exist at that world. And it seems to make good sense to say that not all the properties

[8]Even without this assumption, (1) is equivalent to (1?). Assume (1). Take any world w with at least one object; since for every property either it or its negation is instantiated in w, we can apply (S5) to derive (1?). Now assume (1?), and let w be any world. If w has at least one object, then we may again use (S5) and (1?) to derive that the set of perfectly natural properties and relations at w is a Q-base for w. On the other hand, if w does not have at least one object, then the same result holds, for the set of qualitative properties and relations at w is then the null set, which trivially supervenes on every set. So, (1) follows.

and relations that there are exist *at* the actual world, just as it makes sense to say that not all the possible individuals exist at the actual world.

Though I cannot define 'exists at', I can say a bit about it. The properties that exist at a world should be closed under the (infinitary) Boolean operations. They should also be closed under structural combinations (see the end of this section). Moreover, if a perfectly natural property is instantiated at a world, then it exists at that world.

If this notion of *existence at* makes sense, then (1) does not collapse to (1?), for it is not true, given this notion of *existence at*, that for every property and every world, either the property or its negation exists at the world.

Let us return to the elaboration of Conception 1. Conception 1 says more than (1). (1) merely says that the set of perfectly natural properties and relations at a world is a Q-base for that world. But each world has many Q-bases. For example, by the reflexivity of supervenience (see (S4)), the set of any world's qualitative properties and relations will be a Q-base for that world; moreover, if S is a Q-base for w, then by (S1) any superset of S will be as well. The set of w's perfectly natural properties should be the "most fundamental" Q-base for w. As Lewis says in the quotation above, there are only enough perfectly natural properties to characterize the qualitative character of the world *without redundancy*. But what is meant by this?

If S and S' are each Q-bases for w, and S' is a proper subset of S, then S seems less fundamental than S'. We do not need to make reference to all the properties in S to characterize w qualitatively. Call S a *minimal Q-base for* w iff S is a Q-base for w and no proper subset of S is a Q-base for w. We might think to cash out the non-redundancy intuition as follows:

(1′) for any world w, the set of perfectly natural properties and relations at w is a **minimal** Q-base for w.

However, there is a complication, which we can call the "problem of minimality". Suppose we have some properties or relations that are, intuitively, equally natural. Yet, given some of them, we can define the others. Suppose for example that P, Q, and R are perfectly natural, and yet P is definable from (and hence supervenes on, by (S3)), $\{Q,R\}$. This will violate (1′). For suppose that S is the set of perfectly natural properties at w, and suppose that P, Q, and R are members of S. Since $\{P\}$ supervenes on $\{Q,R\}$, by (S1) and (S2), anything that supervenes on S supervenes on $S-\{P\}$ as well. Hence S is not a minimal Q-base for w.

Here is a concrete instance of the problem of minimality. We need a perfectly natural asymmetric relation. The relation *temporally earlier than* seems a likely candidate. But if *earlier than* is perfectly natural, then surely its converse, the relation *later than*, is also perfectly natural. But since *earlier than* and *later than* are interdefinable, by (S3) each supervenes on the other. Thus, the argument in the previous paragraph refutes ($1'$). This is an instance of what we can call "the problem of permutation", a subproblem of the problem of minimality.

Let us introduce the concept of a "permutation" of a relation. *Earlier than* is a permutation of *later than*. If relation B is the temporal *betweenness* relation, holding between x, y, and z iff x is earlier than y which in turn is earlier than z, then the following relation would be a permutation of B: the relation holding between x, y, and z iff y is earlier than x which in turn is earlier than z. Generalizing, let R be an n-place relation. R' is a *permutation* of R iff i) R' has n places, and ii) there is a one-one function f from $\{1,..,n\}$ onto itself such for any possibilia $x_1 \dots x_n$, $R(x_1 \dots x_n)$ iff $R'(x_{f(1)} \dots x_{f(n)})$.

We might try to protect ($1'$) from the problem of permutation by revising the definition of 'minimal Q-base' as follows. A set S would be taken to be a minimal Q-base for world w iff S is a Q-base for w, and for every S' that is a proper subset of S and a Q-base for w, $S - S'$ contains only permutations of members of S'. Intuitively, this revision says that the only redundancy allowed in the set of perfectly natural properties and relations at a world is the redundancy of containing both a relation and one of its permutations.

I do not think that this revision would be *ad hoc*. Suppose I describe the world by saying that a certain event e_1 happened earlier than another event e_2. Suppose you say instead that e_2 happened later than e_1. In a sense, we said the same thing—our sentences had the same factual content. There is a single fact about the world that we described in different ways. I am quick to grant that there is a clear sense in which we do *not* say the same thing. On one sense of the word 'proposition', we expressed distinct propositions: the proposition that e_1 is before e_2, and the proposition that e_2 is later than e_1.[9] But there is also a sense of 'proposition' according to which the propositions that e_1 is before e_2 and the proposition that e_2 is later than e_1 are identical: the sense of 'proposition' on which necessarily coextensive propositions are always identical. So, allowing R and also a permutation of R in the set of perfectly natural properties and relations at a world seems not too distant

[9] A discussion with Fred Feldman was helpful here.

from the intuition that this set shouldn't be "redundant".

This revision to the definition of 'minimal Q-base' would be unnecessary if relations were always identical to their permutations. Timothy Williamson argues for a special case of this thesis in his intriguing 1985 paper "Converse Relations": that a binary relation is always identical to its converse. His defense of this view presumably carries over to a defense of the stronger view that a relation is always identical to each of its permutations. Though I believe that in his paper Williamson answers the obvious objections to this thesis, I will stop short of adopting this defense against the problem of permutations.[10]

Even if we could solve the problem of permutation with the revision to the definition of 'minimal Q-base', we would still have the more general problem of minimality. Though I have no detailed example of which I am certain, I cannot rule out the possibility of a violation of (1′).

For example, perhaps there are a number of spatial relations that are perfectly natural:

point x is linearly between points y and z

points x and y are equidistant from point z

segment x is the segment between points y and z

segment x is congruent to segment y

segment x is longer than segment y

etc.

It seems open that some of these are definable in terms of others. Perhaps there are numerous proper subsets of this set of relations, each sufficient to define the entire set. And yet all the listed items seem equally natural.

In conversation, Phillip Bricker has emphasized a formal analogy between the problem of minimality and the familiar existence of multiple bases for the definition of the truth functions. We can start with negation and conjunction, and define disjunction, material implication, and the rest. We can instead begin with disjunction and negation. And there are other bases. None of these choices seems the most "natural". Perhaps there is an analogy in the realm of natural properties.

[10] If this route were taken, my assumption governing the abundance of relations from section 2.1 would have to be revised.

Thus, I am unwilling to affirm (1′). The problem of minimality remains unsolved, and I come to a disappointing conclusion of my discussion of the relation between fundamentalness and supervenience.

Next, I want to discuss *microphysicality*. In our world at least, microphysicality and fundamentalness seem to go hand in hand. For example, *having spin 1/2* seems more microphysical than, and also more fundamental than, *redness*. I think we should entertain, then, the following claim:

(2) If property or relation P is more microphysical than Q, then P is more fundamental and hence more natural than Q

But (2) contains a term whose meaning is less than transparently clear: 'more microphysical'. Note the following example: let $P =$ *being an electron that is a part of a friend of George Bush*; let $Q =$ *redness*. P is clearly less natural than Q, and yet those objects that instantiate P are electrons, whereas the objects that instantiate Q are macroscopic objects. The moral is that 'P is more microphysical than Q' had better not mean merely that P's instances are smaller than Q's instances.

Instead of attempting a definition of 'more microphysical', I would rather consider a more precise principle. The following principle flows naturally from the intuition that greater microphysicality entails greater fundamentalness, and hence greater naturalness according to Conception 1:

(3) If object x has proper parts, then some proper part of x has a property P that is more natural than any property of x

A possible ambiguity should be pointed out. The property *having unit negative charge* seems like a perfectly natural property. But we often say that macroscopic objects have a "negative charge of one" if their net charge is equal to that of an electron. Does this force us to give up (3)? No. We should distinguish the property of *having unit negative charge* from the property *having **net** unit negative charge*. Net charge and charge may be related as follows:

x has *charge* r iff x has *net charge* r and no proper part of x has any *net charge*

x has *net charge* r iff the sum of the *charges* of parts of x is r

Macroscopic objects have net charge but not charge. As I see it, charge is more fundamental and hence more natural (on Conception 1) than net charge.

Principle (3) has the virtue of circumventing the problem of defining 'more microphysical'. It also has the virtue of bringing out possible objections to the connection between microphysicality and naturalness. Imagine a very bland possible world called "Vanilla".[11] I call Vanilla "bland" (and "Vanilla") because the only intrinsic properties that exist at Vanilla are those that involve mass.[12] Properties such as the properties of charge, spin, charm, etc. are alien to Vanilla. An objector to (3) might argue as follows. Consider an object x from Vanilla that has proper parts. Suppose that x has mass m. It seems that x's most natural property would be *having mass m*. The objector would ask: what proper part of x could have a property that is more natural than this? The only possible candidates are other mass properties, and surely every mass property is as natural as every other.

I regard this as a formidable objection to (3), but I do not think it is conclusive. The defender of (3) could make a distinction analogous to my distinction between *unit negative charge* and *net unit negative charge*. Let us suppose that Vanilla can be decomposed exhaustively into (mereological) atoms—objects with no proper parts. In other words, Vanilla contains no "atomless gunk".[13] There is, then, a certain subset of the mass properties instantiated at Vanilla that we might call the *micromass* properties. Micromass is had only by atoms, whereas mass is had by atoms and non-atoms alike. More carefully,

> x has micromass $m =_{\text{df}} x$ has mass m and x is a mereological atom.

(Indeed, by analogy to 'net charge', mass might be more properly called 'net mass'.) It could be argued that the micromass properties are more fundamental and natural than the other mass properties. If this is correct, then every non-atom x of Vanilla will have a proper part with a property more natural than any of x's properties. The part is any of x's atoms, and the property is that atom's micro-mass. So principle (3) yields the correct result in this case.

[11] I thank David Cowles and Phillip Bricker for helpful discussions about Vanilla.

[12] I assume the spatiotemporal *relations* apply to the objects at Vanilla.

[13] An object is atomless gunk iff it contains no atoms; i.e. iff it has no parts (proper or otherwise) that lack proper parts.

The evaluation of this response is a tricky affair, but I will not pursue it further, for I think that (3) is in trouble on independent grounds. (3) entails the following:

(3′) Perfectly natural properties are had only by mereological atoms—that is, by objects without proper parts.

But (3′) seems overly bold. Couldn't there be a world in which a certain perfectly natural property *P* holds independently of the perfectly natural properties of parts of its instances? *P* might, for example, be an irreducibly mental property.[14] We believe the macroscopic qualitative properties of the *actual* world to depend on the more fundamental properties of microscopic entities, but I see no reason why this must be true *of necessity*. Thus, at such a world, an object with proper parts could have a perfectly natural property. Hence, (3′) is false. The problem with (3′) (and also (3)) is that it makes a contingent feature of our world into a necessary feature. Insofar as I understand (2), I reject it also. The connection between naturalness and microphysicality is contingent.

So much for the relation between naturalness and microphysicality. Next I want to look at the naturalness of "combinations" of perfectly natural properties and relations. Intuitively, the perfectly natural properties and relations should be "simples", not combinations. But more precision is needed.

First, Boolean combinations.[15] Say that a property is *conjunctive* iff it is the conjunction of two other properties (similarly for 'disjunctive' and 'negative'). We must not be tempted by the following:

(?) If *P* is a conjunctive, disjunctive, or negative property then *P* is not perfectly natural

for almost *every* property is conjunctive, disjunctive, and negative.[16] Consider the property *having unit negative charge*. It seems likely that this property is perfectly natural, and yet it is the conjunction of *having unit negative charge or being a part of a green thing* and *having unit negative charge or not being a part of a green thing*. Also, it is the disjunction of *having unit negative charge and being part of a green thing* and *having unit negative charge*

[14]Phillip Bricker made this point.

[15]I intend what I say about Boolean and structural combinations to carry over to relations.

[16]I say "almost" because the universal property had by everything is not conjunctive, and the empty property had by nothing is not disjunctive.

and not being part of a green thing, and it is the negation of *not having unit negative charge*.

However, I do think that Boolean combinations of distinct *perfectly natural* properties are less fundamental and hence less natural than the originals. I cannot derive this result. But on Conception 1 there are only supposed to be enough perfectly natural properties to characterize things "without redundancy". Including $P \wedge Q$ as well as P and Q among the perfectly natural properties seems to go against this spirit. Moreover, the property *having unit negative charge and having spin ½* seems, intuitively, less fundamental than both *having unit negative charge* and *having spin ½*. Physics textbooks mention only the latter.

There are non-Boolean ways to "combine" properties and relations to form properties. Supposing *unit positive charge* to be a perfectly natural property and *being ten feet from* to be a perfectly natural relation, we have the following:

> *being such that something has unit positive charge*
>
> *being ten feet from something*
>
> *being ten feet from something with unit positive charge*

These properties should not turn out perfectly natural. However, this needs no special provision. These properties are not intrinsic and are thus ruled out by principle (0). Clearly, each can differ between perfect duplicates.

There are, however, non-Boolean combinations that are not ruled out by (0): *structural* combinations. Given certain properties and relations, we can combine them to form a property had by an object when its *parts* instantiate those properties and relations in a certain way.[17]

This is best introduced by example. Begin with the properties of *electronhood*, *protonhood*, and a relation *bonded*. The property *being an Hydrogen atom* is had by an object x iff x is composed of a proton and an electron bonded to each other. In such a case, we say that *being a Hydrogen atom* is a "structural combination" of *electronhood*, *protonhood*, and *bonded*.[18]

More precisely, say that P is a structural combination of the properties and relations in set A iff P is denoted by some sentence of the following form:

[17]This discussion could be generalized easily to apply to structural *relations*.

[18]For a related example, see Lewis (1986*a*, p. 27).

the property of being the fusion of n possible objects $x_1 \dots x_n$ such that $\phi(x_1 \dots x_n)$

where $n \geq 1$ and ϕ is a conjunction containing only i) conjuncts with purely mereological terms, and i) conjuncts of either of the following forms:

$$Fx_i$$

$$Rx_{i_1} \dots x_{i_m}$$

where i, and $i_1 \dots i_m$ are between 1 and n (inclusive), F expresses some property in A, and R expresses some m-place relation in A.[19] Say that P is a *proper* structural combination of the members of A if P is a structural combination of the members of A, but P is not a structural combination of the members of any proper subset of A. Henceforth, by "structural combination" I mean "proper structural combination".

The importance of structural combinations of perfectly natural properties and relations is that they are intrinsic. This seems intuitively correct, and is a straightforward consequence of the definition of 'intrinsic' offered in section 4.2.2.[20] So, principle (0) does not rule out the possibility that structural combinations of perfectly natural properties are themselves perfectly natural. And yet I think that such structural combinations are *not* perfectly natural. Again, I cannot prove this. It simply seems to me that the property of *being a Hydrogen atom* is less fundamental than *electronhood*, *protonhood*, and *bonded*.

Thus, I accept:

(4) Boolean combinations of other perfectly natural properties, and structural combinations of other perfectly natural properties and relations, are less natural than those perfectly natural properties and relations, and hence are not perfectly natural

[19]This definition could be generalized by allowing infinitary definition. I do not claim that this definition accords exactly with D. M. Armstrong's usage of 'structural universal'. Notice that since I do not require n to be greater than 1, I count the conjunction of P and Q as a structural combination of P and Q.

[20]Consider a pair of perfect duplicates. Their parts have the same perfectly natural properties and stand in the same perfectly natural relations. Thus, they could not differ with respect to any structural combination of perfectly natural properties and relations. Therefore, such properties are intrinsic.

I would like to be able to derive (4) from a general principle stating that the perfectly natural properties are "non-redundant", but I failed to state such a principle when I failed to solve the problem of minimality. (4) must therefore stand on its own. Indeed, by making this claim, I intend to partially specify what I mean when I say that the perfectly natural properties and relations are the "most fundamental" properties and relations.

3.2.2 A Complication

In this section I want to consider the possibility of a world at which there are no perfectly natural properties. This will require modification of principle (1).

It seems possible that the world might be endlessly complex. Consider a possible world that we may call "Onion". At each level, physical properties decompose into still more basic physical properties. There are macroscopic physical properties, micro-physical properties, micro-micro-physical properties...and so on forever.[21]

Let me describe Onion in detail. Intuitively, every fairly natural property at Onion is a structural combination of more fundamental properties and the spatiotemporal relations. More carefully: i) Onion has perfectly natural *relations*: the spatiotemporal relations. ii) for all intrinsic properties P_1 and P_2 at Onion, either P_1 is at least as natural as P_2, or P_2 is at least as natural as P_1 (or both). iii) there is an intrinsic property P at Onion such that for every intrinsic property Q at Onion that is at least as natural as P, Q is a structural combination of properties at Onion that are at least as natural as Q and perhaps spatiotemporal relations.

Here is a more concrete way my description of Onion could be true. For any intrinsic property P at Onion that is at least as natural as the one mentioned in clause iii) above, the following sort of description is true. P is the property of being composed of two parts separated by n feet, one with some property P_1, the other with some property P_2, where P_1 and P_2 are at least as natural as P. P_1 in turn is the property of being composed by two parts separated by $\frac{n}{4}$ feet, one with some property P_3, the other with some property P_4, where P_3 and P_4 are at least as natural as P_1. P_3 in turn is the property of being composed of two parts separated by $\frac{n}{16}$ feet, one with some property P_5, the other with some property P_6, where P_5 and P_6 are at least as

[21] Armstrong considers such a world in Armstrong (1978*b*, pp. 67–8).

natural as P_3. And so on.

I believe that such a world is possible. First, there seems to be no inherent contradiction in the idea. Second, we seem to be able to conceive what such a world would be like in some detail. Such a world would be like an onion, with the layers going on forever. Imagine scientists at such a world engaged in a futile quest to discover a "bottom" to the complexity. Every property they discover would be found to be composed of still more basic properties. Finally, the possibility of endless complexity seems like an open possibility even for the actual world. Once we thought that atoms were mereological atoms. For all we knew, the smallest bits of matter were atoms of Hydrogen, Helium, etc.—hence, we called them "atoms". Then we learned that these "atoms" have proper parts (protons, neutrons, and electrons), the properties of which determine the properties of atoms. Still later we learned that protons and neutrons have still more basic parts (quarks), the properties of which determine the properties of protons and neutrons. Perhaps this will go on forever.

The importance of Onion is that no perfectly natural properties exist at Onion, at least on the Conception 1 conception of naturalness. For suppose that P_1 is a perfectly natural property at Onion. By (0), P_1 is intrinsic. Let P be an instance of the existential quantifier in clause iii) of the description of Onion. By clause ii), one of P and P_1 is at least as natural as the other; since P_1 is perfectly natural, P cannot be more natural than P_1, so P_1 is at least as natural as P.[22] So by iii), P_1 is structurally composed of properties that are at least as natural as it (and perhaps spatiotemporal relations). These latter properties cannot be more natural than P_1 since P_1 is perfectly natural, so they are as natural as P_1. Thus they are perfectly natural.[23] But this violates principle (4) from section 3.2.1, which says that structural combinations of perfectly natural properties are not perfectly natural.

This spells trouble for principle (1). Principle (1) entails that N, the set of Onion's perfectly natural properties and relations, is a Q-base for Onion. But if there are no perfectly natural properties instantiated at Onion, then N just contains the spatiotemporal relations, and hence is clearly *not* a Q-base for Onion. For surely there could be two worlds satisfying my description of Onion, each with objects standing in exactly the same spatiotemporal rela-

[22] Here I rely on principle (e) from section 3.2.3.

[23] Principle (f) from section 3.2.3 guarantees the acceptability of the inference from 'P is equally as natural as a perfectly natural property' to 'P is perfectly natural', whether we mean strong or weak perfect naturalness (see section 3.2.3 for this distinction).

tions, but with objects instantiating different intrinsic properties. Of course, in each world these intrinsic properties will divide endlessly into more basic intrinsic properties, but surely there could be different intrinsic properties of this kind.

It will not do to restrict (1) by saying that *if* the set of w's perfectly natural properties is *nonempty*, then the union of it and the set of that world's perfectly natural relations is a Q-base for w. Consider the following possibility. World w has two regions: one region resembles Onion in having no perfectly natural properties, but the other region *does* have perfectly natural properties.

The solution I propose is to weaken principle (1) as follows:

(1a) If the set N of perfectly natural properties and relations at world w is such that for every qualitative property or relation P at w, some member of N is at least as natural as P, then N is a Q-base for w

To see the idea behind (1a), consider how it applies to Onion. N, the set of perfectly natural properties and relations at Onion, contains only spatiotemporal relations. Is it the case that for every qualitative property P at Onion, some relation in N is as or more natural than P? I think not, because of considerations that arise at the end of the next section. Consider properties from Onion's infinite sequence of increasingly natural properties. Because of principle (e) from the next section, none of these properties could be *equally* as natural as any member of N. So each would have to be less natural. I find this implausible, and I motivate this intuition at the end of the next section. So, Onion doesn't satisfy the antecedent of (1a), and hence we do not get the incorrect result that N is a Q-base for Onion.

Notice that I have not argued for the existence of a world with no perfectly natural *relations*. The existence of a possible world like Onion, but in which the *relations* come in a never-ending sequence of increasing naturalness, seems more controversial. I do not deny the possibility of such a world; I merely do not assert its existence. Onion has importance beyond (1a). We will return to it.

3.2.3 Relative Naturalness According to Conception 1

In section 3.2.1 I focused mainly on perfect naturalness. I turn now to what Conception 1 has to say about relative naturalness. Our basic relation is *as* *or more natural*. I symbolize this relation "$\geq$". We have made use of the relations *equally as natural as* ("$\approx$") and *more natural than* ("$>$"). These may be defined in a familiar way:

Definition: $P \approx Q =_{\text{df}} P \geq Q \wedge Q \geq P$
Definition: $P > Q =_{\text{df}} P \geq Q \wedge {\sim}Q \geq P$

Conception 1 equates naturalness with fundamentalness, but what does relative fundamentalness amount to? As elsewhere in this chapter, I have no analysis, only intuitive pictures. I will first attempt to illuminate relative naturalness and fundamentalness by using an intuitive notion of "distance". Then I will clarify the relation further by discussing its formal properties.

Again, I draw on a quotation from David Lewis: (1986*c*, p. 61)

> Some few properties are *perfectly natural*. Others, even though they may be somewhat disjunctive or extrinsic, are at least somewhat natural in a derivative way, to the extent that they can be reached by not-too complicated chains of definability from the perfectly natural properties.

To elaborate the idea in this passage I invoke an intuitive notion of "distance" of a property from a set of properties (suggested by Phillip Bricker). Property or relation P is "at least as close" to a given set as property or relation Q if P can be at least as directly and non-disjunctively defined from that set's members as Q (that is, if there is no definition of Q from S that is more direct and non-disjunctive than every definition of P from S). Finite definability in terms of a set makes for closeness; if a property can only be defined using infinitary means, this makes for greater distance. Disjunctive definitions make for more distance than non-disjunctive ones of equal length. (I make no attempt to say how these factors weigh off against each other in determining closeness.)[24]

[24]The closeness relation deserves its name because of features it shares with spatial distance. A property is maximally close to itself (actually, to its unit set). Distance from a set was, in essence, defined as distance "along the shortest path" when I said that P is as close as Q to S if it *can* be defined from S at least as directly and non-disjunctively (the "paths" here are particular definitions). Moreover, Phillip Bricker pointed out that there is a sort of triangle inequality for particular paths. Suppose that property P is definable via path P from set S, and that Q is definable via path P' from $S \cup \{P\}$. Hence, Q is definable from S via what we might call the "concatenation" of paths P and P': path "$P + P'$". Think of $P + P'$as being a definition of Q according to the following directions: "first define P from S using definition P. Then *keep going* using definition P' to get Q". The triangle inequality says that the shortest path from S to Q must be at least as short as $P + P'$.

We can define 'closer than' and 'equally as close as' as we defined 'more natural than' and 'equally as natural as' above.

The working idea is that relative fundamentalness is relative distance from N, the set of perfectly natural properties and relations.

Consider the properties of the actual world as an example. I suppose that the most fundamental properties of the actual world are the charges, masses, spins, and the quark charms, flavors, etc., and that these are members of N. Likewise, I assume that the spatiotemporal relations are the most fundamental actual relations, and that they are also members of N. Of the actual properties, closest to N (excepting the members of N themselves) are properties like *being a proton*, *being a neutron*, and *being an electron*, for these can be defined directly from N. As the distance from N increases, properties get less fundamental. In decreasing order of fundamentalness: *having unit positive charge*, *being a proton*, *being a hydrogen atom*, *being a water molecule*. Properties like *blueness* and *greenness* will presumably admit of no finite definition (although they presumably supervene on N) and will thus be further from N. And even further than these will be Nelson Goodman's properties of *grueness* and *bleenness*, in virtue of their disjunctiveness.[25] At the bottom are properties that do not even supervene on N at all, and hence cannot be defined from N at all, for a property is not definable by any means, finitary or infinitary, from a set on which it does not supervene.[26]

So, the following principle initially seems attractive:

(5) property or relation P is at least as natural as property or relation Q iff P is at least as close to N as Q is

We have the following corollaries:

(5′) $P > Q$ iff P is closer to N than Q is

(5″) $P \approx Q$ iff P is equally as close to N as Q is

However, (5) is incorrect. In worlds like Onion without perfectly natural properties, (5) breaks down. Properties at Onion are not supervenient on N at all, and hence not definable from it at all. So no property from Onion will

[25] See Goodman (1955, p. 74). Let t_0 be a certain fixed time. An object x is grue at time t iff $t < t_0$ and x is green at t, or $t \geq t_0$ and x is blue at t. Similarly, x is bleen at t iff $t < t_0$ and x is blue at t, or $t \geq t_0$ and x is green at t.

[26] This follows from an infinitary version of (S3), which I do not prove.

be closer to N than any other. But surely some properties at Onion are more natural than others. This violates (5′).

There is another difficulty with (5). Every perfectly natural property is equally as close to N as every other, since each perfectly natural property is a member of N and is therefore maximally close to N. (5″) implies, therefore, that every perfectly natural property and relation is equally as natural as every other. I worry that this result is too strong, and so I offer the following weakened version of (5), which takes into account both of (5)'s problems:

(5a) For any set N of pairwise-equally natural perfectly natural properties and relations, and any properties or relations P and Q that supervene on N, $P \geq Q$ iff P is as least as close to N as Q is.

Why not hold that all perfectly natural properties and relations are equally natural? A stronger version of (5a) would have this consequence:

(5a?) For any properties or relations P and Q that supervene on N, the set of all perfectly natural properties and relations, $P \geq Q$ iff P is at least as close to N as Q is.

Well, I do not want to claim that the view that perfectly natural properties and relations are equally natural is false, but neither do I want to rely on its truth. I would like to stay neutral with respect to this view, since I feel some inclination towards the view that some perfectly natural properties are *incomparable* in terms of naturalness with some other perfectly natural properties. Consider, for example, perfectly natural properties that are *alien* to our world. I see no good reason to say that those properties are equally as natural as our world's quark properties.

One potential reason for claiming that an alien perfectly natural property P is equally as natural as, say *unit negative charge*, is that each one has a maximal degree of naturalness. We might put this point by saying that P and *unit negative charge* play identical roles in the structure of naturalness. But I don't think this is a good reason for claiming that P and *unit negative charge* are equally natural. The fact that P and *unit negative charge* are each maximally natural does not *formally* entail that they are equally natural,[27]

[27]It is crucial here that by 'perfectly natural' I mean *weakly* perfectly natural (see below for the distinction). Strongly perfectly natural properties would indeed have to be equally natural.

for the field of a transitive and symmetric relation can have the structure of a forest, with different "trees" having "trunks" with bottoms (the perfectly natural properties and relations) that are nonetheless incomparable in terms of that relation. Perhaps there is a sense in which the degree of naturalness P has "relative to its tree" is equal to the degree of naturalness *unit negative charge* has relative to *its* tree (in spelling out this notion in detail, the "tree" of a property or relation R would be taken to be the set of properties R' *comparable* to R (that is, such that either $R \geq R'$ or $R' \geq R$)). But the original question involved comparisons of naturalness *directly*, and not "relative to trees".[28]

So, to maintain neutrality about the question of whether all perfectly natural properties are equally natural, I endorse (5a) rather than the stronger (5a?). (5a) falls sadly short of offering fully general necessary and sufficient conditions for the $\geq$ relation. It leaves many questions unanswered. Suppose property P supervenes on one set of pairwise equally natural properties, whereas Q supervenes on another; suppose further that neither P nor Q supervenes on the set associated with the other property. Is either P or Q as or more natural than the other? (5a) is silent. Worse, it gives us no answer to the question of when one perfectly natural property is equally as natural as another. Still, it has value, for it relates the *more natural than* relation to distance-from-a-set in a special case—with this I must be content. The answer that (5a?) gives to the last question (namely, the answer that perfectly natural properties are *always* equally natural) seems to me to be incorrect—better not to take a stand than to take a dubious stand.

(5a) does not concern properties from Onion, for those properties, presumably, do not supervene on any set of perfectly natural properties. Intuitively, however, the notion of closeness should apply to properties from Onion in the following way. If we consider a set of all the properties from Onion with a certain degree of naturalness or more, then the *rest* of the properties at Onion should be more or less natural depending on how close they are to that set. This idea may be specified as follows. For any properties or relations P, Q, let $S(P,Q)$ be the set of properties and relations that are *more*

[28] Suppose an actual property P is perfectly natural, but is not comparable in terms of naturalness with some alien perfectly natural property Q. Given certain plausible recombination principles, it follows that there is some world at which *both* P and Q are instantiated. So, at some worlds, properties that are incomparable in terms of naturalness will be instantiated. I owe this observation to Phillip Bricker. It seems to me that there is nothing objectionable in this.

natural than both P and Q.

(5b) For any properties P and Q, if P and Q both supervene on $S(P,Q)$, then $P \geq Q$ iff P is at least as close to $S(P,Q)$ as is Q.

(5b) seems to have correct consequences for Onion's properties. Given a pair of properties P and Q from Onion, $S(P,Q)$ will contain an infinity of properties of which P and Q are structural combinations. It seems that distance from this set will determine the relative naturalness of P and Q. (5b) also seems to have acceptable consequences for worlds containing perfectly natural properties. When we compare, say, *blueness* to *grueness* in terms of naturalness, (5b) instructs us to look at S(*blueness*, *grueness*). This set will contain the perfectly natural properties and relations of the actual world, and all properties definable from them that are more natural than both *blueness* and *grueness*. It seems plausible that *blueness* is closer to this set than *grueness*. Thus, (5b) has the correct result: *blueness* is more natural than *grueness*.

We must include in (5b) the requirement that P and Q supervene on $S(P,Q)$ for various reasons. For one, if P and Q are each perfectly natural, then $S(P,Q)$ will be empty; if every property is as close to the null set as every other (as seems plausible), it would then follow that perfectly natural properties are always equally natural. Or suppose that neither P nor Q is perfectly natural, and yet P and Q are incomparable in terms of naturalness. It seems plausible that no property is more natural than both P and Q, and so $S(P,Q)$ would again be the empty set.

The purpose of (5b) is to help us fix on the notion of relative naturalness. However, it employs that very notion in doing so. While this perhaps limits its value, it does not make it worthless. (5b) is a substantive principle governing naturalness, and is therefore worth asserting in an attempt to illuminate naturalness.

So much for the intuitive picture of relative fundamentalness. Next, let us focus on the formal properties of our three naturalness relations. Here are some principles governing these notions:

(a) $\geq$ is reflexive and transitive

(b) $\approx$ is reflexive, transitive, and symmetric

(c) $>$ is transitive and irreflexive

(d) $\forall P, Q \sim(P > Q \wedge P \approx Q)$

(e) $\forall P, Q[P \geq Q \leftrightarrow (P > Q \vee P \approx Q)]$

(f) $\forall P, Q, R\{(P \approx Q) \rightarrow [(P \geq R \leftrightarrow Q \geq R) \wedge (R \geq P \leftrightarrow R \geq Q)]\}$

I assume (a) as a constraint on the intended interpretation of '≥'; the rest follow from (a) and the definitions of '≈' and '>'.

Equally as important as principles we should accept are those that we should reject. I think we should reject the following principle:

Connectedness $\forall P, Q\,(P \geq Q \vee Q \geq P)$

Let us return to Onion, the world of endless complexity. There are no perfectly natural properties at Onion, only an infinite sequence S of properties of increasing naturalness (this sequence need not increase monotonically, but for every member of the sequence there is another more natural member). Let us compare the naturalness of the members of this sequence and P, one of our world's perfectly natural properties. Either it is, or it is not the case that P is more natural than each property in S. One of the following two alternatives holds:

i) $\forall Q \in S\, P > Q$

iia) $\exists Q \in S \sim P > Q$

(The '∼' here stands for sentence negation, *not* property negation). Because of the definition of '>', iia) is equivalent to the following:

iib) $\exists Q \in S\,(\sim P \geq Q \vee Q \geq P)$

But no member of S could be at least as natural as P. For suppose $Q \in S$ and $Q \geq P$. Since P is perfectly natural, Q cannot be more natural than P; hence, by (e), $P \approx Q$. But there are properties in S that are more natural than Q, and so by (f) more natural than P, contradicting P's perfect naturalness. So iib) is materially equivalent to:

ii) $\exists Q \in S\,(\sim P \geq Q \wedge \sim Q \geq P)$

But ii) contradicts Connectedness, so if we can rule out i), then Connectedness is false.

I think that i) is false. Unfortunately, it is difficult to argue rigorously for this since neither (5a) nor (5b) apply here. I must rest with an intuitive judgment that it is not the case that P is more fundamental than all of the properties from Onion. After all, the properties from Onion bear no interesting supervenience relations to P. Why would they all be less fundamental? I find it more intuitive to claim that P is incomparable in respect of naturalness with some properties at Onion.

I deny i). But it might be objected that *by definition* P is more natural than all the properties at Onion since it is perfectly natural whereas the latter are not. A perfectly natural property, according to the objection, is one that is at least as natural as every other property or relation.

To evaluate this objection, we must resolve an ambiguity in the term 'perfectly natural' that we have not yet discussed. We should distinguish between strong and weak perfect naturalness:

P is *strongly* perfectly natural $=_{df} \forall Q\, P \geq Q$

P is *weakly* perfectly natural $=_{df} \sim\exists Q\, Q > P$

If by 'perfectly natural' we mean 'strongly perfectly natural', then i) is indeed true by definition. For if P is strongly perfectly natural, then (by (e)) for every property $Q \in S$, either $P > Q$ or $P \approx Q$, and as we argued above, P cannot be equally as natural as any property from Onion.

But of course I do not grant that P is strongly perfectly natural. When I assumed that P is perfectly natural, I meant that it is *weakly* perfectly natural. And since the very point I am arguing for is that P is *not* more natural than all of Onion's properties, to assert that P is *strongly* perfectly natural would beg the question against me.

In fact, I don't think there are *any* strongly perfectly natural properties. A strongly perfectly natural property P would have to be at least as natural as every property at Onion, and I don't think there are any such properties. P could not be from Onion (since for every such property there is one more natural). On the other hand, I reject the existence of a property, P, alien to Onion that is at least as natural as each of Onion's properties. P would have to be more natural than each of Onion's properties, because if it were equally as natural as one, it would be less natural than some other property at Onion (remember principle (f)). And I reject P's being *more* natural than each of

Onion's properties for the same reasons I had for rejecting i). Properties alien to Onion do not seem intuitively to be more fundamental than all of Onion's properties. After all, the properties at Onion do not supervene on such alien properties.

In this dissertation, by 'perfectly natural' I always mean 'weakly perfectly natural'.

Let us return to my denial of i). I have responded to one objection: that i) is true by definition since P is perfectly natural. I now want to respond to another objection. Someone might claim that, while i) is not true *by definition*, still it is evident. P has a maximal degree of naturalness, whereas for any property at Onion, there are still more natural properties. Doesn't this make P more natural than every property from Onion?

I don't think so. Perhaps there is a sense in which P is more natural *relative to the actual world* than any property from Onion is, relative to Onion. After all, P occupies a distinguished position in the actual world—unexceeded naturalness. No property from Onion occupies this position. But this has nothing to do with Connectedness. Connectedness concerns comparisons of naturalness directly between properties, not "naturalness relative to worlds".[29]

[29] I have argued that there are no strongly perfectly natural properties. But one might use a thought experiment mentioned to me by Phillip Bricker to argue that no properties are even *weakly* perfectly natural (Bricker did not use the thought experiment for this purpose). Consider a property P of quarks, an alleged (weakly) perfectly natural property. Bricker grants that *in fact*, P is not a structural combination of other properties—it is, in fact, "a simple". But consider another world w at which Ps are always composite objects made up of two parts standing in relation R, each with property Q. The idea here is that P is a structural combination of Q and R *at* w. Mightn't it be plausible to say that Q and R are more natural than P? But if this is correct, then P isn't perfectly natural after all. And this argument seems general.

On an intuitive level, my reply is that the property Bricker imagines that "divides" into more fundamental properties is not P, the same property from the actual world, since "second level" properties like *being a structural combination of* Q *and* R are **essential** properties of those properties that have them.

Now, for more precision. My definition of 'structural combination' entails biconditionals of the form 'P is a structural combination of Q and R iff P *is the property of* ...'. Now, my assumption is that a phrase such as 'the property of ...' picks out a property with the feature expressed by the '...' in *every world*. This sort of assumption is commonplace. For example, most people would assume that the phrase 'the property of having both P and Q' picks out a property that is not merely the actual conjunction of P and Q, but rather is *essentially* the conjunction of those properties. To deny this assumption would be to introduce obscurity in the notion of transworld identity for properties.

3.3 Conception 2

3.3.1 Perfect Naturalness

I hope that section 3.2 has helped to fix the notion of naturalness as fundamentalness. I turn now to Conception 2, which construes naturalness in terms of similarity. I remind the reader that the project here is to illuminate the primitive notion of naturalness. Thus, when I relate naturalness to similarity, this is not to be taken as an analysis, but rather as an intuitive aid in grasping just what this notion of naturalness is.

I need to say a bit about the relevant notion of similarity. Given the abundant construal of properties I accept, every two objects share infinitely many properties, and differ with respect to infinitely many more, since every object is a member of infinitely many sets of possibilia, and likewise fails to be a member of infinitely many sets of possibilia. But not all these shared properties count as *genuine* similarities. Some properties have this feature: when they are shared between two objects, this counts as a genuine intrinsic similarity. Other lack this feature. For short, some properties *make for similarity*. (Let us stipulate that a property P such that there are no *two* possible objects x and y such that Px and Py does *not* make for similarity).

I remind the reader of my section 1.3 remarks on the contextual dependence of the word 'similarity'. In some contexts one might say that two people are similar in that each was born in California. This is *not* the sort of similarity I mean. The notion here is one of "objective" intrinsic similarity.

At the very least, to make for similarity, a property must be intrinsic. But not all intrinsic properties make for similarity. Disjunctions and negations of properties that make for similarity will usually not make for similarity, and yet disjunctions and negations of intrinsic properties are intrinsic. Which intrinsic properties make for similarity? According to Conception 2, it is exactly the perfectly natural properties that make for similarity:

The objection, as I construe it, involves a rejection of this assumption. It runs as follows. i) First, we redo the definition of 'structural combination' to make it world-relative—P is a structural combination of such and such properties *relative to world* w iff P is the property of …*at* w. ii) Then we claim that P (the alleged perfectly natural property from the actual world) is a structural combination of Q and R at some other world w. iii) finally, we claim that it follows that P is less natural than Q and R.

I object to step ii), for I claim that if P is a structural combination of some properties at one world, then it is essentially so. So, if P were a structural combination of some other properties at w, it would have to be at the actual world as well—but by hypothesis it is not.

(6) A property or relation is perfectly natural iff it makes for similarity

(In this section I will focus on properties, although I intend the discussion to carry over to relations. But what does it mean to say that a relation *R* makes for similarity? To every *n*-place relation *R* there corresponds a property: *being the fusion of n objects that stand in R*. To say that *R* makes for similarity is to say that this property makes for similarity in the sense already introduced for properties.)

The first thing to note is that part of Conception 2 must be rejected, for it is at odds with Conception 1. Principle (4), which flows from the intuition behind Conception 1, implies that conjunctions of distinct perfectly natural properties are not perfectly natural. But (6) implies that conjunctions of consistent perfectly natural properties *are* perfectly natural, for if two consistent properties make for similarity, then their conjunction will as well.

(6) also seems to violate the other half of principle (4), which states that structural combinations of perfectly natural properties and relations are not natural. Suppose properties *P* and *Q* make for similarity. Surely the property *being the fusion of two objects x and y such that x is 1 millimeter from y and x has P and y has* Q also makes for similarity (if it is not the impossible property).

The most dramatic difference between the Conceptions, however, emerges when we consider *intrinsic profiles*. Intrinsic profiles are maximally specific intrinsic properties; they are shared by perfect duplicates, and perfect duplicates alone (see sections 4.2.2 and 4.2.3 for more on intrinsic profiles). Clearly, intrinsic profiles make for similarity, and are therefore perfectly natural according to (6).

The situation is different with Conception 1. Intrinsic profiles are typically structural combinations of perfectly natural properties and relations. To see this, consider a very simple object *x*. Let *x* have only two proper parts that stand in just one perfectly natural relation: *being 1 millimeter apart*; one has perfectly natural properties *P* and *Q*; the other has perfectly natural properties *R* and *S*. Then, the intrinsic profile of *x* is the property of *being composed of two mereological atoms separated by 1 millimeter, one of which has P and Q, the other of which has R and S*. This is *x*'s intrinsic profile since an object has this property iff it is a duplicate of *x*.[30] And notice that it is a structural combination of *P*, *Q*, *R*, *S*, and the relation *being 1 millimeter apart*.

[30] See the definition of 'duplicate' in section 4.2.1.

Thus, the intrinsic profiles of certain objects are structural combinations of the perfectly natural properties of their atomic parts, and the perfectly natural relations between those parts. (Exceptions occur for objects with no atomic parts, and also for objects from Onion.) Principle (4) implies that the intrinsic profile of such an object is not perfectly natural.

Moreover, it is clear that the intuition behind Conception 1 bans intrinsic profiles from being perfectly natural. Consider the property that specifies the intrinsic nature of the Eiffel Tower in complete detail. Surely, this extremely specific property would not be counted as fundamental furniture of the universe. Rather, the properties of the Eiffel tower's small parts—charges, spins, masses, etc.—are more fundamental.

The Conceptions can incorporate elements from each other. Although Conceptions 1 and 2 are incompatible, each can accept a part of the other. When I argued that Conceptions 1 and 2 are incompatible, I appealed to the fact that (6) makes making for similarity *sufficient* for perfect naturalness. But it is consistent with Conception 1 that making for similarity, while not sufficient, is nonetheless *necessary* for naturalness. So I build the following thesis into Conception 1:

(7) All perfectly natural properties and relations make for similarity

(7) joins (4) in ruling out perfectly natural properties that are negations of perfectly natural properties. Suppose that *unit positive charge* is perfectly natural. It is plausible that sharing of the negation of this property does not count as a similarity. Surely, two objects that fail to have unit positive charge need not resemble each other in the slightest. (Be sure not to confuse the negation of *unit positive charge* with the property *unit negative charge*. Most things have neither.) (7) also implies that disjunctions of perfectly natural properties will not *in general* be perfectly natural, although it seems consistent with (7) that the disjunctions of *some* pairs of distinct perfectly natural properties are perfectly natural. Suppose, for example, that P and Q are two perfectly natural "very similar to each other" in the following sense: intuitively, every instance of P resembles closely every instance of Q. Perhaps we'd then want to say that when objects share $P \vee Q$, this counts as a similarity.

Since Conception 1 now contains (7) and Conception 2 contains (6), it follows that the set of perfectly natural properties and relations as conceived by Conception 1 is a subset of the perfectly natural properties and relations as conceived by Conception 2. (Moreover, we can say "*proper* subset" since Conception 2 allows perfectly natural conjunctions and structural combinations of perfectly natural properties.)

(Note: (7) is all that David Lewis clearly affirms in the quotation at the beginning of this chapter, so I do not attribute Conception 2 to him. But he does express agnosticism on the question of whether structural combinations of perfectly natural properties are perfectly natural (1986*c*, p. 62), so it may be that he does not reject Conception 2. I cannot tell for sure.)

So Conception 1 can incorporate part of Conception 2. Likewise, Conception 2 can incorporate part of Conception 1. I have in mind principle (1), which states that the set of perfectly natural properties at a world is a Q-base for that world.

In section 3.2.3 I argued that, given Conception 1 of naturalness, Onion is a world at which there are no perfectly natural properties. This required me to take back (1) and substitute the more complicated (1a). But given (6) as a characterization of perfect naturalness, we can take these moves back. Any object at Onion has an intrinsic profile—its most specific intrinsic property,[31] and it was argued above that intrinsic profiles are perfectly natural, given (6). Onion, then, does have Conception-2 style perfectly natural properties. Thus, I build (1) into Conception 2 of perfect naturalness.

3.3.2 Relative Naturalness According to Conception 2

In section 3.3.1 I discussed only perfect naturalness. We must now consider how to strengthen Conception 2 to apply to the *more natural than* relation.

First, I want to note an extreme method, which is actually a departure from the Conception 2 of the previous section, rather than an extension of it. We must shift to a different similarity notion, the relation of one property *making for more similarity* than another. Then we give up (6) in favor of characterizing *more natural than* directly: one property is more natural than another iff it makes for more similarity than the other.[32] (Weak) perfect naturalness is then defined as unexceeded naturalness.

The extreme version of Conception 2 differs starkly from Conception 1, and also from what I called "Conception 2" in the previous section. Ac-

[31] Principles (C2) and (A1) from the next chapter together imply that every object has an intrinsic profile.

[32] From his remarks on the upper end of naturalness in Armstrong (1989*b*, p. 24), Armstrong seems to conceive of primitive naturalness in this way.

cording to the new version of Conception 2, *only* intrinsic profiles will be perfectly natural, for intrinsic profiles make for the highest possible degree of similarity—duplication.[33] Properties like charge, spin, mass, etc.—which are perfectly natural according to (6) as well as Conception 1—will *not* be perfectly natural on this extreme version of Conception 2.

What I will henceforth call "Conception 2" is *not* the extreme view. Rather, it contains (6) as a characterization of perfect naturalness, and uses the section 3.2.3 notion of distance from a set to characterize relative naturalness. Thus, Conception 2 accepts principles (6) and (5a):

(6) A property or relation is perfectly natural iff it makes for similarity

(5a) For any set N of pairwise-equally natural perfectly natural properties and relations, and any properties or relations P and Q that supervene on N, $P \geq Q$ iff P is as least as close to N as Q is.

In the previous section, I showed how Conception 1 could incorporate a part of Conception 2 that dealt with perfect naturalness—principle (7). One might hope for a similar thing for relative naturalness. Specifically one might hope that the comparative similarity relation, which was employed in stating the extreme version of Conception 2, could constrain Conception 1's *more natural than* relation.

Unfortunately, I doubt that this is possible. The only potential principles of this kind seem to be the following:

(?) If P makes for more similarity than Q, then P is more natural than Q.

(??) If P is more natural than Q, then P makes for more similarity than Q

But each contradicts the claim of Conception 1 that conjunctions of perfectly natural properties are less natural than their conjuncts.

[33] It is crucial here that we are concerned only with *intrinsic* resemblance. If we allowed general qualitative resemblance (thereby allowing qualitative relational properties to affect resemblance), then the maximum degree of resemblance would not be duplication, but something stronger: what Lewis calls *indiscernibility.* See Lewis (1986*c*, p. 63).

3.4 Choosing a Conception of Naturalness

We have a primitive, the relation *more natural than*, and two main competing Conceptions of that primitive.[34] Let me summarize them.

On Conception, 1, 'more natural than' means 'more fundamental than'. The most natural properties are the most fundamental properties—at our world, the fundamental properties of physics. As well as being most fundamental, these properties make for similarity. Boolean and structural combinations of perfectly natural properties are not perfectly natural. For properties in worlds like the actual world, if one property requires a more lengthy or disjunctive definition from the perfectly natural properties than another, then it is less natural. In many worlds, the set of perfectly natural properties and relations is a Q-base for that worlds. But some worlds, like Onion, do not contain *any* perfectly natural properties, only a sequence of properties with ever-increasing naturalness. The set of perfectly natural properties at such a world is *not* a Q-base for it.

On Conception 2, 'perfectly natural' means 'makes for similarity'. Conception 2's perfectly natural properties include the perfectly natural properties of Conception 1, and more besides. While negations and disjunctions of perfectly natural properties are *not* perfectly natural, conjunctions and structural combinations of perfectly natural properties *are* perfectly natural.[35] As for other properties, they are more or less natural depending on how directly and non-disjunctively they may be defined in terms of the perfectly natural properties. At any world, the set of perfectly natural properties and relations is a Q-base for that world.

We also have a dark horse in the competition: the extreme version of Conception 2. On this version, 'at least as natural as' means 'makes for at least as much similarity'. So, the more intrinsically specific a property is, the more natural it is. The perfectly natural properties are the intrinsic profiles; properties like charge, mass, etc. lag far behind.

[34]It should be noted that by accepting a Conception, one does more than single out a notion of naturalness. The Conceptions relate naturalness to many different concepts, thus relating the latter concepts to each other. For example, by accepting the existence of the Conception 1 notion of perfect naturalness, one would thereby assert that the most fundamental properties make for similarity.

[35]Exceptions: as noted in section 3.3.1, there *may* be a case for some perfectly natural disjunctions of perfectly natural properties. And if a conjunction or structural combination of perfectly natural properties is the impossible property, then it is not perfectly natural.

One could accept *three* notions of naturalness, one for each Conception. But economy would have us choose. I choose Conception 1. What reasons could there be for favoring one Conception over another?

The less we understand a primitive, and the less we can illuminatingly characterize it, the worse it is to accept. Each Conception has deficiencies on this score. First, the notion of closeness to a set, which was used to characterize *more natural than* for Conceptions 1 and 2, is suspect. I said that disjunctiveness and length of definition made for greater distance. But how do these factors play off each other? Let N be the set of perfectly natural properties. Is a simple disjunction of two members of N closer to or farther from N than a more complicated structural combination of members of N? I don't know.

Conception 1 employs the notion of fundamentalness, and our grasp of this notion might be doubted. The paradigm cases for comparisons of fundamentalness come from the actual world: *redness* is less fundamental than *being a hydrogen atom* which is less fundamental than *having unit negative charge*. But we must abstract away from various features of the actual world to form the general concept of fundamentalness. For example, this sequence of increasingly fundamental properties from the actual world also seems to be one of increasingly "microphysical" properties, and I claimed in section 3.2.1 that microphysicality and fundamentalness are only contingently coincident. Moreover, the properties in the example are all physical. Can we be sure that we have a general concept of fundamentalness?

Likewise, Conception 2's notion of *making for similarity* seems questionable. First, the very notion of objective similarity is dark to some. Moreover, even though I do not suppose the all-or-nothing notion of *making for similarity* to be *defined* in terms of an underlying notion that comes in degrees, it still seems fair to ask: how much similarity must a property guarantee to make for similarity? For example, does *crimson* make for similarity? What about *redness?* What about still more broad color ranges?

Consider a world that we might call "Color". Color is somewhat like Onion in that it has a long sequence of properties that divide into other properties (but unlike Onion, this sequence comes to an end—see below.) In Color, properties (beyond a certain point of naturalness) are determinables of more specific determinates. (I call this world "Color" because the properties are like colors—determinables of determinates.) So, for example, P is a disjunction of $P_1 \ldots P_n$. But each of the P_i is a disjunction of Q_is, each of which are disjunctions…until finally we reach properties that are maximally specific: intrinsic profiles. We may stipulate that as we continue along this sequence, the properties make for more and more similarity.

To accept Conception 2, we must accept that at some point along the sequence, the properties begin to make for similarity. Unless by 'makes for similarity' we mean 'makes for *perfect* similarity', the cutoff point must be before the end of the sequence. But surely this is arbitrary. There is no privileged cutoff point before the endpoint. So the primitive notion of perfect naturalness, on Conception 2, is suspect.

If we do place the cutoff point at the only non-arbitrary place, the endpoint, then we will have collapsed Conception 2's notion of perfect naturalness into that of the extreme version of Conception 2 from section 3.3.2. For according to the extreme version of Conception 2, the perfectly natural properties are those that make for the maximum amount of similarity, namely, the intrinsic profiles. By 'Conception 2', I will continue to mean the more moderate version which allows perfectly natural properties that do not make for *perfect* similarity. This version is the one that is subject to the objection.

It might be replied that *making for similarity* is vague, but that this does not mean that the notion is ill-understood. We understand plenty of vague notions just fine.[36] I find this response objectionable for two main reasons. First, vagueness in the concept of naturalness might translate into vagueness of concepts that are analyzed in terms of naturalness, and some of these do *not* seem vague. For example, in section 4.2.1, I give an analysis of duplication in terms of naturalness. In chapter 9 I give an account of physical *distance* in terms of naturalness. David Lewis characterizes notions of lawhood, causation, and materialism in terms of naturalness in "New Work for a Theory of Universals".

Secondly, I find it misleading to call the dubiousness of *makes for similarity* "vagueness". This gives the impression that the only trouble is that we cannot find a *precise* cutoff point—we have a rough idea of how much similarity a property must ensure to count as making for similarity; we just can't say *exactly* how much. This isn't the right picture at all. We have a continuum of amounts of similarity ensured. At one end are properties that ensure perfect similarity. At the other end there are properties that ensure no

[36]I do not mean to suggest the existence of vague properties. When I say that a notion is vague, what I mean is that we have a phrase such that it is vague exactly which property is expressed by it. To me, the idea of real, non-linguistic "vagueness in the world" is unintelligible.

similarity. There are *no* distinguished halfway points, not even distinguished halfway *regions*. So, it seems to me that the "looseness" in Conception 2's perfect naturalness is severe, and thus it is unfit to take as primitive. We wouldn't understand the primitive.

How does this affect Conception 1? I related Conception 1 perfect naturalness to the all-or-nothing notion of similarity by accepting (7). Since I have argued that the notion of making for similarity is ill- understood, (7) can do its part of illuminating naturalness less well than we might have thought. Fortunately, (7) is not at the heart of Conception 1. The heart of that Conception is the notion of fundamentalness, which is unaffected by trouble with similarity.

Notice that the extreme version of Conception 2 is not affected by this problem, since it doesn't appeal to an all-or-nothing notion of similarity. All it appeals to is the relational notion *makes for at least as much similarity as*.

My first reason to prefer Conception 1 over Conception 2, then, is that Conception 2's notion of *making for similarity* is ill-understood and arbitrary seeming. I have a second reason for preferring Conception 1 that involves the nature of its primitive notion. Primitive naturalness on Conception 1 is a more "natural" and "unified" notion than on Conception 2. *Perfect* naturalness on Conception 2 involves similarity. But *relative* naturalness on that conception seems entirely different—it was illuminated by appeal to distance from N, the set of perfectly natural properties and relations. Distance from N seems to have nothing to do with similarity.

In contrast, I think there is a certain unity to the overall explanation of Conception 1 naturalness. Throughout I invoked the intuitive notion of *fundamentalness*. The most natural properties are the most fundamental properties, and a property is at least as natural as another when it is at least as fundamental. True, to explain relative naturalness I appealed to distance from a set. But this, it seems to me, *just was* an explanation in terms of relative fundamentalness.

The *extreme* version of Conception 2 has this same unity, since it explains both relative naturalness and perfect naturalness in terms of similarity. So again, the extreme version of Conception 2 fares better than its more moderate cousin.

After considering the nature of a primitive itself, we may consider the nature of the resulting theory. One consideration is the simplicity of that resulting theory. On the surface, the honors here go to Conception 2, because of the existence of Onion, the world of endless complexity. On Conception

1, there are no perfectly natural properties at Onion; this required jettisoning the simple principles (1) and (5). Moreover, it complicates other theory besides; see section 4.2, especially the definition of 'duplicate'.

However, I think that this is not a real advantage for Conception 2. Onion represents a genuine complication in the theory of naturalness, and Conception 2's way around it is cosmetic. As we saw, Conception 2 patches over the problem by its dubious all-or-nothing primitive *makes for similarity.*

The extreme version of Conception 2 affords a very simple definition of 'duplicate'. As we noticed in section 3.3.2, perfectly natural properties on this view will be, simply, intrinsic profiles. Objects are duplicates iff they share intrinsic profiles (see principle (C2) from section 4.2.3); moreover, each object has exactly one intrinsic profile (this, too, follows from (C2)—assuming as I do that necessarily coextensive properties are identical). Therefore, every object has exactly one perfectly natural property on this Conception, and duplicates are those objects that have the same perfectly natural property.

The final and most important consideration in deciding on a Conception is the strength of the resulting theory. Here, I have a programmatic reason for preferring Conception 1. I suspect that Conception 1 is more powerful than Conception 2, and more powerful than the extreme version of Conception 2. (This latter claim is important, since the extreme version of Conception 2 has, so far, high marks in this competition.)

Some of the tasks accomplishable by Conception 1 seem beyond the reach of both versions of Conception 2. For example, the very notion of a most fundamental property is intuitive. It is a natural and common view of physics that its goal is to seek out a complete description of the world in terms of the most fundamental properties. We have here a concept that should be acknowledged by ontology. Conception 1 acknowledges it (albeit by taking it as a primitive). But it is hard to see how Conception 2 can reconstruct this notion. According to Conception 2, the set of perfectly natural properties contains many properties in addition to the most fundamental properties. How could we sort out the most fundamental? And according to the extreme version of Conception 2, the perfectly natural properties are intrinsic profiles, which are not very fundamental at all. In each case, we are left without the notion of the most fundamental properties.

On the other hand, I suspect that Conception 1 can do all the work of Conception 2. The theoretical complications introduced by Onion, while irritating, seem surmountable (see section 4.2 in addition to principles (1a), (5a), and (5b) from this chapter). And I have hopes that *making for similarity* can be analyzed using Conception 1 naturalness, at least after the "looseness" in the notion is resolved. We have two clear sufficient conditions for *P*'s making for similarity, however the looseness is resolved: *P*'s being perfectly natural, and *P*'s being an intrinsic profile. We know that consistent conjunctions and structural combinations of properties that make for similarity themselves make for similarity. Here, however, the presence of Onion complicates matters. A full discussion of this issue is beyond the scope of this dissertation.

(A final aside: at this point we can make good on an earlier promise. In section 3.1 I said that naturalness is best thought of as coming in degrees. Now we can see why. Whether we accept Conception 1 or Conception 2, if we accepted only an all-or-nothing notion of naturalness, we would be committed to an unappealing arbitrary cutoff point. For Conception 1, it would mean accepting an arbitrary cutoff point of fundamentalness in the sequence of properties from Onion. For either the extreme or the moderate version of Conception 2, it would be an arbitrary cutoff point of amount of similarity ensured in the sequence of properties from Color.[37])

3.5 Conclusion

I have chosen Conception 1. But I must admit that my grasp of the notion of naturalness is uneasy. The project of clarifying naturalness was fraught with difficulty—I seemed to get as many negative results as positive. The possibility of Onion, the world of endless complexity, further complicated matters. In the next chapter I analyze *duplication*, *intrinsicality*, and other notions in terms of naturalness, but why not forget about naturalness altogether, and take one of those notions as a primitive instead? The leading candidate, I think, is duplication. We could go straight to duplication, rather than taking

[37] Another impetus for degrees of naturalness is that they seem to be required for one of the applications of naturalness: in determining the content of thought. See Lewis (1983*b*, pp. 370–7) and Lewis (1984, pp. 227–8). But the acceptability of Lewis's solution has been questioned by Phillip Bricker in conversation, at least given Conception 1 of naturalness. The worry is that naturalness conceived as fundamentalness might not match up with eligibility for being thought about. Perhaps a property that clearly is not expressed by one of our predicates is more fundamental than the intuitively correct candidate.

the long road through naturalness. In contrast to the notion of a natural property, the notion of duplication, to me at least, is perfectly clear.

This would be a bad idea, for we need naturalness. It does work that duplication and intrinsicality do not. First, some conception of naturalness seems to be required to analyze *making for similarity* (after the looseness in that notion has been eliminated). Any two objects will share an infinity of *intrinsic* properties; moreover, objects can share a property that makes for similarity without being duplicates. How to analyze similarity in terms of intrinsicality and duplication, then, is at least not obvious. Secondly, as I claimed in the previous section, the concept of "fundamentalness" is an intuitive one that must be acknowledged by ontology. Finally, Lewis's solution to the problem of the content of thought and language requires naturalness.[38]

Naturalness will not go away. But if it could be *analyzed* in terms of intrinsicality or duplication plus concepts of supervenience and the like, then we could have a clear primitive—duplication—and still reap all the benefits of naturalness. It would be nice, then, to have a principle of the form:

> *P* is more natural than *Q* iff _______

where the right side is stated in terms of duplication, supervenience, the part-whole relation, etc. However, I have found such a principle hard to come by.[39] So it seems that we'd better stick with naturalness as our primitive and make sense of it as best we can.

---

[38] See Lewis (1983*b*, pp. 370–7), and especially Lewis (1984). But see the previous note on the content of thought.

[39] For example, in section 8.1 I show how various attempts to define 'natural' in terms of 'supervenience' and 'qualitative' fail. These attempts fit the schema in the text since 'qualitative' is definable in terms of 'duplicate'—see section 4.2.2.

Chapter 4

Intrinsicality, Duplication, and other Notions

In this chapter I lay out a theory of intrinsicality, duplication, and other notions based on the notion of naturalness. In section 4.1 I investigate the claim that 'intrinsic' and 'duplicate' are interdefinable; it turns out that they are, given the intended interpretation of those terms. So, if we succeed in defining either we thereby define both. In sections 4.2.1 and 4.2.2, following the strategy of David Lewis, I use naturalness to analyze these two notions, as well as the notions of internal and external relations, intrinsic profiles, and qualitative properties and relations. In particular, it will be seen that the definition of 'duplicate' is not as straightforward as Lewis suggests. Finally, in section 4.2.3 I derive various consequences of the definitions.

4.1 Intrinsicality and Duplication Are "Interdefinable"

Lewis says there is a "circle of interdefinability" between the terms 'duplicate' and 'intrinsic' (1983*a*, p. 197). We may begin with the notion of a duplicate and go on to characterize intrinsic properties as follows:

(D1) P is *intrinsic* iff P can never differ between duplicates

Or, we may begin with intrinsicality, and define 'duplicate':

(D2) x and y are *duplicates* iff they have the same intrinsic properties

We have given two equivalences relating intrinsicality and duplication. To say that these notions are "interdefinable" via these two equivalences is to say that whichever of the equivalences we regard as a definition, the other may be derived as a consequence of that definition.

Suppose we begin with one of the equivalences as a definition. In what sense must we be able to "derive" the other? The material conditional whose antecedent is (D1) and whose consequent is (D2) is no theorem of formal logic; neither is its converse.

Background principles constraining the interpretation of 'intrinsic' and 'duplicate' must be used. I intend to investigate just what assumptions about the notions of duplication and intrinsicality must hold in order for 'intrinsic' and 'duplicate' to indeed be interdefinable. These assumptions must be weak in order for us to legitimately regard 'intrinsic' and 'duplicate' as being interdefinable via (D1) and (D2). In fact, I think that the assumptions must be *analytic* —true in virtue of the intended meanings of 'intrinsic' and 'duplicate'. Fortunately, the necessary assumptions are indeed analytic.

First some terminology. When x and y are duplicates, I will say "$\text{Dup}(x,y)$". When x and y share all properties in a set S (that is, when $\forall P \in S\,(Px \leftrightarrow Py)$) I will say that x and y are S-indiscernible ("$S\text{-Ind}(x,y)$"). Let I be the set of intrinsic properties.

Suppose we take 'duplicate' as undefined. We then regard principle (D1) as a definition of the term 'intrinsic'. The task is to derive principle (D2). It is trivial to show one direction of the equivalence: if x and y are duplicates, then it follows from (D1) directly that they have the same intrinsic properties.

The other direction is less trivial: if x and y share all intrinsic properties, then they are duplicates. The assumption we need is:

(A1) *duplication* is an equivalence relation

For suppose that $I\text{-Ind}(x,y)$, and consider the property *being a duplicate of* x; (i.e. $\lambda y(\text{Dup}(x,y))$).[1] By (A1), the *duplication* relation is reflexive, so x has this property. If we can show that this property is intrinsic, then since $I\text{-Ind}(x,y)$, y has the property as well, and hence $\text{Dup}(x,y)$.

The proof that *being a duplicate of* x is intrinsic uses (A1); specifically, the transitivity and symmetry of *duplication*. Recall that we are presently

[1] By 'λy($\text{Dup}(x,y)$)' I mean the property had by object y iff y is a duplicate of x, *not* the property had by object y iff y is a duplicate of the counterpart of x in y's world. See section 7.1.3 for more on this distinction.

regarding 'intrinsic' as being defined by (D1), so we must show that this property can never differ between duplicates. So suppose Dup(z, w); we show that z has the property *being a duplicate of* x iff w has this property. Suppose that z has the property; i.e. Dup(x, z). By transitivity, Dup(x, w), so w has the property *being a duplicate of* x as well. On the other hand, if w has the property, then Dup(x, w); by symmetry, Dup(w, x); by transitivity, Dup(z, x); by symmetry, Dup(x, z); hence, z has the property.

Assumption (A1), we have seen, is sufficient for deriving (D2) when we assume (D1). Indeed, it is necessary as well. (D2) trivially implies (A1), since the relation *having the same intrinsic properties* is an equivalence relation.

(A1) is not news. Anyone who understands the concept of duplication knows that it is an equivalence relation. So, it seems that when we take 'duplicate' as a primitive and define 'intrinsic' as in (D1), an analytic assumption governing the behavior of our primitive concept guarantees the desired relation between duplication and intrinsicality.

Now let us take 'intrinsic' as the primitive and regard (D2) as the definition of 'duplicate'. We must derive (D1) from (D2). As before, one direction is easy: if P is intrinsic then it follows trivially from (D2) (now taken to be a definition) that P can never differ between duplicates.

The other direction is less trivial. We must show that any property that can never differ between duplicates is intrinsic. The conjunction of the following two assumptions is necessary and sufficient for this derivation:

(A2) If property P is necessarily coextensive with an intrinsic property then P is intrinsic.

(A3) The set of intrinsic properties is closed under negation and (infinitary) conjunction[2]

Notice that (A2) is a direct consequence of an assumption from chapter 2: that necessarily coextensive properties are identical. Notice also that an equivalent way of stating (A3) is that I is identical to its Boolean closure.

We need more terminology. Let us write "NC(P, Q)" when $\forall x(Px \leftrightarrow Qx)$ (that is, when properties P and Q are necessarily coextensive). Let us write "$I(P)$" when P is intrinsic. For any set of properties S, let BC(S) be the

[2]Phillip Bricker pointed me in the direction of A3. Notice that (A3) makes necessary and impossible properties intrinsic.

Boolean closure of S—the smallest superset of S that is closed under negation and (infinitary) conjunction.

For brevity I use the concept of *strong supervenience*. For any sets A and B, A strongly supervenes on B ("sv(A,B)") iff $\forall x \forall y[B\text{-Ind}(x,y) \rightarrow A\text{-Ind}(x,y)]$. It is well known that:[3]

(K) if $\text{sv}(A, \text{BC}(B))$ then $\forall P \in A\, \exists Q \in \text{BC}(B)\ \text{NC}(P,Q)$

First, we need to show that (A2) and (A3) together are sufficient for the derivation. Let P be a property that never differs between duplicates; we must show that P is intrinsic:

1.	$\forall x \forall y[\text{Dup}(x,y) \rightarrow (Px \leftrightarrow Py)]$	Ass.
2.	$\forall x \forall y[I\text{-Ind}(x,y) \rightarrow (Px \leftrightarrow Py)]$	1, (D2)
3.	$\text{sv}(\{P\}, I)$	2
4.	$I = \text{BC}(I)$	(A3)
5.	$\text{sv}(\{P\}, \text{BC}(I))$	3, 4
6.	$\exists Q \in \text{BC}(I)\ \text{NC}(P,Q)$	5, (K)
7.	$\exists Q \in I\ \text{NC}(P,Q)$	4, 6
8.	$I(P)$	7, (A2)

Next, we show that (A2) and (A3) are individually necessary for deriving (D1), when we assume (D2). That is, we show that the conjunction of (D1) and (D2) implies each. First, (A2):

1.	$\text{NC}(P,Q)$	Ass.
2.	$I(Q)$	Ass.
3.	$\forall x \forall y[\text{Dup}(x,y) \rightarrow (Qx \leftrightarrow Qy)]$	2, (D1)
4.	$\forall x \forall y[\text{Dup}(x,y) \rightarrow (Px \leftrightarrow Py)]$	1,3
5.	$I(P)$	4, (D1)

[3] See Kim (1984, p. 49). This paper introduced the concept of strong supervenience. The formulation of strong supervenience I use is from Kim (1987, p. 317).

Next, (A3). We must show that $I = \text{BC}(I)$. Trivially, $I \subseteq \text{BC}(I)$. It remains to show $\text{BC}(I) \subseteq I$:

1.	$P \in \text{BC}(I)$	Ass.
2.	$\forall x \forall y[\text{BC}(I)\text{-Ind}(x,y) \rightarrow (Px \leftrightarrow Py)]$	1
3.	$\forall x \forall y[\text{BC}(I)\text{-Ind}(x,y) \leftrightarrow I\text{-Ind}(x,y)]$	Fact[4]
4.	$\forall x \forall y[I\text{-Ind}(x,y) \rightarrow (Px \leftrightarrow Py)]$	2, 3
5.	$\forall x \forall y[\text{Dup}(x,y) \rightarrow (Px \leftrightarrow Py)]$	4,(D2)
6.	$P \in I$	5, (D1)

Are these assumptions weak enough so as not to jeopardize the claim that 'intrinsic' and 'duplicate' are interdefinable via (D1) and (D2)? I think they are. Just as assumption (A1) was a consequence of the intended meaning of 'duplicate', assumptions (A2) and (A3) are consequences of the intended meaning of 'intrinsic': a property "had purely in virtue of its instances' intrinsic natures". Of course, there may be other conceptions of intrinsicality according to which (A2) and (A3) fail. On these conceptions, perhaps 'intrinsic' and 'duplicate' are not interdefinable.[5]

When we take 'duplicate' as a primitive and define 'intrinsic' via (D1), we can derive (D2) using (A1), a consequence of the intended meaning of 'duplicate'. When we take 'intrinsic' as a primitive and define 'duplicate' via (D2), we can derive (D1) using (A2) and (A3), consequences of the intended meaning of 'intrinsic'. Thus, the assertion that 'duplicate' and 'intrinsic' are interdefinable via (D1) and (D2) is justified.

4.2 Duplication and Beyond

In the present section I use naturalness to define various terms I use in this dissertation: 'duplicate', 'intrinsic' 'intrinsic profile', 'qualitative', etc. Some of these definitions have already been discussed. I will then note various consequences of the definitions.

[4]It is a fairly straightforward consequence of the definition of 'Boolean closure' that for any set S, S-Ind(x,y) iff BC(S)-Ind(x,y). See Paull and Sider (1992, Appendix, part one).

[5]See section 7.1.2.

4.2.1 Duplication

Following David Lewis, I begin by defining 'duplicate' in terms of 'natural'. But I do not accept Lewis's definition (1986*c*, p. 61):

> ...two things are *duplicates* iff (1) they have exactly the same perfectly natural properties, and (2) their parts can be put into correspondence in such a way that corresponding parts have exactly the same perfectly natural properties, and stand in the same perfectly natural relations.

I state Lewis's definition as follows. For simplicity, in this chapter I count properties as 1-place relations and take a one place sequence $\langle x \rangle$ to be simply x.

(DP1) x and y are *duplicates* $=_{\mathrm{df}}$ there is a one-one correspondence f between x's parts and y's parts such that for any perfectly natural relation R and for any sequence $\vec{x}$ of parts of x, $R(\vec{x})$ iff $R(f(\vec{x}))$

In what follows I would like to discuss various aspects of (DP1). First, I would prefer to add the condition that the correspondence f be a *part-whole isomorphism* (see section 3.2.1). This requirement would be redundant if the part-whole relation is perfectly natural, but since I am not sure of this, I think it should be built in explicitly.

Secondly, Lewis suggests that the definition might be simplified if we allowed that *structural* combinations of perfectly natural properties are perfectly natural. The simplified definition says that x and y are duplicates iff they have exactly the same perfectly natural properties (Lewis, 1986*c*, pp. 61–62). Without structural perfectly natural properties, we cannot simply require that duplicates share all perfectly natural properties. For example, if the perfectly natural properties in the actual world are properties involving charge, spin, charm, etc., then actual macroscopic objects will not have *any* perfectly natural properties. Any two macroscopic objects would vacuously be duplicates. But if there are structural perfectly natural properties, then differences in the perfectly natural properties and relations of the *parts* of macroscopic objects will translate into differences in the structural perfectly natural properties of those macroscopic objects. Hence, the possibility of simplification of (DP1). However, in 3.2.1 I argued that structural combinations of perfectly natural properties are not perfectly natural, and so I do not simplify the definition of 'duplicate'.

In fact, considerations raised in section 3.2.2 by Onion, the world of endless complexity, show that (DP1) is too simple. Consider any two objects from Onion whose parts stand in the same perfectly natural relations (that is, spatiotemporal relations, as these were stipulated to be the only perfectly natural relations at Onion). Recall that no perfectly natural properties are instantiated at Onion, so no part of either object has any perfectly natural properties. According to (DP1), it follows that these objects are duplicates. But surely they need not be. As we divide either of the objects up into smaller and smaller parts, there will be a sequence of increasingly natural properties had by these parts. The two objects could have entirely different sequences.

(DP1) implies that any two "spatiotemporal isomorphs" from Onion are duplicates, so it must be modified. For x and y to be duplicates we ought to require more than that the parts of x and y share perfectly natural properties and relations. Objects from Onion have properties that come in a never-ending sequence of naturalness. We ought to require of duplicates from Onion that, roughly, at some point along this sequence their parts begin to share all properties and relations, and continue to do so for the rest of the sequence.

However, it is not a trivial matter to make this intuition precise. We might think to proceed as follows. Let the variables 'R', 'R'', and 'R''' range over *all* relations—including 1-place relations as per our convention for this chapter.

(DP2?) x and y are *duplicates* $=_{\text{df}}$ there is a part-whole isomorphism f between the set of x's parts and the set of y's parts such that $\forall R(\exists R' \geq R)(\forall R'' \geq R')$, for every sequence $\vec{x}$ of parts of x, $R''(\vec{x})$ iff $R''(f(\vec{x}))]$

To get an idea of how (DP2?) is supposed to work, consider first how (DP2?) applies to a simple case, a case that the original definition (DP1) handles correctly. Let x be composed of two mereological atoms x_1 and x_2; similarly, let y be the fusion of two atoms y_1 and y_2. Suppose that neither x nor y has any perfectly natural properties. Suppose further that x_1 and x_2 stand in perfectly natural relation R, as do y_1 and y_2; also suppose that all four atoms have perfectly natural property P. Suppose further that none of the four atoms have any other perfectly natural properties, nor do they stand in any other perfectly natural relations. Finally, suppose that none of these objects have any Onion-like properties or relations that divide endlessly into increasingly natural properties and relations.

In this case, (DP2?) has the intuitively correct result: that x and y are duplicates. For let f be the following part-whole isomorphism between the set of x's parts and the set of y's parts:

$$\begin{aligned} f(x_1) &= y_1 \\ f(x_2) &= y_2 \\ f(x) &= y \end{aligned}$$

It is clear that the right hand side of (DP2?) holds for this choice of f. The parts of x and y have the same perfectly natural properties and stand in the same perfectly natural relations, under this function f. So, if the first variable 'R' on the right side of (DP2?) is assigned any property/relation Q such that some perfectly natural property or relation Q' is at least as natural as it, then we satisfy the rest of the right hand side by assigning Q' to the existentially quantified variable 'R''. On the other hand, if 'R' is assigned some property or relation Q such that *no* perfectly natural property is at least as natural as it, then 'R'' may simply be assigned Q. Why? Because in this case, Q, and everything as or more natural than Q, will be Onion-like properties or relations that decompose endlessly into increasingly natural properties and relations, and, by stipulation, the objects in question instantiate no such properties or relations.

Secondly, consider in sketch how (DP2?) applies to two Onion-like objects x and y. These objects will have no perfectly natural properties, nor will their parts. But intuitively, if there is some degree of naturalness such that the parts of x and y share all properties and relations of that degree of naturalness or greater, then x and y turn out duplicates. More carefully, the sufficient condition for duplication is that for *any* property or relation R, there is some property or relation at least as natural as R such that every property or relation as least as natural as it is shared by the parts of x and y. On the other hand, speaking intuitively again, if no matter how high a degree of naturalness we choose, the parts of x and y still differ with respect to some properties at least that natural, then x and y do not turn out duplicates.

However, I doubt that (DP2?) gives an acceptable sufficient condition for duplication. Suppose that x and y are *not* duplicates, and suppose for simplicity that each is a mereological atom. Suppose further that the only intrinsic differences between x and y are those entailed by the fact that x has, while y lacks, a certain intrinsic property P. Now suppose that P is the conjunction of two properties Q_1 and Q_2, and that both x and y have Q_1, a perfectly natural property. It follows that x has, while y lacks, Q_2. If Q_2 is not as or more natural than Q_1, then what we have said is consistent with the right hand side of (DP2) being satisfied, despite the stipulation that x and y are not duplicates, for when we assign P to the variable 'R', we may assign Q_1 to the variable 'R''.

Of course, if the right hand side of (DP2) is true, then it applies to Q_2. But can we not satisfy the right hand side of (DP2) in exactly the same way as above? Let Q_2 be the conjunction of two properties Q_3 and Q_4, where Q_3 is a perfectly natural property that both x and y have (it follows that x has, while y lacks, Q_4). When Q_2 is assigned to 'R', we may satisfy the right hand side of (DP2) by assigning Q_3 to 'R''. The process then will be repeated for Q_4, for one of Q_4's conjuncts, for one of those conjuncts…forever. The following diagram pictures these endlessly conjunctive properties:

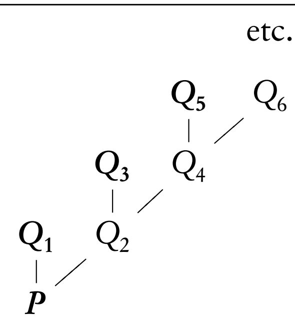

- **boldface** indicates perfectly natural properties
- lines (—) move upwards from properties to more natural conjuncts

Figure 4.1: Endlessly conjunctive properties

So, we have a pair of non-duplicates x and y that, it seems, satisfy the right hand side of (DP2?).

This example is coherent only if we can claim both that $\forall Q' \geq Q_1\, Q'x \wedge Q'y$, and that y lacks Q_2, Q_4, Q_6, etc., and we can make these claims only if we claim that none of Q_2, Q_4, Q_6…are equally as natural as Q_1 (clearly, none can be more natural since Q_1 is perfectly natural). I think that we *can* make such a claim. In fact, we could make it in two ways. Q_2, Q_4, Q_6, etc. could be each be *incomparable* in respect of naturalness to Q_1, or each of these properties could be *less* natural than Q_1. I have no argument that either case is possible, but neither do I have an argument that both cases are impossible. Thus, the acceptability of (DP2?) is at least cast into doubt.

Think of sequences of increasingly natural properties as paths heading upward (the upward direction is the direction of increasing naturalness). Intuitively, (DP2?) says that if above every property or relation there is *some* path along which x and y agree on all properties and relations, then x and y must be duplicates. We have seen that this may be too lenient. The example seems to show that we should make some requirement about *every* path.

This may be accomplished with the following set of definitions:

> C is a *chain* $=_{\text{df}}$ C is a sequence[6] of properties such that:
>
> a) every member of C is more natural than every previous member of C (that is, naturalness of members of C increases monotonically), and
> b) there is no Q that is not a member of C, but is more natural than every member of C

Notice that every chain is either infinite, or has a perfectly natural property or relation as its final member (or both). Notice also that an upper segment of a chain is itself a chain.

Let C be a chain, A and B be sets of objects, and f be a part-whole isomorphism between A and B. We define the notion of A and B *agreeing on chain C under function f* (for short: "agree(A,B,C,f)"):

> *agree*(A,B,C,f) $=_{\text{df}}$ for every property or relation R from C and every sequence $\vec{x}$ of members of A, $R(\vec{x})$ iff $R(f(\vec{x}))$

Another preliminary definition. Let A and B be sets of objects:

> f is a *duplication isomorphism between A and B* $=_{\text{df}}$ f is a part-whole isomorphism between A and B such that every chain has some upper segment or other C such that agree(A,B,C,f)

[6] I mean to allow α-sequences for any ordinal α, with the exception that I do not allow a null sequence.

We are finally in a position to give our definition of 'duplicate':

(DP2) x and y are *duplicates* $=_{\text{df}}$ there is a duplication isomorphism between the set of x's parts and the set of y's parts

The idea behind (DP2) is similar to the idea behind (DP2?), so the motivating remarks for (DP2?) apply to (DP2). However, the revised definition avoids the problem with (DP2?). Let X be the set of x's parts, and Y be the set of y's parts; the chain $\langle Q_2, Q_4, Q_6 \ldots \rangle$ has no upper segment C such that agree(X, Y, C, f), since x and y (in that example) differ with respect to every property in this chain. So there cannot be a duplication isomorphism between X and Y.

4.2.2 Other Concepts

In this section I state definitions of 'intrinsic', 'internal', 'external', 'intrinsic profile', and 'qualitative'. For intrinsicality and internal and external relations, I follow Lewis closely. Intrinsicality may be defined in terms of duplication via (D1):

(D1) P is *intrinsic* $=_{\text{df}}$ P can never differ between duplicates

Lewis defines an *internal* relation as one that "supervenes on the intrinsic nature of the relata" (1986*c*, p. 62). This may be restated as follows:

(D3) A relation R (with two or greater places) is *internal* $=_{\text{df}}$ for any sequence of objects $\vec{x}$ such that $R(\vec{x})$, if f is a function such that for any x in $\vec{x}$, x is a duplicate of $f(x)$, then $R(f(\vec{x}))$

He then defines an external relation to be a non-internal relation R that "supervenes on the intrinsic nature of the composite of the *relata* taken together…" (1986*c*, p. 62). We may elaborate on this as follows. An example of an external relation might be the relation *being ten feet from*. The idea is that if x is ten feet from y, then any duplicate of the *fusion* of x and y will need to contain two parts, a duplicate of x and a duplicate of y, separated by ten feet. Here is the definition:

(D4) R is *external* $=_{\text{df}}$ i) R is a relation (with two or more places) that is not internal and ii) for any two objects x and y, any duplication isomorphism f between the set of x's parts and the set of y's parts, and any sequence $\vec{x}$ of parts of x, $R(\vec{x})$ iff $R(f(\vec{x}))$.

It may not be evident that (D4) captures the intuitive notion of an external relation expressed above. However, (D4) is equivalent[7] to:

(D4′) Relation R is *external* $=_{\text{df}}$ i) R is a relation (with two or more places) that is not internal, and ii) for any sequence of objects $\vec{x}$ such that $R(\vec{x})$, and for any object y that is a duplicate of x, the fusion of the members of $\vec{x}$, and for any duplication isomorphism f between the set of parts of x and the set of parts of y, $R(f(\vec{x}))$

(D4′) seems closer to the intuitive characterization of externality: it says that a non-internal relation is external if, when it holds between some objects, it holds between "corresponding" parts of any duplicate, y, of the fusion, x, of those objects (the "corresponding" parts here are the images of the original objects under any duplication isomorphism between the parts of x and the parts of y). I work with (D4) rather than (D4') because it is simpler.

Next we turn to the concept of an intrinsic profile, a maximally specific intrinsic property. This concept may be analyzed as follows:

(D5) P is an *intrinsic profile* $=_{\text{df}}$ P is an intrinsic property such that for every intrinsic property Q, either P entails Q or P entails $\sim Q$

Finally, we give a definition of 'qualitative'. Qualitative properties and relations are "non-haecceitistic" properties. Intrinsic properties should turn out to be qualitative, but so should non-intrinsic properties such as *being ten feet from something red*. Roughly, a relation is qualitative iff one can determine whether or not it holds between some objects whenever one knows all the intrinsic facts about the possible world containing those objects.

Lewis says that (1986*c*, p. 63):

[7]Proof: We need a lemma, which follows immediately from the relevant definitions:

Lemma if f is a duplication isomorphism between A and B, C⊆A, and $g = f$ restricted to C, then g is a duplication isomorphism between C and $f[C]$.

That (D4) entails (D4′) is evident. For the other direction, suppose R is non-internal, that f is a duplication isomorphism between the set of x's parts, X, and the set of y's parts, Y, that $\vec{x}$ is a sequence of parts of x, and that $R(\vec{x})$. Let x' be the fusion of the members of $\vec{x}$, let X' be the set of parts of x', let $y' = f(x')$, let Y' be the set of parts of y', let g be f restricted to X'. By mereology, $X' \subseteq X$, so by the Lemma, g is a duplication isomorphism between X' and $f[X']$. Since f is a part- whole isomorphism, $f[X'] = Y'$, so by (DP2), x' and y' are duplicates. Hence, by (D4′) $R(g(\vec{x}))$. But $g(\vec{x}) = f(\vec{x})$, so $R(f(\vec{x}))$. Without loss of generality, (D4) follows.

> Extrinsic qualitative character, wherein duplicates may differ, consists of extrinsic properties that are, though not perfectly natural, still somewhat natural in virtue of their definability from perfectly natural properties and relations.

I would alter this definition in two ways. First, let us move from "definability" to (global) supervenience, if for no other reason than to make it explicit that the kind of definability in question is infinitary.[8] Thus, as a first approximation, the qualitative properties and relations are those that supervene (globally) on the set of perfectly natural properties and relations. Unfortunately, this leaves out some qualitative properties and relations because of the complications introduced by Onion. Fortunately, we can avoid this problem by defining the qualitative properties to be those properties that supervene on the set of *intrinsic* properties and *external* relations. We thereby make use of the machinery in (DP2) to get around the Onion problem, since the definitions of 'intrinsic' and 'external' make use of duplication. Thus, I accept:

(D6) Property or relation R is *qualitative* $=_{\text{df}}$ R supervenes on the set of intrinsic properties and external relations

4.2.3 Consequences of the definitions

The consequences of our definitions are an important test of their adequacy. These consequences show that the definitions preserve the features we expect the *definiens* to have, prior to analysis.

First, I note that principle (0) from section 3.1:

[8] Although Lewis is not explicit on this point, it is clear that by 'definability' Lewis means to allow infinitary definition. First, he uses the word 'indiscernible' in Lewis (1986*c*, p. 63) for things that have exactly the same qualitative properties: surely one wouldn't call two things "indiscernible" if they differed with respect to some property that is capable of being defined from the perfectly natural properties and relations, even if the definition must be infinite. Secondly, in Lewis (1986*c*, pp. 62–63), Lewis says that two objects have the same '*intrinsic* qualitative character' iff they are duplicates. That is, the intrinsic qualitative properties are exactly the intrinsic properties. However, *intrinsic* properties needn't be finitely definable from the perfectly natural properties and relations: all that is required of intrinsic properties is that they never differ between duplicates. Infinite Boolean combinations of intrinsic properties, for example, are intrinsic (see (A3) of this chapter). So, *intrinsic* qualitative properties needn't be finitely definable from the perfectly natural properties and relations; hence, it would be odd if Lewis intended such a restriction on the *extrinsic* qualitative properties.

(0) Every perfectly natural property is intrinsic

is a consequence of our definitions.[9] Suppose property P is perfectly natural, and let x and y be any two duplicates. By (DP2), there is a duplication isomorphism, f, between the set X of x's parts and the set Y of y's parts. So every chain has an upper segment C such that X and Y agree on C under f. But $\langle P \rangle$ is a chain, and has just one upper segment: itself. So X and Y agree on C under f. But since every object is a part of itself, $x \in X$ and $y \in Y$; since f is a part-whole isomorphism, $f(x) = y$. Hence, x has P iff y has P. Since P cannot differ between duplicates, then, it is intrinsic by (D1).

A similar result holds for perfectly natural *relations*:

(C1) every perfectly natural relation is either external or internal

Let R be any perfectly natural relation. We will prove that R satisfies clause ii) of (D4)—thus, R is either internal or external. Suppose that f is a duplication isomorphism between X, the set of parts of x, and Y, the set of parts of y, and that $\vec{x}$ is a sequence of parts of x such that $R(\vec{x})$. As in the proof of (0), $\langle R \rangle$ is a chain with only one upper segment, so X and Y agree on $\langle R \rangle$ under f. Therefore, $R(f(\vec{x}))$. Our conclusion follows without loss of generality.

Next, I show that (A1), (A2), and (A3), the conditions from section 4.1 under which 'duplicate' and 'intrinsic' are interdefinable via (D1) and (D2), follow from our definitions.

First, we prove (A1): *duplication* is an equivalence relation. Reflexivity and symmetry are immediate; transitivity remains. Suppose x and y are duplicates as are y and z. By (DP2), there is a duplication isomorphism f between the set X of x's parts and the set Y of y's parts, and a duplication isomorphism g between Y and the set Z of z's parts. Now, $f \circ g$ is clearly a part-whole isomorphism between X and Z. We will show that every chain has some upper segment C such that agree$(X, Z, C, f \circ g)$; it then follows by (DP2) that x and z are duplicates. Let C' be an arbitrary chain. In virtue of f's existence, C has an upper segment C'' such that agree(X, Y, C'', f); similarly, C has an upper segment C''' such that agree(Y, Z, C''', g). If $C'' = C'''$ then let $C = C'' = C'''$; if $C'' \neq C'''$ then one is an upper segment of the other; in this case let C be the smaller of C'' and C'''. Either way, C is an upper segment of C'. Since agree(X, Y, C'', f) and agree(Y, Z, C''', g), we have agree$(X, Z, C, f \circ g)$. Q.E.D..

[9] Another consequence is a more general version of (0): every chain has an upper segment in which all properties are intrinsic.

Next, (A2): anything necessarily coextensive with an intrinsic property is intrinsic. If P is intrinsic, then it can never differ between duplicates, by (D1). But then any property necessarily coextensive with P can never differ between duplicates, and hence is intrinsic by (D1). Notice that I do not rely on my chapter 2 assumption that necessarily coextensive properties are identical.

Lastly, (A3): I, the set of intrinsic properties, is closed under negation and infinitary conjunction. Let $P \in I$. Since P can never differ between duplicates, neither can $\sim P$; hence $\sim P \in I$. Now let the members of A be in I, and let Q be the conjunction of the members of A. Q can never differ between duplicates, and hence $Q \in I$. For suppose otherwise: let Qx, $\sim Qy$, where x and y are duplicates. Since Q is the conjunction of the members of A, x has every member of A, whereas y does not have every member of A. Thus, some intrinsic property—a member of A—differs between duplicates. *Contradiction.*

Since (A1) is a consequence of the definitions, so is (D2), the principle that objects that share intrinsic properties are duplicates. For we showed in section 4.1 that (D1) and (A1) together entail (D2).

Next we prove two principles for intrinsic profiles. First:

(C2) P is an intrinsic profile iff the set of P's instances is a maximal set of possible duplicates

where by a "maximal set of possible duplicates" I mean a set S such that i) every member of S is a duplicate of every other member of S, and ii) S is closed under the relation *duplicate of* (remember the possibilist quantifiers!).

First suppose that the set S of P's instances is a maximal set of possible duplicates. P can therefore never differ between duplicates since if x and y are duplicates then either both or neither are in S. Moreover, let P' be any intrinsic property, and let S' be the set of P''s possible instances. Either S overlaps S' or it does not. If the latter, then P entails $\sim P'$. If the former, then $S \subseteq S'$ and hence P entails P'. For suppose otherwise: let $x \in S, x \in S', y \in S, y \notin S'$. Since $x, y \in S$, x and y are duplicates. But since $x \in S'$ but $y \notin S'$, P' differs between duplicates. But P' is intrinsic. *Contradiction.* Thus, P is intrinsic and for every intrinsic property P', P entails either P' or $\sim P'$. P, therefore, is an intrinsic profile.

For the other direction, suppose that P is an intrinsic profile and let S be the set of P's (possible) instances. Let $x, y \in S$ and let P' be any intrinsic property had by either x or y. Since x and y each has P, an intrinsic profile, the other must have P' as well. So, x and y have the same intrinsic properties, and hence are duplicates by (D2). Next, note that if $x \in S$, any duplicate of x must be in S—otherwise P, an intrinsic property, would differ between duplicates. Thus, S is a maximal set of duplicates.

Next, we prove (C3):

(C3) For any object x, the conjunction of x's intrinsic properties is an intrinsic profile

Let P be the conjunction of A, the set of x's intrinsic properties. By (A3), P is intrinsic. Let Q be any intrinsic property—we show that either P entails Q or P entails $\sim Q$; it follows that P is an intrinsic profile. Either x has Q or x has $\sim Q$. If the former then $Q \in A$ and hence P entails Q. If the latter, then since $\sim Q$ is intrinsic (by (A3)), $\sim Q \in A$, so P entails $\sim Q$.

Finally, we note that the following principles are consequences of the definition of 'qualitative':

(C4) every intrinsic property is qualitative

(C5) every internal or external relation is qualitative

It is trivial that all intrinsic properties and external relations are qualitative, since principles (S1) and (S4) from chapter 3 together entail that a property or relation supervenes on any set containing it. Internal relations remain. Let $I =$ the set of intrinsic properties; let $E =$ the set of external relations, and let R be an internal relation; we show that R supervenes on $I \cup E$. Let f be an $I \cup E$-isomorphism between $D(w)$ and $D(w')$, the domains of possible worlds w and w'. Suppose $R(\vec{x})$, where $\vec{x}$ is a sequence of members of $D(w)$. Since f is an $I \cup E$-isomorphism it is an I-isomorphism, so for any x in $\vec{x}$, x shares all intrinsic properties with $f(x)$. Thus, by (D2), for any x in $\vec{x}$, x is a duplicate of $f(x)$. Since R is internal, it follows that $R(f(\vec{x}))$.

We have shown that our definitions entail many principles that seem intuitively true of our notions. This is evidence that the definitions are correct.

Chapter 5

Naturalness and Arbitrariness

I claim that we should admit the notion of naturalness. Naturalness comes in degrees, but in this chapter I consider only perfect naturalness, which I call simply "naturalness".

David Lewis tentatively proposes taking naturalness as a primitive (1986*c*, pp. 63–4; 1983*b*, p. 347). I have discussed and accepted this proposal in chapter 3. But Lewis proposes the conjunction of this claim with his "class nominalism", according to which relations are classes of ordered pairs of possibilia. I will argue that this combination of views is defective.

After introducing Lewis's proposal in section 5.1, I argue in section 5.2 that we should reject it because of issues involving the nature of *ordered pairs* (and ordered classes in general). Then, in section 5.3, I briefly examine a puzzle due to Kripke and Wittgenstein. I argue that a proposed solution of the puzzle using the natural/nonnatural distinction fails, at least if numbers are viewed, as they often are, as arbitrarily constructible entities.

5.1 Primitive Class Naturalism

Class nominalism is a reductive theory of the abundant properties and relations. The basic entities countenanced by the class nominalist are classes and concrete particulars ("objects"). Properties and relations are then constructed out of these entities. The properties are identified with the *classes* of objects. The n-place relations are identified with the classes of ordered n-tuples of objects, where these n-tuples are constructed from classes according to any one of the various common devices. An n-tuple *instantiates* an n-place relation R iff it is a member of R. An object instantiates a property P iff it is a member of P. For example, the property *redness* is simply the class of all red things; the relation *being ten feet from* is the class of all ordered pairs $\langle x, y \rangle$ where x is an object that is ten feet from y, another object. For a class nominalist, then, properties and relations are not primitive entities—they are constructed from primitive entities using the methods of the set theorists.

Class nominalism is most plausible it is coupled with *possibilism*: the acceptance of non-actual possible objects. Otherwise, properties are identified with the classes of their *actual* instances; thus, distinct properties that just happen to be coextensive—for example, the properties of *being George Bush* and *being the President of the United States in 1992*—turn out identical. I will, therefore, address my remarks to this version of class nominalism.

If every class of possibilia counts as a property, then the natural properties are an exceedingly small minority of the properties. Every possible object has infinitely many properties, since there are (presumably) infinitely many possible objects and hence infinitely many sets of possible objects containing a given possible object. In contrast, on the conception of naturalness that I accept ("Conception 1" from chapter 3), macroscopic actual objects, for example, do not have *any* (perfectly) natural properties—only their subatomic parts do. And it is plausible that even these have only a few natural properties: charge, spin, mass, etc.

We have a question: how do we account for naturalness in ontology? I have discussed some of these options in chapter 1; the one I want to focus on now is the option of **primitive naturalism**. The primitive naturalist takes the distinction between naturalness and unnaturalness as a primitive distinction that is incapable of reductive analysis. What is more, the distinction is in no way a matter of convention—it is an objective fact about the world that some properties and relations are natural, while others are not.

Lewis holds class nominalism and entertains primitive naturalism.[1] I call the conjunction of these two views "primitive class naturalism", or "PCN" for short. Thus, PCN is the view that i) reduces properties and relations to classes of objects and 'tuples of objects, and ii) takes the natural/non-natural distinction as an objective primitive distinction among properties and relations. I will argue in the next section that PCN is an unacceptable

[1]See Lewis (1986*c*, section 1.5). Anthony Quinton defends a related view in Quinton (1958).

combination of views.

5.2 Against Primitive Class Naturalism

5.2.1 Pairs and Relations

According to PCN, a binary relation is a class of ordered pairs. For example, the relation *being ten feet from* is the class of all ordered pairs $\langle x, y \rangle$ of possible objects where x is ten feet from y. But what are ordered pairs?

Ordinary classes are unordered. When we write the name of a class thus: "$\{x, y\}$", the order in which we write the members is insignificant, for the class $\{x, y\}$ is identical to the class $\{y, x\}$. Classes are individuated solely by their membership. For this reason, we do not identify, for example, the relation *taller than* with the class of all classes $\{x, y\}$ where x is taller than y. Presumably, we would go on to claim that the relation *shorter than* is the class of all classes $\{x, y\}$ where x is shorter than y. But since for any x and y, $\{x, y\} = \{y, x\}$, these two classes would be *identical*, and hence we would have identified two distinct relations: *taller than* and *shorter than*.

In contrast, ordered pairs are ordered. When we write the name of an ordered pair thus: "$\langle x, y \rangle$", the order is significant, for $\langle x, y \rangle$ is *not* identical to $\langle y, x \rangle$. Ordered pairs are individuated by their membership *and* the relative order of those members. That is, ordered pairs obey the following identity condition:

$$\langle x, y \rangle = \langle z, w \rangle \text{ if and only if } x = z \text{ and } y = w.$$

But the notion of an ordered pair is *not* an undefined notion of typical class theories. Ordered pairs are *constructed* from classes by any one of a number of methods, each of which preserves the identity condition just mentioned. To reap the benefits of having ordered pairs in our ontology, all we need is the concepts of standard class theory and a little ingenuity.

One method for constructing ordered pairs ("method" for short) was introduced by Wiener. For any x and y, Wiener identifies $\langle x, y \rangle$ with the class: $\{x, \{y, \varnothing\}\}$. Given the class $\{x, \{y, \varnothing\}\}$, we can "recover" the information of which is the first member of the pair and which is the second. We know, for example, that the second member is y since y is the element paired with the null-class. The more common method is due to Kuratowski, who identifies $\langle x, y \rangle$ with $\{x, \{x, y\}\}$. But there are countless others—any method that yields one ordered pair for any two objects and obeys the stated identity condition will do.

These methods should not be viewed as conflicting theories of the nature of "the ordered pairs" conceived as a sort of entity with which we have prior acquaintance. Rather, they are proposals for using certain classes to do the work we require of "ordered pairs".

Or rather, most of the work. Some of the work asked of the ordered pairs cannot be done by their stand-ins. Constructions fail at certain tasks because they are constructions. Since each method generates the appropriate ordered pairs equally well, no one method generates classes that deserve the title of "*the* ordered pairs" any more than any other method. True, some methods are more common in mathematical and philosophical writing while others are not even used at all, but these differences among methods have no ontological significance, in a sense to be discussed below. I will argue that this fact means trouble for PCN.

5.2.2 Benacceraf's Argument

First, however, let us rehearse informally a famous argument due to Paul Benacerraf (1965). The multiplicity of methods has an analog in number theory: the familiar reduction of numbers to classes can be carried out in many ways. We might let zero be the null class, and let $n+1$ be the unit class of the class identified with n. This is the strategy of Zermelo. Or, we might let the null class again be zero, but let $n+1$ be the union of n and $\{n\}$, as did von Neumann. Or, we could employ any one of countless other strategies.[2]

Each such strategy provides adequate surrogates for the numbers. Each can mathematically do the job, in the sense that the resulting "numbers" provably have class-theoretic properties that are exactly isomorphic to all the number-theoretic properties that numbers are supposed to have. The strategies may differ one from another in terms of convenience, but they seem to be "ontologically on a par". No strategy generates "the real and true" numbers.

What this means, argues Benacceraf, is that we cannot think of the numbers as a single determinate class of objects, fixed once and for all. Which classes are "*the* numbers"? Is the number 2, for example, $\{\{\varnothing\}\}$, as Zermelo

[2]See Quine (1963, pp. 81–5). See Benacerraf (1965) for an account of the desiderata a method for constructing the natural numbers must satisfy.

would have us believe, or is it $\{\varnothing, \{\varnothing\}\}$ as it is for von Neumann? It cannot be both, for $\{\{\varnothing\}\} \neq \{\varnothing, \{\varnothing\}\}$. But surely it cannot just be one of them, since it would be arbitrary to grant one strategy the privileged status of generating the true numbers. So neither can be *the* number 2. In Benacceraf's words, this means we should not view the numbers "as objects".

That is, if we wish numbers to be reduced to other entities. As Linda Wetzel and Michael Resnik have pointed out, if we believe in numbers as objects in their own right, in no need of reduction, then we escape Benacceraf's argument (Wetzel (1989, part one); Resnik (1980, p. 231)). But many shy away from such a bold response.

This multiplicity of constructions of the natural numbers, argues Benacceraf, does not matter for most contexts in which we use concepts of number. But how can this be, if numbers are "not objects"?

An attractive answer is given by the "structuralist". Take the natural numbers, for example, and ask yourself what is distinctive about them. The structuralist answer is: structure alone. What is distinctive about the natural numbers is that there is an initial element (0), and a relation called the "successor relation" such that each number has a unique successor, the initial element is the successor of no element, etc. But exactly what the elements are does not matter, so long as there is enough of them, and there is a relation among them to play the role of the successor relation. A line of people that continued forever, with one person first in line, could be the numbers. Any "ω-sequence"—any one- one function from ω, the class of finite von Neumman ordinals, onto *any* domain has the appropriate structure. The number to which $\varnothing$ is mapped plays the role of 0; the number to which $\{\varnothing\}$ is mapped plays the role of 1, etc. The image of the ordinal successor relation plays the role of number- theoretic *successor*; the image of $\subseteq$ plays the role of $=$, etc.

Any ω-sequence, then, is an adequate surrogate for the sequence of natural numbers. A number-theoretic sentence ϕ, for the structuralist, will be interpreted as a universally quantified claim:

for any ω-sequence f, $\phi(f)$

where $\phi(f)$ is the result of replacing the number-theoretic predicates, constants, and functors in ϕ with appropriate constructions in terms of f. The quantifiers in ϕ that supposedly range over numbers are in ϕ(f) allowed to range over the range of f. Since each method faithfully preserves number-theoretic properties, if ϕ is an "ordinary" number-theoretic sentence then the truth-value of $\ulcorner \forall f \phi(f) \urcorner$ will always match the truth value we would have intuitively assigned to ϕ. And this, despite the fact that we quantify over no numbers.

By an "ordinary" number theoretic sentence, I mean roughly a sentence that concerns only the *structural* properties of numbers. For non-ordinary sentences, structuralism gives somewhat odd results. Consider:

(1) Ted Sider is distinct from the number 2

Structuralism interprets (1) as the claim that for any ω-sequence, the third member of the sequence is not identical to me (the first member is identified with zero, remember). But this claim is not true, since I am the third member of many ω-sequences (consider: $\varnothing, \{\varnothing\}, \text{Ted}, \{\{\varnothing\}\}, \{\{\{\varnothing\}\}\}, \dots$). So structuralism implies that (1) is false, and this is certainly not an intuitive result. (Of course, structuralism does not imply that (1′) is true:

(1′) Ted Sider is identical to the number 2

since there are many ω-sequences whose third member is *not* me.)[3]

Is (1) in fact true? Common sense would presumably answer "yes". But the structuralist answer is that to think that (1) is true is to think of numbers as being (particular) objects. Since there is no one object that uniquely deserves the name 'one', this would be a mistake.

5.2.3 An Argument against Primitive Class Naturalism

Suppose for the moment that, contrary to Benacceraf, "numbers are objects". Without worrying overmuch about just what this comes to, we may simply take it to imply that either (1) or (1′) is true. If this were the case, then structuralism would be false. The question I want to pursue is whether

[3]As an alternative to the "Ramsification" approach in the text, a structuralist could view number theoretical sentences as being *partially interpreted*. On this view, there is no one intended interpretation of the language of number theory; rather there are a number of acceptable interpretations (one for each ω-sequence). A sentence is true if true on all acceptable interpretations, false if false on all acceptable interpretations. The number theoretic terms are equivocal over their interpretations on the various ω-sequences. Presumably, if a sentence is true on some of the interpretations but false on others, it lacks truth value. My point would still hold: (1) would not be true.

PCN runs into a difficulty of this kind. Put intuitively, the question is: does primitive class naturalism require that relations "be objects"? I think it does.

In this section, I will give my argument against primitive class naturalism. The problem with PCN is, in a nutshell, that it applies one of its primitive concepts (naturalness) to entities (ordered pairs) that are neither primitive entities PCN accepts, nor are they uniquely constructible in terms of such entities. The argument is related to arguments given by D. M. Armstrong and Peter Forrest.[4]

Let us continue for the moment at an intuitive level. The question of which relations are natural is not a "structural" question about the set of all relations. Rather, given a particular natural relation, it would seem that PCN requires there to be a fact *about that relation*, namely, that it is natural. Indeed, the question of whether a given relation is natural is a bit like the question of whether I am identical to the number 2—to answer either question we must be able to single out a particular entity and ask a question about *it*. By countenancing primitive naturalness, the class nominalist runs into the same difficulty that would beset the structuralist if numbers were "objects".

Let's look at this in more depth. PCN employs a primitive distinction between natural and non-natural relations. But relations are identified with classes of ordered pairs of possibilia. These classes of pairs are in turn identified with classes of classes of possibilia since pairs are constructed from classes. So, PCN invokes a primitive concept of naturalness, which applies to classes of classes of possibilia. What is more, this concept is taken to be an objective feature of the world.

Suppose I ask whether a given relation is natural. I have the concept of naturalness as applied to classes of possibilia, on one hand, and a question—Is *being ten feet from* natural?—on the other. How do I get my answer to that question from facts about naturalness? Remember that the relation *being ten feet from* will be identified with different classes, depending on which method for constructing the ordered pairs we use. Do I look at the class of Kuratowski pairs $\langle x, y \rangle$ where x is ten feet from y, and see whether the concept of naturalness applies to it? Do I look at Wiener pairs? Some other kind of pairs?

In outline, my argument against PCN runs as follows. I consider three

[4] See Armstrong (1986, pp. 86–7); Armstrong (1989*b*, pp. 29–32); Forrest (1986, pp. 90–1). Kitcher (1978) discusses related issues involving ordered pairs.

possibilities for formulating PCN and attempt to show that each is flawed. As near as I can tell, I consider all plausible formulations. The conclusion is that there is no acceptable formulation of PCN. The possible formulations are distinguished by the way they answer the question—Is relation R natural?—based on the way the concept of naturalness applies to classes of classes of possibilia.

Throughout this section I assume that ordered pairs are indeed arbitrarily constructible, and not in the primitive ontology of PCN. In the next section I will consider the result of giving up this assumption.

Possibility 1: One method

One possibility would be that there is one method of constructing ordered pairs such that the natural relations as constructed according to that method are all and only the classes with the property of naturalness. Let us suppose that method is Kuratowski's. So, the natural relations are sets of Kuratowski pairs.

Suppose we want to know whether the relation *being ten feet from* is a natural relation. The present proposal instructs us to consider the class of all pairs $\langle x,y\rangle$ where x and y are possible objects separated by ten feet, constructed according to Kuratowski's method. If and only if this class has the property of naturalness, our answer is yes. And it is crucial that we used Kuratowski's method—if we had used another method, say Wiener's, in constructing the class of pairs $\langle x,y\rangle$ where x is ten feet from y, we might have gotten the wrong answer! The class obtained in each case is different—perhaps exactly one has the property of naturalness.

I reject this answer since, I will argue, it contradicts the claim of primitive naturalism that naturalness is objective, and not a matter of convention.

Given a certain convention, I need a test for what is a matter of that convention. As an example, let us look at the convention of naming. My place of birth is not a matter of the convention of naming, for the proposition that Ted was born in New Haven has the same truth value as it would have had if our convention of naming were different, and 'Jeff' were my name. Contrast this with the proposition that Ted's name has 3 letters, which has a different truth value than it would have had if my name were 'Jeff' instead of 'Ted'. The moral: the number of letters in my name is a matter of the naming convention.

I take it as a premise that the choice of a method for constructing ordered pairs is conventional, just as the choice of what to name a person is conventional. This is quite plausible. Surely, the various methods for constructing ordered pairs are mathematical tricks for using classes to explicate concepts involving order. Surely, there is no ontological seriousness in our choice of whether to use this method or that method in constructing pairs.

Since the choice of a method for making pairs is conventional, propositions that change their truth value when we change our methods for making ordered pairs are about matters of convention—specifically, they are about matters of the pair-making convention. One proposition that does *not* change its truth value in this way is the proposition that the *earlier- than* relation is asymmetric. Suppose, for example, that this proposition is true, when we make pairs according to Kuratowski's method. This means that (i) the set of sets $\{x,\{x,y\}\}$ such that x is earlier than y is such that for any x and y, if $\{x,\{x,y\}\}$ belongs to it, then $\{y,\{y,x\}\}$ does not belong to it. Now, let us make pairs Wiener's way instead. We must show that (ii) the set of sets $\{x,\{y,\varnothing\}\}$ where x is earlier than y is such that for any x, y, if $\{x,\{y,\varnothing\}\}$ is in the set then $\{y,\{x,\varnothing\}\}$ is not. Clearly, ii) follows from i). (From (i) it follows that for any x, y, if x is earlier than y then y isn't earlier than x, and (ii) follows from this.)

Since the truth value of the proposition that the *earlier-than* relation is asymmetric does not vary when we change methods, this shows that the asymmetry of the *earlier-than* relation is not conventional (at least with respect to the convention of making ordered pairs). It should be intuitively clear that most propositions about relations will not vary in truth value when we switch methods of making pairs, since when we talk about relations we don't say things that depend on the quirks of a particular method. Typically, when we make assertions about ordered pairs, all we assume about the nature of pairs is that they obey the identity condition listed above (pairs are individuated by their members and their order). Of course, a proposition like the proposition that the transitive closure[5] of the *earlier-than* relation contains $\varnothing$ *will* vary in truth value when we switch methods (it will be true when we use Wiener's method; false when we use Kuratowski's method). This just goes to show that whether or not the transitive closure of the *earlier-than* relation contains $\varnothing$ is a matter of convention.[6]

---

[5] x is in the transitive closure of a set y iff x bears the ancestral of the membership relation to y—intuitively, if x is a member of y, or a member of a member of y, or a member of a member of a member of y, or ...

[6] I am, of course, assuming the truth of Class Nominalism for the moment: that relations

Back to naturalness. Clearly, the proposition that the *earlier-than* relation is natural will vary in truth value when we switch methods. Only sets of Kuratowski pairs have the property of naturalness. So, naturalness turns out to be a matter of convention, on the interpretation of PCN that I have called "Possibility 1". Thus, this version of PCN is contradictory, since it is built into the statement of PCN (via the statement of Primitive Naturalism) that naturalness is "objective", and not a matter of convention.

Possibility 2: All methods

Here is an alternative proposal. Instead of there being one distinguished method, surely *every* method is on a par. One way to develop this thought is as follows. In possibility one, the property of naturalness was had by all and only the classes that counted as natural relations according to *one* method. According to possibility 2, a class of classes counts as natural just when it corresponds to a natural relation under some method or other for constructing ordered pairs.

Suppose I ask: is *being ten feet from* natural? On this approach, to answer this question I must first choose a method for making ordered pairs. Think of this as an arbitrary choice of a language to discuss ordered pairs—I may choose any method I like. Suppose I choose some method M. I must now construct the class of pairs $\langle x, y \rangle$, where x is ten feet from y, according to method M. Now, the answer to my original question is *yes* iff this class has the property of naturalness. And since our naturalness property now applies to any class of classes that corresponds to a natural relation under *any* method, the answer to the question will be the same whatever method I choose.

The problem here is that too many relations will turn out natural, because a class of pairs that corresponds to a certain relation under one method of constructing ordered pairs may correspond to an entirely different relation under another method.

Suppose the relation *being ten feet from* is a natural relation. The present proposal says that for any method M, the class of M-pairs $\langle x, y \rangle$ where x is ten feet from y must have the property of naturalness. So, using the Kuratowski method, the following class S has the property of naturalness:

are sets of ordered pairs. If they are not, then obviously (6) won't vary in its truth value when we switch methods of constructing pairs.

$$S = \{\{x, \{x, y\}\} : x \text{ and } y \text{ are possibilia that are ten feet apart}\}.$$

Now let u and v be two objects that are ten feet apart, and let u' and v' be two other objects that are not ten feet apart. We define a new method X for constructing the ordered pairs. Intuitively, method X is just like Kuratowski's method, save that the pairs $\langle u, v\rangle$ and $\langle u', v'\rangle$ are swapped. Method X may be defined as follows:

$$\langle x, y\rangle = \begin{cases} \{u, \{u, v\}\} \text{ if } x = u' \text{ and } y = v' \\ \{u', \{u', v'\}\} \text{ if } x = u \text{ and } y = v \\ \{x, \{x, y\}\} \text{ otherwise} \end{cases}$$

Now consider the following relation R. R is just like *being ten feet from*, save that u and v do *not* stand in R, and u' and v' *do* stand in R. Plainly, R is not a natural relation. But the present proposal has the consequence that R is natural. The present proposal says that to ask whether a given relation is natural one must first choose a method for constructing pairs. Let us choose method X. That proposal says next that the relation is natural iff the class of ordered pairs of possibilia that stand in that relation, as constructed according to that method, has the property of naturalness. Well, the class of ordered pairs of possibilia that stand in R, as constructed according to X, is exactly S! The tricky clauses in the definitions of relation R and method X "cancel" each other, and we get exactly the same class identified with *being ten feet from* by the Kuratowski method. But S has the property of naturalness, and so R turns out to be natural. The present theory must be disposed of.[7]

[7] Phillip Bricker suggested a hybrid of Possibilities 1 and 2 that is worth discussing. Naturalness, on this view, is a property had by a set of sets just when it counts as a natural relation under some *natural* method for constructing ordered pairs. To ask whether a relation is natural, on this view, one must first choose a natural method M for making pairs, and then consider whether the set of M-pairs that stand in the given relation has the property of naturalness. This view escapes the difficulty I posed for Possibility 2 since method X is, allegedly, an unnatural method for constructing the pairs.

I think that this view runs into problems with conventionality, just like Possibility 1. This view instructs us to choose only *natural* methods to make pairs. But if we choose only natural methods, I think this is itself merely a matter of convention. So, if the truth value of

(*) *being 10 feet from* is a natural relation

switches when we move from Kuratowski's method to method X, then, I say, naturalness is a

Possibility 3: Naturalness is a relation between classes and methods

In light of the difficulty with possibility 2, perhaps we should keep track of which method we are using when we evaluate the naturalness of a class. Class S from the last section should count as a natural relation, considered as a class of Kuratowski pairs, but considered as a class of X-pairs (recall my method X), it shouldn't be natural.

As stated, this proposal is unacceptable. To say that a class is natural "considered as" a class of Kuratowski pairs, but non-natural considered another way surely cannot be taken seriously. Naturalness is supposed to be an objective feature of the world, so if a class has the property of naturalness, then it has that property regardless of how it is considered.

Of course, there is a serious proposal in the neighborhood. Perhaps the property of naturalness isn't a *property* at all, but rather a *relation*. When we say a class is natural considered as a class of Kuratowski pairs, what we mean is that the class bears the naturalness relation to the Kuratowski method. Class S will count as natural *with respect to method X*, but not with respect to the Kuratowski method. That is, S will bear *naturalness* to method X, but not to the Kuratowski method.

On this view, what is taken as a primitive by PCN is a relation, the naturalness relation, which holds between a class and a method iff the class counts as a natural binary relation when interpreted as a class of pairs constructed according to that method. Let us return to our question: is the relation *being ten feet from* natural? On the present theory, we must first choose a method M for constructing ordered pairs. Given this choice, we form the class of all M-pairs $\langle x,y\rangle$, where x is ten feet from y. On the present theory, the answer to our original question is *yes* if and only if this class bears the naturalness relation to our chosen method M.

To evaluate this proposal, we should draw a distinction we have overlooked so far. Earlier I said that PCN reduces properties and relations to *classes* —properties are classes of possibilia, relations are classes of 'tuples

matter of convention, The following argument shows that (*) does indeed switch truth value in this way. I assume that relation R from the text is unnatural, that the relation *being 10 feet from* is natural, and that Kuratowski's method is a natural method. Since Kuratowski's method is a natural method and *being 10 feet from* is natural, set S from the text is natural. Let S' be the set of Kuratowski pairs $\langle x,y\rangle$ that stand in R. Since R is unnatural, S' does not have the property of naturalness. Therefore, the hybrid theory has the consequence that (*)'s truth value switches when we move from Kuratowski's method to method X. For S' is also the set of X-pairs $\langle x,y\rangle$ where x is ten feet from y.

of possibilia. This is not quite right. Most relations are analyzed in this way, but not all. The relation ∈ of class membership, for example, is not analyzed as a *class* of anything, for this would be circular. In building an ontological theory, the champion of PCN appeals to a very few primitive properties and relations, (e.g. class membership) and to a group of entities (e.g. possibilia, classes). Most relations are thereby analyzed, but not all. Class membership never gets analyzed. A traditional terminology calls the primitive entities "ontology", and the primitive concepts "ideology" (Quine, 1951). Most relations, such as the relation *being taller than*, find a place in ontology, but ∈ is a part of ideology.[8]

When presenting Possibilities 1 and 2, I assumed that naturalness was a property, but I never said whether the property was a part of ontology or ide-

[8] Actually, one class nominalist, David Lewis, does not definitely accept ∈ in his ideology; in section 2.7 and the appendix of Lewis (1991) he considers the prospects of eliminating ∈ in favor of the *part-whole* relation. I ignore this complication and take class nominalists to be accepting ∈ as a primitive notion.

This view that Lewis considers in parts of Lewis (1991), structuralism in set theory (SST), faces our Benacceraf-style difficulties in new ways. A "singleton function" assigns to any object its singleton (if it has one). According to (SST), mereology does double work; singleton *functions* are constructed mereologically, and also mereology gives us the rest of the set-theoretic hierarchy, since on Lewis's view a set is the mereological sum of its singletons. For our purposes, the crucial feature of this view is the claim that we do not grasp any one singleton function. Set theory is not the theory of *the* singleton function, "but rather the general theory of all singleton functions" (Lewis, 1993, p. 15). Any set-theoretic claim ϕ is tacitly a universally quantified claim:

for any singleton function f, $\phi(f)$

where $\phi(f)$ is ϕ with its references to the membership relation replaced with appropriate constructions involving the variable f.

Now we can see the problem. There are many functions that assign singleton sets to entities. Sets are not renounced. But there is no one unique singleton of David Lewis. Some singleton functions assign one set to be his singleton, other functions allow other sets that privilege. A property, according to Lewis, is a *set* of possible individuals. And the natural properties will be a select few of these sets. But according to which singleton function? A certain set S that is a natural set (say, the set of red things) according to one singleton function may be, according to another singleton function, an entirely different and perhaps nonnatural set (the set of items purchased by me, say). Clearly, our problem has arisen again, and now it applies to natural *properties* as well as natural relations. The problem arises in another way as well: singleton *functions* are construed mereologically by Lewis, and these functions can be constructed in various equally adequate ways. See Lewis (1991, Appendix).

ology. That is, I never said whether PCN identified the property with a class. This is because the objections there did not depend on this distinction. But there is a special objection waiting if the naturalness relation of Possibility 3 is a part of *ontology*.

If *naturalness* is a class of ordered pairs then we have our trilemma all over again. Which method for constructing the pairs shall we use? To use only one would make naturalness a matter of convention. Instead, we could use every method: the class corresponding to the naturalness relation contains a subset for each method of constructing ordered pairs. No good, for this would generate spurious naturalness relations as we saw above. As a last resort we might appeal to a relation R between classes and methods: class S bears R to method M iff S is the naturalness relation, when interpreted as a class of ordered pairs according to M. This, of course, is the beginning of a vicious regress: what method is used for constructing R's ordered pairs? The naturalness relation, then, cannot be taken as a class of ordered pairs.

So instead let us take Possibility 3 to place the naturalness relation in ideology. Naturalness is a relation, but is no class of ordered pairs. This would escape the previous objection. But in its place I have another objection. The problem comes when we inquire into the nature of one of the *relata* of the proposed naturalness relation.

Naturalness is supposed to be a relation between classes and **methods**. I ask what a "method" for constructing ordered pairs is. I suppose many answers could be given. The Kuratowski method, for example, might be identified with a two-place function f, where $f(x,y) = \{x, \{x,y\}\}$. But the Kuratowski method could just as well be taken to be a different function $g : g(x,y) =$ the characteristic function of $f(x,y)$. And in either case, the functions will surely be construed as classes of ordered pairs. Since the ordered pairs are capable of multiple constructions, the functions will be as well. Or mightn't the Kuratowski method be taken to be the *sentence* 'let $\langle x,y\rangle$ be $\{x,\{x,y\}\}$'? Something else corresponding to the "directions" given by that sentence?

Let's face it. There is no one group of entities in the ontology of PCN that uniquely deserves the name 'the method of Kuratowski'. Rather, there are a number of constructions, each capturing the relevant properties of Kuratowski's method. But these constructions reintroduce the old difficulty. Our naturalness relation is supposed to apply to methods—the Kuratowski method, for example. But the Kuratowski method can be constructed according to many "meta-methods", so we have our trilemma all over again.

Can we use just one meta-method to construct the method of Kuratowski? No—surely the choice of a meta-method is conventional, and would therefore lead to conventional naturalness. Can we use all methods to construct the method of Kuratowski, thus allowing the naturalness relation to relate any method that counts as Kuratowski's method under any meta-method? No (some entity might be interpretable as more than one method). Can we let Kuratowski's method be a relation between constructions and meta-methods of interpreting those constructions? No; we will have the same problem all over again on the next level when we inquire into the nature of meta-methods.

As near as I can tell, I have considered and rejected all possibilities for formulating PCN. I conclude that no version of PCN is a viable theory, at least under the present assumption that ordered pairs are not in the primitive ontology of PCN.

The following summary may help us see the forest through the trees. An ontologist presents a total picture of reality in terms of some primitive entities and primitive concepts. After giving a theory of the workings of the primitive entities and concepts, constructed entities may be introduced, but first we must have the theory of the primitives. The Class Nominalist accepts possibilia and classes as the primitive entities, and class membership as the primitive concept. PCN then adds a primitive notion of naturalness. Naturalness must be a property or relation. If naturalness is a property, then what primitive entities have it? We rejected two possibilities here: classes of ordered pairs for one distinguished method (Possibility 1), and classes of ordered pairs for *each* method (Possibility 2). We rejected the first because it made naturalness conventional, the second because it made an obviously nonnatural relation, my concocted relation R, natural. On the other hand, if naturalness is a relation, then we must ask whether that relation belongs to ontology or to ideology. In the first case, we ran up against an infinite regress. In the second case, where naturalness was construed as a primitive binary relation, we noted that its relata are *methods*; since methods are arbitrarily constructible from the primitive entities of PCN (as are meta- methods, and meta-meta-methods...), we found no acceptable entities for the naturalness relation to relate.

Notice that other views about naturalness are unaffected by my argument. For example, if relations are taken to be primitive entities, rather than classes of ordered 'tuples, then there is no trouble. Even if we take relations to be classes of 'tuples, the argument of this chapter does not apply if, rather than taking naturalness as a primitive, we analyze it in terms of universals.[9] With these real entities doing the work of distinguishing between natural and unnatural relations, we need not say that one class is privileged as *the* class of natural ordered pairs. When we discuss relations, we must first pick a method M for constructing the ordered pairs. This done, a natural relation is a class S such that there is some dyadic universal such that S is the class of all pairs $\langle x, y \rangle$ (constructed according to M) where x and y instantiate that universal. At any time, we can switch to another method, because the universals are "out there", independent of our methods of constructing ordered pairs, ready to make the distinction among whatever entities we should propose as the relations.

5.2.4 Response 1: "The Argument Proves Too Much"

In the next two sections I will consider responses to my argument. The first response derives from a worry that my argument, if sound, would prove too much. In philosophy and mathematics, the concept of an ordered pair is frequently invoked. Surely, the fact that there is no uniquely distinguished method for constructing the pairs does not mean that these uses of pairs by philosophers and mathematicians are faulty! But couldn't my argument be extended to show these uses illegitimate? And if not, then what is so special about the use of pairs made by PCN?

I agree that most uses of pairs are legitimate. The case of PCN *is* special. What is special about PCN is that it grants significance to the particular objects to which pairs are reduced, by applying primitive naturalness to classes of pairs. In fact, it seems to me that my argument applies whenever a Class Nominalist postulates a non-conventional primitive property or relation of relations. In contrast, everyday uses of pairs are legitimate because on those uses no such significance is granted—rather, the pairs are only used to get a job done.

Let me illustrate my point with what I take to be a representative example of an ordinary use of ordered pairs. Binary relations—construed as classes of ordered pairs—are often taken by formal semanticists to be the meanings of dyadic natural language predicates. Imagine an objector asking: "what set, exactly, is the meaning of 'taller than'? What method do you use to make the pairs?" Presumably, we will answer that it doesn't matter—any method

[9] See, however, my reservations about this project in section 6.2.

is fine. "But wait!", the objector will say. "This means that what 'taller than' means is a matter of the convention of making ordered pairs. Granted, word meaning is conventional in that we can choose to make words mean anything we like. But word meaning isn't a matter of *this* convention, the convention of making ordered pairs. When we switch methods for making pairs, 'up' doesn't start meaning down. On to possibility 2 ...".

I think we should stop this argument short right at this point. We can accept the conclusion that what 'taller than' means is a matter of the convention for making pairs, in the following sense: exactly what entity is assigned as the semantic value for 'taller than' depends on the convention for making pairs. The conventionality conclusion that is unacceptable for PCN is, I claim, innocuous in the present case.

This conventionality would be unwelcome only if we required semantic theory to provide answers to questions like "what entity exactly is the meaning of 'taller than"', as if such questions are pre-theoretical and demanding of an answer. But we can claim that semantic theory seeks to provide a systematic account of certain intuitions speakers have about their language, intuitions perhaps about the *truth conditions of sentences* or about the *validity of inferences*. To accomplish this task, semantic theory is free to employ whatever "internal concepts" it chooses, and this choice is a matter of convention. What entities it uses in its modeling of language is not part of the "ultimate output" of the theory. They are "artifacts of the model". David Lewis agrees:[10]

> Semantic values may be anything, so long as their jobs get done. Different compositional grammars may assign different sorts of semantic values, yet succeed equally well in telling us the conditions of truth-in-English and therefore serve equally well...Likewise, different but equally adequate grammars might parse sentences into different constituents, combined according to different rules.

So, we only require our semantic theory to generate the correct results with respect to validity and truth conditions. We don't care about what particular entities are assigned as the meanings of predicates. Hence, we don't mind when we learn that it is conventional what object is assigned to be the semantic value of 'taller than'.

[10] See Lewis (1980, pp. 5,6).

Of course, there is a sense in which the meaning of 'taller than' is not a matter of the pairs convention. 'Taller than' does not begin to express the relation of being *shorter than*, no matter what convention for making pairs we choose. But my argument does not contradict this intuition. For example, whether or not the following sentences are true is *not* a matter of the convention of making ordered pairs:

(7a) The sentence 'Ted is taller than Mike' is true if and only if the sentence 'Mike is shorter than Ted' is true

(7b) If a sentence $\ulcorner \alpha$ is taller than $\beta \urcorner$ is true (where α and β are proper names) then the sentence $\ulcorner \beta$ is taller than $\alpha \urcorner$ is false

(7c) The sentence 'Ted is taller than Mike' is true if and only if Ted is taller than Mike

(7d) 'is taller than' applies to Ted and Mike (in that order)

I suggest that it is claims like these that we have in mind when we think that what words mean is not a matter of the pairs convention. The phrase 'what a word means' is ambiguous. Are we discussing the entity that is assigned to be that word's semantic value? Then we are discussing something that is (partly) a matter of the convention of making pairs. Or are we discussing meaning relations between words as in (7a), formal properties of a word's meaning as in (7b), or relations between words and the world as in (7c) and (7d)? Then what we are discussing is *not* a matter of the pairs convention.

Similar responses, I think, will answer most worries of the effect of the multiplicity of methods on uses of relations. The general idea is that we have a particular task at hand. To accomplish this task, we need the effect of entities that behave like ordered pairs, but we do not care what particular entities we use. Kuratowski pairs would be fine, but so would Wiener pairs. We don't care how the job gets done, so long as it gets done.

5.2.5 Response 2: Primitive Ordered Pairs

The second response I will consider is analogous to Wetzel and Resnik's response to Benacceraf's argument. That argument does not touch a believer in numbers as primitive, irreducible entities. In the present case, the PCN theorist could countenance a new class of primitive entities, the ordered pairs, and introduce a primitive ordered-pair forming operation. For any objects u and v, there is exactly one entity that is the true ordered pair of u and v. Granted, there are many class-theoretic structures that are isomorphic to the true ordered pairs, but these are not *reductions* of the ordered pairs. The property of naturalness, on this view, is had just by classes of true ordered pairs.

The horn of my trilemma that would no longer be unanswerable would be the first one. If ordered pairs are counted among the primitive constituents of the world, then the choice of a method isn't conventional after all. One method is distinguished, only now it is misleading to call it a method of *construction*: it is the "method" that identifies $\langle x, y\rangle$ with the true ordered pair of x and y.

I take it that such a response can indeed escape my argument. Thus, my argument should be taken as an argument against a particular version of primitive class naturalism—that version that accepts the orthodox view of ordered pairs according to which pairs are *not* primitive entities. However, I do have some comments on the theory of primitive ordered pairs.

First, we should note that primitive ordered triples will have to be accepted as well. In class theory, ordered triples are typically constructed from ordered pairs, once the ordered pairs have been constructed. We could construe the ordered triple of u, v, and w as $\langle\langle u, v\rangle, w\rangle$, but this is an arbitrary choice; we could have picked $\langle w, \langle u, v\rangle\rangle$ or $\langle u, \langle v, w\rangle\rangle$ instead. Perhaps some choices are more convenient than others, but surely this has no real ontological significance. So the trilemma above would apply to three-place natural relations. Primitive triples, therefore, had better be accepted. Similar arguments would require that that for every n, if there are any n-place natural relations then primitive n-tuples must be accepted.

So it seems that PCN will require primitive n-tuples together with an n-tuple formation operator for every number n such that there are natural n-place relations. We should notice that these additions would diminish the appeal of PCN, and specifically that of class nominalism. Class nominalism attracts us by its simplicity. We need only believe in possibilia and class theory, and in return we are promised ontological heaven. Now we find ourselves postulating new entities and operations. The price seems a high price to pay for the natural/nonnatural distinction. Perhaps we should shop elsewhere.

So, when we consider this enhanced version of PCN, we should not forget that PCN has competitors. The theory of universals, for example, purports to give an analysis of naturalness, although it pays the price of invoking universals. Then, there is a view that I favor: the combination of primitive naturalism with the claim that properties and relations are primitive entities.[11] We need to ask whether the addition of primitive 'tuples to PCN tips the scales of philosophical price against PCN in favor of its rivals.

Finally, let us note that one particular (potential) defender of PCN cannot take the way out offered by primitive ordered pair formation. David Lewis defends the following principle: "the 'generating' relation of a system should never generate two different things out of the very same material. There should be no difference without a difference in content."[12] As Peter Forrest (1986, 90–1) and D. M. Armstrong (1986, 86–7) have pointed out, primitive ordered- pair formation would violate this principle. To the extent that this principle is attractive, this new version of PCN is unattractive.

5.3 Naturalness and "Kripkenstein"

The natural/nonnatural distinction has been thought to help solve Kripke's (1982) version of a puzzle of Wittgenstein's. Surely, by using the word 'plus', I express the *plus* function, which assigns to any two numbers their sum. But my behavior seems to be consistent with my meaning something different by 'plus': *quus*. Quus is a function that assigns to x and y their sum if each is less than some suitably chosen number; seventeen otherwise. If the chosen number is high enough, argues Kripke, nothing about my behavior could constitute my meaning plus and not quus when I say 'plus'. And surely my behavior determines what I mean. It seems we ought to conclude that I do not mean plus any more than quus. But I *do* mean plus and not quus.

Lewis (1983*b*, pp. 375–77) proposes we solve the puzzle using naturalness. Plus, but not quus, is a natural function. Natural functions are *prima facie* more eligible to serve as meanings for our words. This is why I mean

[11] Of course, primitive properties and relations are no more economical in themselves than primitive 'tuples. They earn their keep, however, in eliminating the need for possibilia—standard reductions of possibilia and possible worlds invoke primitive properties and relations, or similar entities. See section 6.3.1.

[12] Lewis (1991, p. 38). Lewis wields this principle against structural universals in Lewis (1986*a*) and discusses it at length in Lewis (1991, section 2.3).

plus and not quus. Nothing about my behavior determines this. I may never have heard of naturalness. It is simply a fact about reference that reference goes to the most natural candidate—in this case, plus.

The importance of this example is that the notion of naturalness is applied to numbers. But now a new problem arises. It will be analogous to the problem of the previous section, only now it will generate a difficulty for all theories of naturalness, not just for PCN.

There is a difficulty with taking certain number-theoretic functions as being natural functions, if numbers are constructions from classes and are not primitive entities in their own right, for to countenance numbers as entering into natural relations is, I say, to treat numbers "as objects".

Suppose certain numeric functions are distinguished as the natural ones. Suppose further that numbers are constructions from classes. So certain class-theoretic functions are natural. We need to know to which method for constructing the numbers these classes correspond. There are three possibilities as before.

Possibility one: there is one privileged method. This I reject since the fact that the choice of a method for constructing the natural numbers is conventional would make naturalness as applied to numbers conventional, which in turn would make it a matter of convention whether I meant plus or quus by 'plus'. That seems wrong.

Possibility two: naturalness is a property that applies to any class- theoretic function that counts as a natural number-theoretic function under some method for constructing the numbers. Here, whenever we choose a method M for constructing the numbers, the natural number-theoretic functions are those class-theoretic functions that have the property of naturalness and are, according to M, number-theoretic functions. This is no good, since it generates too many natural functions. For example, we could construct the numbers in von Neumann's way, except with two and seventeen reversed. The class-theoretic function f that was the plus function according to von Neumann will have the property of naturalness. But interpreted according to the new method, f will count as some bizarre number theoretic function. So this bizarre number theoretic function will turn out natural: an incorrect result.

Possibility three: naturalness is a *relation* that holds between methods for constructing the natural numbers from classes and class- theoretic functions (i.e. relations). This is rejected on the grounds that "methods" are arbitrarily constructible entities, and so the difficulty is re-introduced. The Zermelo method for constructing the numbers might be conceived as the appropriate function from the class ω of finite ordinals into the class Z of Zermelo numbers (i.e. $\{\varnothing, \{\varnothing\}, \{\{\varnothing\}\} \ldots\}$). Or it might be conceived of as a similar function with a different domain (some well- ordered class isomorphic to ω). Or it might be a two-place function f from $\omega \times \mathrm{Z}$ into $\{0, 1\}$ (e.g. where α is the nth ordinal, $f(\alpha, s) = 1$ if s is the class Zermelo identifies with n; 0 otherwise). Or a similar two-place function from $\omega \times \mathrm{Z}$ into $\{\mathrm{T}, \mathrm{F}\}$. Maybe it is the "rule" by which we specify the Zermelo numbers, whatever that is (i.e. something corresponding to the directions: "let 0 be $\varnothing$, and let $n+1$ be $\{n\}$").

Since "the Zermelo method" is arbitrarily constructible, we have our old trilemma. Is there one entity that is the method of Zermelo? No. Is the method of Zermelo all of these entities? No. Is the method of Zermelo a relation involving meta-methods? No.

It should be clear that this argument is not particular to numbers. It generalizes whenever entities, like numbers, are capable of multiple equivalent constructions. I conclude, therefore, that only bona fide entities can enter into natural relations. "Entities" that are mere constructions from other objects and capable of multiple equivalent constructions cannot enter into natural properties or relations. More precisely, while most talk of mere constructs can be suitably paraphrased, talk of mere constructs having natural properties and entering into natural relations cannot be suitably paraphrased, and must therefore be abandoned.

There is no problem to the solution of the Kripkenstein puzzle, of course, if numbers are taken as primitive entities. Furthermore, there may be some other solution to the puzzle that invokes naturalness, but not of number-theoretic functions. But if numbers are viewed as mere constructions from classes, then the solution to the puzzle in terms of natural functions is untenable.

Chapter 6

Properties, Universals, and Naturalness

The time has come to discuss a class of competitors to primitive naturalism: "sparse theories of universals", especially a theory suggested by the writing of D. M. Armstrong. As I have indicated in chapter 1, it has been thought to be possible to use sparse universals to analyze naturalness. In this chapter I present these theories and examine the possibilities for analyzing naturalness in terms of universals. The outlook for a fully general analysis is not good. Finally, I criticize Armstrong's contention that immanent universals are superior to transcendent universals.

6.1 Sparseness vs. Abundance; Immanence vs. Transcendence

In this section I characterize various theories of universals by focusing on two questions. The first is the question of whether universals are sparse or abundant. The second is the question of whether universals are immanent or transcendent.

I will approach these questions by considering how David Lewis's D.M. Armstrong answers them. Let me explain. In his two volume work *Universals and Scientific Realism*, Armstrong laid out a theory of universals that was vague and inexplicit on the question of immanence. David Lewis interprets Armstrong in a certain way in *On the Plurality of Worlds*. Since then, Armstrong has been more explicit: in his book *A Combinatorial Theory of*

Possibility. However, I find Armstrong's version of his own theory difficult to understand, and therefore difficult to critically evaluate. I will therefore confine my remarks to Lewis's clearer Armstrong, who I will (somewhat unfairly) persist in calling "Armstrong".[1]

The terms 'universal' and 'property' are often used interchangeably, but I need to separate them when discussing sparse universals. I reserve the words 'property' and 'relation' for entities that obey the abundance assumption given in chapter 2: one property for every class of things; one n-place relation for every class of n-tuples. On some views of universals, universals are identifiable with the properties and relations, but for now let us forgo this simplification.

The distinction between universals and particulars is familiar. Particulars include people, planets, protons, etc. Universals, such as **roundness**, **unit positive charge**, etc. are *instantiated* by particulars. So far, no surprises. But I want to look at two main differences between Armstrong's theory of universals and traditional theories. Armstrong's universals are *sparse*, and they are *immanent*, or "wholly present" in their instances.

Universals are often taken to be *abundant*, rather than sparse. Exactly what we take the term 'abundant' to mean is somewhat open. A universal may be admitted for every meaningful predicate of ordinary language. We might have a universal for "every way a thing could be", on some suitable interpretation of this phrase. Abundant universals would surely be closed

---

[1] Armstrong (1989*a*) introduced a non-mereological relation called "constituency". Objects, properties, and relations are constituents of the states of affairs ("thick particulars") that involve them. Also, properties and relations are constituents of their structural combinations. What is relevant to our purposes is that Armstrong can characterize the claim that universals are *immanent* as follows: universals are *constituents* of the (thick) particulars (i.e. states of affairs) that instantiate them.

This affects various points in my discussion. In section 6.3.2, for example, I ask whether transcendent universals explain resemblance better than immanent universals. My answer is no, but it is based on Lewis's interpretation of 'immanence'. Perhaps Armstrong's conception of immanence fares better!

The reason I do not consider Armstrong's own conception of immanence is that we are told so little about constituency. Armstrong tells us that universals are not *parts* of their instances (1989*a*, p. 41), and yet they are "present as a whole in" their instances (1978*a*, p. 68). What, then, is this relation of constituency? Before I know more about it, I cannot evaluate a view based on it.

Note also that the real Armstrong's view has distinctive features along other dimensions. For one, Armstrong rejects uninstantiated universals (1978*a*, p. 113). For another, he says that universals and particulars are "abstractions from" states of affairs (1989*a*, p. 43).

under the Boolean operations.

For our purposes here, let us construe abundance broadly. Say that a property P and a universal U *correspond* iff the set of U's possible instances is identical to the set of P's possible instances. I take the claim that universals are abundant to be the claim that for every property there is a corresponding universal. Recall my chapter 2 abundance assumption for properties: one for every class of things.

In contrast, on a "sparse" view of universals, only a select minority of properties have corresponding universals.[2] Like 'abundant', 'sparse' is schematic. I will characterize sparseness in roughly the same way that Armstrong does, although no precise criterion is given. One non-negotiable constraint on the sparseness of universals for Armstrong is the following:

(S) A property (relation) has a corresponding universal only if that property (relation) makes for similarity

(See section 3.3.1 on the notion of making for similarity.)

As an example of the application of (S), suppose that **redness** and **blueness** are both universals.[3] This seems consistent with (S), since shared color seems to be a genuine similarity. But the sharing of the property *redness-or-blueness* does *not* seem to count as a genuine similarity. A red object could be quite dissimilar to a blue object. Thus, (S) prohibits a universal of **redness-or-blueness**. Similar remarks apply to negations of universals. In fact, for Armstrong there are no disjunctive or negative universals (1978*b*, pp. 19–29).

For Armstrong, the properties that correspond to universals are an elite minority of the entire set of properties.[4] Likewise, the perfectly natural properties are an elite minority of all the properties. This analogy suggests the possibility of analyzing perfect naturalness in terms of sparse universals. In the next section, I will discuss this analysis.

There is more to be said about sparseness, but I postpone that discussion until section 6.2. Let us move to the second feature of Armstrong's

[2]This idea runs throughout Armstrong (1978*a*,*b*). See, for example, Armstrong (1978*b*, pp. 9–12).

[3]Armstrong would not accept a universal **redness**—see Armstrong Armstrong (1978*b*, p. 117).

[4]However, Armstrong would not put matters this way, since he does not accept the existence of the abundant properties—instead, he accepts "propositional predicates" (1978*a*, pp. 3–6).

view. His universals are *immanent*—that is, they are "wholly present" in their instances. Armstrong contrasts immanent universals with *transcendent* universals (1978*a*, chapter 7). Transcendent universals, or "Forms", are supposed to be "separate" from their instances, to which they are linked by the relation of "participation".

We have what looks like a disagreement: are universals *immanent*—"wholly present" in their instances—or are they *transcendent*—"separate" from their instances? Let us call the view that universals are immanent "IU", and the view that universals are transcendent "TU". So far, my characterization of the dispute over immanence is too vague. What is it for a universal to be "wholly present" in its instances? What is it for a universal to be "separate" from its instances?

First, the issue is mereological. The claim that universals are wholly present in their instances would seem to mean at least that universals are *parts* of their instances. If **redness** is a universal, then any given red thing has the universal **redness** as a part. Thus, if x and y are both red, then x and y *overlap*. That is, they have a part in common; namely, the universal **redness**.

"Separate" too may be taken mereologically. In this sense, when the defender of TU claims that "universals are separate from their instances", what is meant is that universals are (mereologically) disjoint from their instances.

According to IU, then, a universal is a part of each of its instances. Since universals are supposed to be "wholly" present in their instances, it might be thought that we must add that the *whole* of the universal is a part of each instance, i.e. that every part of the universal is a part of the instance. However, this would be redundant since the phrases:

> x is a part of y
>
> every part of x is a part of y

are equivalent.

Armstrong intends his claim that universals are immanent to mark a profound metaphysical difference between his view and views that deny this claim. For example, Armstrong criticizes Platonism, a version of TU, for its failure to locate an object's properties ("Forms") *in* that object (1978*a*, p. 68):

> Is it not clear that *a*'s whiteness is not determined by *a*'s relationship with a transcendent entity? Perform the usual thought- experiment and

> consider *a* without the form of Whiteness. It seems obvious that *a* might still be white…
>
> It is important to see that this argument succeeds only against a doctrine of *transcendent* universals. It would fail…against the view that the Form is something *present as a whole* in *a*…

But I will now argue that *if* there is nothing to "wholly present" beyond the present mereological claim, *then* there is no genuine metaphysical significance to the question: are universals immanent or transcendent?

Consider the universal **redness**, and consider a particular x that is red. According to TU, **redness** is not a part of x. But now consider the fusion of x and all of its universals. **Redness** is indeed a part of this entity. Why not consider this fusion the particular? In what follows I will develop this intuition.

TU countenances particulars; we refer to and quantify over these objects in everyday language. It also countenances universals in which those particulars participate; these may be predicated of particulars in ordinary language. Moreover, no particular overlaps any universal.

Consider any particular x that participates in universals $U_1, U_2, \ldots$. Let y be the fusion of x and the universals $U_1, U_2, \ldots$. Let us call y a *thick particular* —the fusion of a particular with all of the universals in which it participates. We may also introduce a new term 'participates*': a thick particular participates* in a universal iff the original particular from which it was derived participates in that universal.

A hybrid form of TU is "HTU". According to HTU, it is *thick* particulars rather than ordinary particulars that are properly called "particulars"; it is thick particulars that we refer to and quantify over in ordinary language; we predicate of them the universals in which they participate*.

Imagine a vocal advocate of IU berating a defender of TU: "You are mistaken in your belief that universals are transcendent. Change your ways; admit that universals are *in* their instances." Then imagine the reply: "I accept your criticism. I was wrong in accepting TU. I now accept HTU. Universals are indeed parts of their instances—universals are contained in the particulars (that is, thick particulars) that participate (that is, participate*) in them."

The "concession" by the defender of TU is clearly no concession at all. He has changed no metaphysical views; all he has changed is what he calls "particulars". Before he took the term 'particular' to refer to objects without their universals, he now takes that term to refer to fusions of objects and their universals. But he has accepted no new entities, and no new facts about entities he accepts.

Thus, if immanence were merely mereological, then there would be no real metaphysical dispute over immanence.

We can illustrate the situation with a picture. Let us begin with a particular, x, that instantiates universals U_1, U_2, ...To begin on neutral ground, let us mereologically *subtract* U_1, U_2, ...from x; call the result TP. TP may be called a "thin particular" since it contains no universals that it instantiates as parts.[5] According to TU, TP is simply x, the object we began with, since U_1, U_2, ...weren't part of x to begin with. According to IU, on the other hand, TP is distinct from x since x contained those universals. Now consider the following diagram:

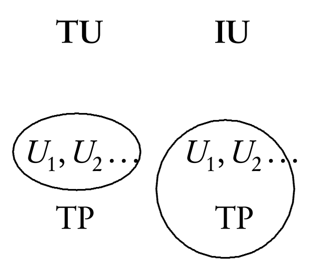

Figure 6.1: Particulars: TU vs. IU

The circles in the diagram represent what the theory in question calls a "particular". The diagram makes it plain that the only difference between TU and IU (as I have so far characterized them) is semantic. Each accepts the same metaphysical picture; they merely draw the circles differently—that is, they differ over the reference of 'particular'.

Surely the dispute over immanence *is* metaphysically significant. What has gone wrong? We have left out a crucial element of the claim that universals are "wholly present" in their instances. David Lewis is clear on this matter. In his presentation of Armstrong's view in *On the Plurality of Worlds*, he says that immanent universals are *located* exactly where their instances are located. And it is not merely that *part* of the universal is co-located with one

[5] Armstrong uses the term '"thin" particular' in Armstrong (1978*a*, pp. 114–5), but in a different (and somewhat mysterious) sense. Lewis mentions what I call thin particulars in Lewis (1986*c*, p. 65).

instance, part of it with another. The *entire* universal shares total spatiotemporal location with *each* of its instances. Thus, universals are "recurrent", or "multiply located" (1986*c*, p. 64). A universal is *all* here, and **also** *all* there.

This added feature provides a genuine metaphysical difference between IU and TU. Transcendent universals are typically taken to exist "outside of" the spatiotemporal world. Let us return to the diagram. The difference would be in *where* I draw the universals U_1, U_2... On the TU part I could draw those universals anywhere, indicating that they are not in space and time. But on the IU part, I need to indicate that U_1, U_2... are *co-located* with TP. The diagram must reflect the difference.

We may construe the opposing view, TU, as taking the opposite position: no universal spatially coincides with any particular. In fact, let us take this to follow from a more basic claim: that universals have no spatiotemporal location whatsoever.

But I want to draw a moral from our consideration of HTU. We should separate out two components of the claim that universals are immanent, or "wholly present" in their instances. One component, the mereological component, is the claim that universals are parts of their instances. The other component, the spatiotemporal component, is the claim that universals share spatiotemporal location with their instances. Likewise, the claim of TU that universals are transcendent, or "separate" from their instances has two components. The mereological component is the claim that no universal overlaps any particular. The spatiotemporal component is that no universal spatially coincides with any particular (in virtue of the fact that universals lack spatiotemporal location). As we have noted, the spatiotemporal components seem to provide the real metaphysical difference between the views.

Here are two questions. Question #1: are universals sparse or are they abundant? Question #2: are universals immanent or are they transcendent? Using these questions we can distinguish *four* possible theories of universals. The first theory, "sparse IU", takes the first option in both questions. Armstrong's theory is a version of sparse IU. Another theory, which we might call "sparse TU", accepts sparse transcendent universals. "Abundant TU" accepts the second answer to both questions. A fourth and rather odd theory, "abundant IU", would accept abundant immanent universals.

Only the first two theories, the ones whose universals are sparse, are of any help in analyzing naturalness. The next section considers the possibility of analyzing naturalness in terms of sparse universals. What I say there applies equally to the first two theories: sparse IU and sparse TU.

6.2 Universals and Naturalness

In the following section I examine the possibility of using a sparse theory of universals to analyze naturalness.

The project in which I am interested is that of using a sparse theory of universals to analyze the concept of naturalness as applied to the abundant *properties* (and *relations*). For example, the universals will be used to pick out which properties are perfectly natural.

As mentioned in chapter 1, Lewis's suggestion is that 'perfectly natural' might be analyzed in terms of universals as follows:

(U1) property (relation) *P* is *perfectly natural* iff it corresponds to some universal

So far, I have only mentioned one constraint on sparseness: (S). This leaves many questions of the form 'Does *P* correspond to a universal?' unanswered. Indeed, there are various ways to develop a sparse universals view. On some ways (U1) might be acceptable; on others it might not.

Armstrong's leading intuition is that what universals there are is an *a posteriori* matter, to be established by "total science". We have reason to postulate only those universals required to explain the genuine similarities in the world.[6]

This conception seems to match with Conception 2 of perfect naturalness from 3.3.1. Conception 2 invokes natural properties to explain resemblance; Armstrong invokes universals for the same purpose. On Conception 2, (consistent) conjunctions and structural combinations of perfectly natural properties and relations are perfectly natural. With one exception, Armstrong too accepts conjunctive and structural universals, where a "conjunctive" universal is one that is the conjunction of others; a "structural" universal is one that is a structural combination of others, called its "constituents" (1978*b*, pp. 30, 68–71). (The exception to Armstrong's acceptance of conjunctive and structural universals is that he does not accept the existence of a universal in a world in which it has no instances (1978*a*, p. 113). For example,

[6] This idea is implicit in much of Armstrong (1978*a*,*b*). See, for example, the introduction to Armstrong (1978*a*) and Armstrong (1978*b*, pp. 7–9).

if something is U, and something else is V, but nothing is both U and V, then universals U and V will have no conjunction (at that possible world). I will ignore this complication in what follows.) So (U1) provides, perhaps, an adequate analysis of Conception 2- style perfect naturalness.

(U1) does *not*, however, provide an adequate analysis of Conception-1 style perfect naturalness. Conceptions 1 and 2 disagree. As we saw in section 3.2.1, conjunctions and structural combinations of perfectly natural properties are *not* perfectly natural. This means that (U1) must be modified, if we are to have an analysis of Conception 1-style perfect naturalness. Otherwise, (U1) would generate perfectly natural properties corresponding to structural and conjunctive universals. But we can restrict (U1):

(U2) property (relation) P is *perfectly natural* iff there is some nonconjunctive, nonstructural universal U such that the set of P's possible instances is identical to the set of U's possible instances

It seems, then, that we can evade the problems caused by conjunctive and structural universals by offering (U2) as our definition of 'perfectly natural', as construed by Conception 1. (From now on, I consider only Conception 1-style naturalness.) But there is more to naturalness than perfect naturalness. I will argue that universals cannot help us in analyzing the *more natural than* relation.

Universals, at first blush, seem only to help with the upper end of naturalness. How to use them to analyze relative naturalness? One strategy would employ conjunctive and structural universals. The intuition behind this strategy is that conjunctions are less natural than their conjuncts, and structural universals are less natural than their constituents. For any property or relation P that corresponds to a universal, denote that universal by "univ(P)" (I assume there are no distinct, necessarily coextensive universals according to the sparse views). I think that it is plausible that the following gives a sufficient condition on the *more natural than* relation:

(C1) For any two distinct properties (relations) P and Q that correspond to universals, if univ(P) is a constituent or a conjunct of univ(Q) then P is more natural than Q

However, this condition does not seem to be necessary. First, suppose the following are *equally natural* and correspond to universals: properties P_1

and P_2, and relations R_1 and R_2. Now, let Q_1 be the property of having two proper parts x and y such that R_1xy and P_1x and P_1y; let Q_2 be the property of having two proper parts x and y such that R_2xy and P_2x and P_2y. Q_1 corresponds to a structural universal that has univ(P_1) and univ(R_1) as constituents. Q_2 corresponds to a structural universal that has univ(P_2) and univ(R_2) as constituents. Intuitively, P_1 is more natural than Q_2. Similarly, P_2 seems more natural than Q_1. P_1, P_2, R_1, and R_2 are all on the same "level" of naturalness, whereas Q_1 and Q_2 seem to be on a different "level". And yet, were the 'if' in (C1) changed to 'if and only if', these intuitive judgments would be overruled. Univ(P_1) is neither a conjunct nor a constituent of univ(Q_2).

It might be thought that we can repair the difficulty as follows. Extend (C1) into a necessary and sufficient condition for the relation *being more natural* than*:

(C2) For any properties (relations) P and Q that correspond to universals, P is more natural* than Q iff univ(P) is a constituent or a conjunct of univ(Q)

Then, analyze the *more natural than* relation as follows:

(C3) For any properties (relations) P and Q that correspond to universals, P is more natural than Q iff there are properties (relations) P' and Q' that correspond to universals and are such that i) P and P' are equally natural, ii) Q and Q' are equally natural, and iii) P' is more natural* than Q'

In the example above, (C3) yields the result that P_1 is more natural than Q_2, for P_1 is equally as natural as P_2, Q_2 is equally as natural as itself, and, by (C2), P_2 is more natural* than Q_2 (since univ(P_2) is a constituent of univ(Q_2)).

Unfortunately, (C3) saddles us with a new difficulty: that of providing an analysis of the *equally as natural as* relation in terms of universals. I do not see how this could go.

There is another, independent difficulty with the project of defining 'more natural than' along these lines. The proposals so far have only concerned properties that correspond to universals. But what about other properties? For example, it seems intuitive to say that the property *being red or green* is more natural than the property *being grue or being identical to George Bush* *or being five feet away from someone with 6 coins in his pocket*. But neither of these properties corresponds to a universal, since each fails the similarity test contained in (S). Therefore, (C3) is silent with respect to them.

An entirely different method for characterizing relative naturalness *would* have consequences for properties that don't correspond to universals. Begin by using (U2) to characterize perfect naturalness, and then use the notion of distance from a set to characterize relative naturalness as was explained in section 3.3.1:

(5) P is at least as natural as Q iff P is at least as close to N as Q is

(where N is the set of perfectly natural properties and relations).

This method is fraught with difficulty. First, it doesn't achieve the goal of analyzing naturalness in terms of universals, for it invokes an additional primitive: distance from a set. That notion was used in chapter 3 to illuminate naturalness, but in the end naturalness was the primitive. In the present case, the goal is to avoid taking naturalness as a primitive, so we really are engaged in analysis. (5) here is being offered as a *definition* of 'at least as natural as' and hence we would need to take as primitive *distance from a set* (unless some way were found to analyze it). In addition to appealing to a new primitive, (5) is wrong. Let us adapt our description of Onion, the world of endless complexity, as follows: no non-structural universals are instantiated at Onion—rather, every universal is a structural universal. So, (U2) implies that there are no perfectly natural properties at Onion. Thus, most properties at Onion will not supervene on N at all (the only exceptions will be properties that supervene on the spatiotemporal relations alone), and (5) implies that any two such properties are equally natural—an unacceptable result.

In section 3.2.3, I showed how to get around (5)'s problems by offering the weaker (5a) and (5b). But (5a) and (5b) are of no use to us here, for they do not give general necessary and sufficient conditions for one property or relation being at least as natural as another. They merely give such conditions in certain special cases. Their restricted nature did not cause problems in chapter 3, since my goal there was not *analysis* of naturalness. However, the present goal *is* analysis.

It seems that, while we can define 'perfectly natural' in terms of sparse universals, we cannot so define 'more natural than'. Thus, using universals we cannot match the power of primitive naturalism. Given the results of section 4.2.1, this is a crucial shortcoming. I argued that the definition of 'duplicate' given by Lewis in terms of 'perfectly natural' was unsatisfactory. A revised version was given, but that version employed 'more natural than'. Since the analysis of duplication is one of the key uses for naturalness, we cannot simply abandon the *more natural than* relation.[7]

6.3 Armstrong's Objections to Transcendent Universals

6.3.1 Two Theories of Abundant Universals

In the previous section I considered the possibilities for using a sparse theory of universals to analyze naturalness. I argued that the project had difficulties. I favor instead the view that naturalness should be taken as a primitive distinction among properties conceived abundantly, where those properties are not reduced to classes of possibilia.

In this section I say briefly why I favor this view. What I say will fall far short of a full justification, but that would (thankfully) be beyond the scope of this dissertation. I want to accept abundant properties, for they seem theoretically important.[8] These abundant properties (relations) could be construed as classes of ('tuples of) possibilia. But in chapter 5 I raised a difficulty for applying primitive naturalness to sets of pairs of possibilia, at least when pairs are constructed arbitrarily from sets. My argument does not apply if primitive 'tuples are accepted. But I do not even wish to construe relations as sets of primitive 'tuples of possibilia, for the simple reason that I am wary of possibilia. I would rather reduce the possibilia to properties and relations than vice versa. So I seem to be stuck with the abundant properties as *sui generis* entities (here is an exception to my usual neutrality on such ontological issues). At least, this is the official view that I want to work with.

(Since I am no longer discussing sparse properties or universals, I now

[7] Lewis (1983*b*) details some of the proposed uses for naturalness, and many involve duplication. Moreover, the solution to the problem of "The Content of Thought and Language" on p. 372 in that paper requires naturalness to come in degrees (and to apply to properties that don't correspond to universals at that). But see note 37.

[8] See, for example, Lewis (1983*b*, pp. 348) and Lewis (1986*c*, 66–67) on the need for an abundant conception of properties.

use the terms 'property' and 'universal' interchangeably. Likewise, I will use 'participates' and 'instantiates' interchangeably.)

One view of abundant properties would have these properties be wholly present in their instances. This is the view I called "abundant IU" at the end of section one, the "odd" view. David Lewis mocks abundant IU in Lewis (1986*c*, p.67):

> ...it is just absurd to think that a thing has (recurring or non- recurring) non-spatiotemporal parts for *all* its countless abundant properties.

This view may be absurd, but *if* it is absurd, then it is the *spatiotemporal* component of "wholly present" that makes it absurd.

To see this, let us return to the hybrid theory HTU that we constructed in section 6.1. A "particular", according to this view, is really a *thick* particular: a fusion of an ordinary particular with all of its universals. A thick particular contains as parts all of the universals in which it participates*. This claim does not become absurd when we add that these universals are abundant. In fact, if one believes in abundant universals and arbitrary fusions of things, then one is *committed* to this conclusion. What I think appears absurd to Lewis is the claim that a particular, whether thick or no, shares total spatiotemporal location with all of its abundant universals.

I want to believe in an abundant conception of universals. But I am unwilling to accept immanent universals. In part this is because of the oddness of abundant IU. Another reason is that the spatiotemporal component of immanence is *prima facie* implausible in its own right. The notion that an entity can be multiply located—all of it is *here* and also all of it is *there*—is one that, other things being equal, we should reject.

If there are arguments that immanent universals are superior to transcendent universals then, assuming we need universals at all, this presumption against immanence would be outweighed. Armstrong offers such arguments. But I will argue in the next section that these arguments fail.

Thus, I take it that the way is clear to accepting abundant TU. Let me summarize my choices with respect to properties and relations in this section. I chose:

i) **abundant over sparse** (since the latter couldn't be used to define naturalness and the former are theoretically useful)

ii) **sui-generis over classes of possibilia** (because of chapter 5 and because I am wary of possibilia)

iii) **transcendent over immanent** (because immanence is prima facie implausible and is not—contra Armstrong—superior to transcendence)

6.3.2 Armstrong's Arguments against Transcendent Universals

In the present section I will critically assess Armstrong's arguments against TU, the view that universals are transcendent. This view is opposed to IU, the view that universals are immanent. The arguments in this section are designed to apply to TU, but not to IU. Therefore, if the arguments are convincing, they would point to a superiority of IU over TU. But I will argue that none of the arguments is convincing. None of the arguments presents an unsavory bullet for the proponent of TU to bite.

The Regress Argument

> ...the Relation regress, first stated, as far as I know, by Ryle [Ryle (1939, 137–8)], appears to be vicious. Particulars participate in Forms. The relation of *participation* is therefore a type having indefinitely many tokens. But this is the very sort of situation which the theory of Forms finds unintelligible and insists on explaining by means of a Form. The theory is therefore committed to setting up a Form of Participation in which ordered pairs consisting of a particular and a first-order Form *participate*.
>
> Once again, however, the problem is reproduced. If this second- order participation is something different in nature from first-order participation, then it requires to be explained by third-order participation, and so *ad infinitum*. But if second-order participation is the same in nature as first-order participation, then the analysis of first-order participation is proceeding in terms of this (first- order) participation in a Form, which is circular (Armstrong, 1978*a*, p. 70).

This argument is quite straightforward. I will not spend much time discussing it, since I believe David Lewis's discussion in "New Work for a Theory of Universals" (pp. 352–355) to be thorough and conclusive. The argument fails because it misconstrues TU. It assumes that the TU-ist attempts to provide an analysis of *all* predication, of all sameness of type. Thus, Armstrong assumes that if for various particulars x and Forms F, x *participates* in F, the TU-ist must postulate a Form of participation. If the goal of TU were to analyze all predication, all sameness of type, then this would indeed follow. But, as Lewis points out (1983*b*, p. 353):

> Doing away with all unanalysed predication is an unattainable aim, and so an unreasonable aim. No theory is to be faulted for failing to achieve it. For how could there be a theory that names entities, or quantifies over them, in the course of its sentences, and yet altogether avoids primitive predication? Artificial tricks aside, the thing cannot be done.

As Lewis points out, Armstrong's own version of IU does not accomplish the task of eliminating primitive predication. The relation that holds between Armstrong's universals and their instances, *instantiation*, is not analyzed.

Thus, I take it that neither IU nor TU seeks to analyze all predication. IU takes the relation of *instantiation* as a primitive; TU takes the *participation* relation as a primitive.

This response to the Relation Regress is the same as the response to another argument offered by Armstrong. He calls it "The restricted third man" argument (1978*a*, pp. 72–73):

> Consider the Forms. Each of them is its own unique self. But they do have something in common. They are different tokens of the one type. They are all Forms. Formhood is a one which runs through this many. So must there not be a Form of Formhood?...
>
> ...Consider the collection of first-order Forms plus the Form of Formhood. The members of this expanded collection have something in common. The different tokens are all of the same type. In consistency, therefore, they must all said to participate in a third- order Form of Formhood. The regress then continues. It is either vicious or, at best, uneconomical.

Now, if this "third-order Form of Formhood" were identical to the original Form of Formhood, then there would be no regress. But Armstrong argues that a property cannot have itself as one of its properties (1978*b*, chapter 19 section VI, and chapter 23 section II). Thus, he claims, the original argument stands.

Without considering whether Armstrong is correct in rejecting the notion of a property instantiating itself, I think we can dismiss this argument in the same way that we dismissed the first regress argument. TU is not engaged in the project of analyzing *all* predication. As well as taking the two-place relation of *participation* as a primitive, I assume that the TU-ist also takes the notion of a *Form* (universal) as a primitive.[9] So the TU-ist will reject the step in the argument where a Form of Formhood is invoked. There is no Form of Formhood—the fact that the various Forms are all Forms is a fact incapable of further analysis.

The Duplication Argument

Another argument is contained in the following passage (Armstrong, 1978*a*, p. 69):

> Suppose that *a* and *b* have quite different properties. According to the theory of transcendent Forms they are *in themselves* exactly the same. Their only differences lie in their relational properties: their relations to a different set of Forms. But may there not be a difference of nature in *a* and *b*, beyond mere numerical difference? Yet this difference the theory of Forms could not account for.

I think we may interpret Armstrong as follows. When Armstrong says that *x* and *y* are "in themselves exactly the same", we may take him to be saying that *x* and *y* are *duplicates*. Thus, the argument seems to be:

Duplication Argument 1

(1) If TU is true, then *a* and *b* are duplicates

(2) *a* and *b* are not duplicates

(3) Therefore, TU is not true

Premise (2) is true by stipulation: we are asked to consider objects *a* and *b* that are not exactly alike. And premise (1) has some plausibility. Imagine a TU-ist trying to claim that *a* and *b* are not duplicates. He might point out that only *a* is white. "But," Armstrong would respond, "according to *you*, **whiteness** is a *relational* property. After all, *a* is white because of its relation to a wholly separate entity: the form of **whiteness**. The "difference" you have pointed out between *a* and *b* is not a difference between the marbles considered *in themselves*. It is a mere relational difference."

[9] Alternatively, Formhood might be analyzed in terms of participation.

As initially appealing as premise (1) may seem, the defender of TU has a response. We must distinguish two senses of 'relational property'. One we may call the "trivial" sense:

> P is relational in the trivial sense $=_{\text{df}}$ for any object x, if P is had by x then this is in virtue of x's relations to objects wholly distinct from x

Clearly, if TU is true then every property is relational in the trivial sense, for according to TU objects do not contain their properties as parts.

The other sense we may call the non-trivial sense. How we define this sense depends on the exact method for drawing the various distinctions that are the subject of this dissertation; I favor using naturalness to analyze all the rest. At any rate, one sufficient condition for a property's being *non*-relational (in a non-trivial sense) is its being *intrinsic*.[10]

What the TU-ist will claim is that it is the non-trivial sense of 'relational' that is relevant to the question of whether *a* and *b* are duplicates. Let us suppose that the property *being white* is an intrinsic property.[11] Objects *a* and *b*, then, are not duplicates, since they differ over an intrinsic property. Granted, this property is relational in the trivial sense. But it is *not* relational in the non-trivial sense since it is intrinsic, and hence we are free to use it in explaining why *a* and *b* are not duplicates. Thus, we can reject premise (1). TU is consistent with *a* and *b* failing to be duplicates.

Suppose Armstrong were to stipulate that he intends to be using 'duplicate' in such a way that relational properties in the trivial sense are the relevant properties. That is, x and y are duplicates in this new sense iff they share all non-relational properties in the trivial sense. We may call this the trivial sense of duplication. The argument now reads as follows:

Duplication Argument 2

(1) If TU is true, then *a* and *b* are duplicates in the trivial sense

[10] Giving a necessary and sufficient condition for a property's being relational is tricky, and I won't attempt it. The trickiness results from the fact that some extrinsic properties seem non-relational; e.g. haecceities like *being George Bush*.

[11] In fact I am not sure that color properties are intrinsic. See section 1.1. However, even if they are not intrinsic, surely objects within a given world that differ in color differ with respect to some intrinsic properties (microscopic properties of surfaces, for example).

(2) *a* and *b* are not duplicates in the trivial sense

(3) Therefore, TU is not true

The TU-ist will presumably grant premise (1) of this version of the argument. But premise (2) will be rejectable. For *a* and *b* to fail to be duplicates in the trivial sense, they would have to differ with respect to some property that is non-relational in the trivial sense. But the defender of TU does not grant the existence of any such property since on that view every property is relational in the trivial sense.

The Argument From Causal Powers

A related argument is the following (Armstrong, 1978*a*, p. 75):

> It is natural to say both that the causal powers of a particular are determined by its properties, and that these powers are determined by the particular's own self and not by anything beyond it. But if the theory of transcendent universals is accepted, a thing's properties are not determined by its own self, but rather by the relations it has to Forms beyond itself.

Armstrong's claim here is that TU is inconsistent with the truth of:

(CP) The causal powers of an object are not determined by objects "beyond" it

From our discussion of the meaning of 'wholly present in', we know to distinguish two senses of 'beyond'. On one sense, (CP) means:

(CP1) The causal powers of an object are not determined by any objects not a part of it.

As I see it, the TU can simply reject (CP1). And furthermore, the rejection of (CP1) need be no embarrassment. HTU entails (CP1), while TU (we may grant), entails the denial of (CP1), and yet HTU is no different metaphysically from TU. So the question of whether (CP1) is true or not is simply the question of whether we mean thick or thin particulars by 'object'. In rejecting (CP1), the TU-ist is making a mere semantical claim, and not going out on a metaphysical limb.

Alternatively, Armstrong may be taken to be claiming that TU is inconsistent with the following truth:

(CP2) The causal powers of an object are not determined by any objects not sharing spatiotemporal location with it

I suppose the defender of TU must bite the bullet and reject (CP2). But I do not see any great disadvantage in doing so. (CP2) does not strike me as being especially compelling.

I must grant that there is a sense of (CP) in which it may be counted as part of common sense. For example, the following seems true:

> My causal powers are not determined by objects "beyond" me. For example, my ability to lift this barbell is independent of the outcome of the Olympic weightlifting championship.[12]

But this intuition seems to be captured by the following principle, which principle may be accepted by the TU-ist:

(CP3) The causal powers of an object are completely determined by what *intrinsic* properties that thing has, together with the laws of nature—that is, two things with the same intrinsic properties in worlds with the same laws must have the same causal powers

Since I have my intrinsic properties just in virtue of the way I am in myself, and not in virtue of my relations to the competitors in the Olympic championship, this seems to capture the intuition in question.

To see that TU does not in any way preclude (CP3), recall the discussion of the Duplication Argument, in which it was argued that universals being transcendent in no way precludes the possibility of a (nontrivial) distinction between intrinsic and extrinsic properties.

The TU-ist should reject (CP1) and (CP2). Granted, there is an intuition that the causal powers of an individual x are determined by what x is like, *considered in itself*. But (CP3) vindicates this intuition. Thus, the TU-ist may do his duty to common sense by accepting (CP3), and may fairly reject (CP1) and (CP2).

[12] The relevant sense of 'determined' here is tricky. For example, I might become depressed upon learning of the outcome of the championship, and thereby become unable to lift the barbell, and in this sense the outcome of the championship can "affect" my ability to lift the barbell. But this would be because the outcome of the championship would *cause* certain changes in my intrinsic properties. My causal powers are *determined* (in the relevant sense) by my intrinsic properties, and not by the outcome of the championship. The precise intuition behind the example of the barbell is contained in (CP3).

The Subtraction Argument

Next we return to a passage from Armstrong considered above (1978*a*, p. 68):

> Is it not clear that *a*'s whiteness is not determined by *a*'s relationship with a transcendent entity? Perform the usual thought- experiment and consider *a* without the form of Whiteness. It seems obvious that *a* might still be white. So *a*'s being white is not determined by *a*'s relation to the Form.

Let *a* be any white thing. Armstrong seems to argue as follows:

Subtraction Argument

(1) Possibly, *a* is white but does not instantiate Whiteness

(2) if (1), then TU is false

(3) Therefore, TU is false

At first glance, the Subtraction Argument looks blatantly question begging. Premise (1) amounts to a flat denial of TU, for according to a TU-ist, to be white *just is* to instantiate Whiteness. Surely, a TU-ist would be within his or her rights in simply rejecting (1).

Can any progress be made? I think so. Armstrong seems to be relying implicitly on the fact that *whiteness* is an *intrinsic* property to justify (1).[13] The idea seems to be that since *whiteness* is intrinsic, *a*'s being white cannot depend on its relations to any object external to *a*. Thus, *a* could be white even if, say, nothing else existed. We can imagine Armstrong appealing to a principle of *isolation*:[14]

(I1) For any possible object x, there is a possible world w containing just a duplicate of x and its parts

[13]See note 11 for a caveat about the intrinsicality of *whiteness*. I thank Phillip Bricker for his help in interpreting Armstrong's justification of premise (1) in terms of recombination principles.

[14]There is parallel but weaker defense of (1) in terms of recombination principles. Rather than considering principles that generate worlds in which *whiteness* does not *exist*, we could consider principles that generate worlds containing duplicates of *a* that do not bear the instantiation relation to *whiteness*. The discussion of this approach would parallel that in the text.

The object guaranteed by (I1) would be white, being a duplicate of *a*, but the form (universal) Whiteness would not exist at that world, since according to TU, universals are not parts of their instances.

However, the defender of TU should feel no compunction over rejecting (I1). It is at the heart of TU to admit necessary connections between an object having a certain nature and it standing in certain relations to external entities. (Indeed, on some versions of TU universals enjoy necessary existence, and thus we would have an immediate violation of (I1).[15]) In place of (I1), a more moderate version may be offered. Let "IP(x)" denote x's intrinsic profile, where x is any object (see sections 4.2.2 and 4.2.3 for information on intrinsic profiles). For any object x, call the proposition that x has IP(x) "the nature of x".

(I2) For any possible object x, there is a possible world w containing only i) an object, y, that is a duplicate of x, ii) y's parts, and iii) any object whose existence is entailed by the existence and nature of an object mentioned in i) or ii).[16]

(I2) does not entail the existence of a world with a white object but no universal of *whiteness* because of clause iii): the existence of *whiteness* is entailed by the nature of the white object, according to TU.

I think that it is not ad hoc to accept (I2) rather than (I1). When a theory accepts the fact that, e.g. the whiteness of an object *consists* in its relation to an external entity (a property), then it is clear that all recombination principles will have to be restricted to take this into account. And this does not contradict common sense, since our reasoning about recombination begins by considering rearrangements of *particulars*, and thus only supports principles like (I2), *not* unrestricted principles like (I1).

Explanation of Resemblance

In this section I want to address an intuition that I think is common to the last three arguments. In each argument, Armstrong claimed that some fact could not be explained by TU. In the Duplication argument, the fact was

[15] We may have another, independent, challenge to (I1): one who believes in sets will hold that it is impossible for an object to exist without its unit set existing. Assuming that sets are never parts of non-sets (Lewis, 1991, p. 7), this will violate (I1).

[16] (I2) is adapted from Paull and Sider (1992).

that objects can differ intrinsically. In the causal powers argument, the fact was that the causal powers of an object are in some sense independent of objects "external" to that object. In the Subtraction argument, the fact was that the nature of an object is independent of things external to that object.

In each case, I argued that TU *can* account for these facts—through the distinctions that are the subject of the dissertation. Thus, in the case of the duplication argument, I argued that the TU-ist can say that ball *a* is *not* a duplicate of ball *b* because *a* has an *intrinsic* property that *b* lacks.

However, in each of these cases, we can imagine Armstrong responding as follows.

> In *name*, you can account for these facts. But your *explanation* of the facts is inferior because your universals are transcendent. You say that ball *a* has an intrinsic property that ball *b* lacks. But this property is not *in* ball *a*, on your view, whereas the property *is* in ball *a* on *my* view. My account is the more satisfying one.

I imagine Armstrong replying in similar fashion to my responses to the other arguments.

The response I imagine Armstrong making, then, concedes that TU has *some* way of accounting for the facts in question. But it holds that the account IU gives of these facts is *superior* to that of TU.

I will focus on the question of whether IU gives a superior account of what it is for objects to differ intrinsically. Let us remember the two separate components of the notion of immanence:

(P) universals are parts of their instances

and

(L) universals share spatiotemporal locations with their instances

I think we can see that (P) does *not* give immanent universals an edge over transcendent universals. Let us return to the hybrid theory HTU. Suppose that (P) represents a superiority of immanence over transcendence in the explanation of intrinsic difference. Since (P) is true on HTU, it would seem that this theory can also claim an explanatory advantage over TU. But surely HTU can claim no explanatory advantage over TU, for these theories do not differ metaphysically in the least.

To say that a universal is part of an instance is just to say that the instance encompasses the universal. This can be true if the universal "reaches into" the instance, as immanent universals do, but it can also be true if the instance "reaches out" to the universal. Surely, switching to the latter description gives us no edge in the explanation of intrinsic difference. But then (P) cannot represent any explanatory advantage, for (P) is accepted by HTU but rejected by TU.

If IU is superior in explaining intrinsic difference, then it must be because of (L). It is because immanent universal are *in* their instances, in a spatiotemporal sense, that they give a better account of intrinsic difference. Indeed, when I form a picture of immanent universals being wholly present in their instances, and this picture seems, offhand, to provide a superior account of intrinsic difference, I think what convinces me is the spatiotemporal coincidence between the universals and the instances in the picture.

However, I think we can form an argument that shows that (L) does not in fact represent explanatory superiority for immanent universals. Consider objects that exist outside of space. Disembodied souls might be examples. I suppose such things are possible. And surely two such objects could have identical total temporal locations. For such a pair of objects, the relation of duplication could presumably hold, but could also fail to hold. But these facts could not be explained by appeal to spatiotemporal coincidence between universals and instances, for such a pair would have identical spatiotemporal locations.

Consider a possible world w containing three disembodied souls: Moe, Larry, and Curly. Let us suppose that all three come into existence at exactly the same moment, and also go out of existence at exactly the same moment. Furthermore, suppose that Moe and Larry are duplicates, whereas Curly is not a duplicate of either. Let us see if a defender of IU is in a better position to explain these facts than a defender of TU. I think he is not. He can say that Moe and Larry instantiate various universals that Curly does not, but the defender of TU can say this as well. It is true that he can say that these universals are *parts* of Moe and Larry, whereas the defender of TU cannot say this. But I have already argued that this is no explanatory help. Finally, the IU-ist can *not* say that there are universals sharing spatiotemporal location with Moe and Larry, but not with Curly. All three have exactly the same total spatiotemporal location, since all three have the same temporal location and none of the three is located in space. Thus, any universal sharing spatiotemporal location with Moe and Larry also shares spatiotemporal location with Curly.

It may be denied that entities that lack spatial location are possible. But this seems bold.[17]

I think this shows that (L) gives immanent universals no explanatory advantage over transcendent universals. If (L) did represent explanatory advantage, then this would presumably be so of necessity. But I have presented a possible case in which IU does not have any explanatory advantage.

If immanent universals have no advantage in explaining facts of intrinsic difference, then why did they *seem* explanatorily superior?. In particular, why did (L) seem to give explanatory superiority?

Consider a transparent glass globe. Now imagine putting a bright light inside the globe. The light shines through. Putting the light *in* the globe causes the globe to take on a property of the light—luminescence. If the light were not *in*, spatiotemporally *in*, the globe, the globe would not be lit. I often imagine immanent universals this way—as lights that illuminate their instances from inside. Or I imagine them like the caloric fluid of 18th century physicists. These physicists thought that heat was a fluid that permeated hot things, and flowed from hot things to cold things. I imagine an immanent universal as a sort of fluid that permeates its instance causing it to have a certain nature. Were the fluid not in the instance, *in* in a spatiotemporal sense, the instance wouldn't be affected. Perhaps our intuition that immanent universals give a superior explanation of intrinsic difference derives from taking some picture like this too seriously.

The Argument from Causal Impotence

Finally, we have the following argument:[18]

> A spatio-temporal realm of particulars certainly exists (it includes our bodies.) Whether anything else exists is controversial. If any entities outside this realm are postulated, but it is stipulated further that they have no manner of causal action upon the particulars in this realm, then there is no compelling reason to postulate them. Occam's razor then enjoins us not to postulate them.

This argument does not merely apply to transcendent universals, but I will only consider this application of the argument. I construe the argument as

[17] Armstrong, for example, accepts the possibility of such entities (1978*a*, p. 119).

[18] Armstrong (1978*a*, p. 130). In fact, all of chapter 12 is relevant.

follows:

Argument from causal impotence

(1) We have no reason to postulate causally inert entities

(2) Transcendent universals are causally inert entities

(3) Therefore, we have no reason to postulate transcendent universals

Armstrong applies Occam's razor to (3) to conclude that we *should not* postulate transcendent universals, but (3) seems to me to be bad enough. Premise (1) is intended to be a plausible principle of epistemology. Armstrong argues for premise (2) at length. His reasons are complex and controversial, but they need not concern us.[19] What I have to say will sidestep the issues he raises.

The argument from causal impotence is not a direct objection to TU. It does not show TU to be incompatible with any commonly accepted principle, nor does it reveal any internal inconsistency in TU. The conclusion of the argument is not that TU is false. Still, I take it that the TU-ist would feel rather uncomfortable accepting the conclusion. The argument must be answered.

The argument from causal impotence is a familiar argument against postulating "abstract" entities of all kinds. An equally familiar response is that the theoretical benefits of postulating causally inert entities overrides any presumption against them. A common analogy: we can best understand the truth of mathematics by postulating a realm of sets (Putnam, 1971, especially chapters 5, 7, and 8). As it is legitimate to postulate sets to make sense of mathematics, so (the response goes) it is legitimate to postulate propositions, properties, transcendent universals, possible worlds, etc. to make sense of

[19]One argument is this. If a thing enters into causal relations, then it must change. Transcendent universals cannot change, so they cannot enter into causal relations (1978*a*, p. 128). If sound, this argument would seem to show that immanent universals cannot enter into causal relations either! For aren't immanent universals unchanging?

A second argument seems to be this. Traditionally, the notion of a God, or a Cartesian soul, acting on Nature has been problematical. If so, then the notion of transcendent universals entering into causal relations must be more problematical (1978*a*, p. 129–130). I do not find this argument very convincing. Moreover, both arguments depend on the concept of a non-event acting on another. See below.

various other data (Armstrong 1978*a*, pp. 128–132, 1989*a*, pp. 7–13; Lewis 1986*c*, pp. 3–5, 1991, pp. 57–59). Abstract objects are postulated for their non-causal explanatory value.

I have nothing new to add to this part of the debate. My only interest is in formulating a clear version of the argument that is not subject to *other* difficulties. Given such a version, I suppose I throw my lot in with those who appeal to the precedent of postulating sets. But first we must find such a version.

We must ask what it is for an entity to be *causally inert*. An intuitive picture of the *opposite* of causal inertness, causal potency, springs to mind. A billiard ball is a prime example of a causally potent object. Unlike numbers, transcendent forms, etc., billiard balls are capable of causing things (typically, movements of other billiard balls). We might think, then, that an object is causally inert iff it is incapable of entering into causal relations.

However, philosophers often hold that causal relations hold fundamentally between *events*. If a billiard ball's striking another billiard ball causes the other billiard ball to move, we might say that "the billiard ball caused the motion". But surely it was the event of the billiard ball's striking the other ball that did the causing.

It is true that we sometimes speak, for example, of a billiard ball causing something. We might call this "common sense" causation. But just as particulars enter into common sense causal relations, so do universals. I might say that heat caused my rash, that greed causes much of the suffering in the world, or that negative charge sometimes causes electromagnetic repulsion. So if common sense causation were the relevant sort of causation, then universals would turn out to be causally potent, and premise (2) of the argument would be false. Let us then consider only causal relations that hold between events.

We need a way to talk about events. For the moment, let's follow Jaegwon Kim in our talk about events. To an event *x's having P*, there corresponds an ordered pair $\langle x,P\rangle$—x an object, P a property (Kim, 1973, pp. 222–226). To an event *x and y's standing in R*, there is likewise an associated pair: $\langle\langle x,y\rangle,R\rangle$. Events, according to Kim, have "constitutive objects" and "constitutive attributes". The constitutive attribute of the first event is P; its constitutive object is x. The constitutive objects of the second event are x and y; its constitutive attribute is R. (Kim would include a third entity: the constitutive *time* —the time of the event's occurrence. Because of my assumption of temporally bound individuals, this is unnecessary). Let us call both the constitutive property and the constitutive objects of an event its *constituents*.

If only events stand in causal relations, then we cannot say that an object is causally inert iff it is incapable of standing in causal relations. On this definition, billiard balls would turn out causally inert! For myself, I feel compelled to postulate the existence of billiard balls, so let us instead consider the following definition:

(CI) x is *causally inert* $=_{\text{df}}$ It would be impossible for there to be an event E that that "involves" x and enters into a causal relation with some other event

This leaves us the problem of saying what it is for a event to "involve" an object.

It is natural to say that:

(I1) event E *involves* $x =_{\text{df}} x$ is a constituent of E

This definition has a result we want: that billiard balls are not causally inert. For suppose that event $E =$ *billiard ball a's striking billiard ball b* causes the event of *billiard ball b's moving*. Billiard ball a is involved in E, for a is a constituent object of E. Moreover, E causes something, and therefore billiard ball a is not causally inert. Surely, such a scenario is possible for any billiard ball. Thus, billiard balls are not causally inert, given the present definitions.

However, the argument is unsound, given this interpretation. Suppose, for example, that U is a universal and that a is a particular. Let event E be the event of *a's having U*. We may further suppose that E causes some other event. U is the constituent attribute of E. Thus, E involves U, and hence by (CI), U is *not* causally inert. Thus, premise (2) in the argument is false.

(I1) allows that when an event *a's having U* causes another event *b's having V*, in addition to the particulars a and b getting "credit" for a causal interaction, the universals U and V do as well. This might be thought to be improper. When a billiard ball's motion causes another billiard ball to move, it might be argued that we think of the *ball*, rather than its properties, as doing the causing. Thus, it might be argued that only the constitutive *objects* of events are causally potent in virtue of causal interactions, not the constitutive attributes.

To implement this intuition, we might revise our definition of 'involves' to read:

(I2) event E *involves* $x =_{\text{df}} x$ is a constituent *object* of E

On this definition, event E in the previous example would not involve U. Provided that universals are never constituent objects of events, events would never involve universals, and hence by (CI) all universals would be causally inert. (I will indeed assume for the sake of argument that universals are never constituent objects of events, although this might be legitimately challenged.) So premise (2) of the argument would be true.

However, this version of the argument from causal impotence is unconvincing. It has the feel of defining one's way to a desired conclusion. (I2) definitionally prohibits universals from being involved in events; therefore, (CI) definitionally makes universals turn out causally "inert".[20] This way of talking is fine, provided we are clear that we are considering a sense of 'causally inert' such that it is *definitional* that universals are causally inert. But now consider what premise (1) of the argument says: that we have no reason to postulate causally inert objects. Since the conclusion of the argument is that we have no reason to postulate transcendent universals, the argument begs the question. It includes a premise to the effect that we have no reason to postulate a certain kind K of entity, where it is true by definition that universals are not of kind K. The defender of universals will simply reject line (1) of the argument.

We need another definition of 'involves'. Let us shift to an alternate conception of events: as regions of spacetime.[21] This view of events gives us a neat definition of 'involves':[22]

(I3) event E involves $x =_{\text{df}} x$ is (wholly) located in E

Thus, the event of a particular billiard striking another would involve (the time slices of) the billiard balls.

On this definition, we get our desired result: transcendent universals are not involved in any events. Moreover, we get another desired result: immanent universals *can* be involved in events. This latter result is desirable

[20]I assume here that universals are essentially universals.

[21]See, for example, Lemmon (1967, pp. 98–99), and Lewis (1986*b*, note 4) for more references on this conception of events. Points analogous to those I make in the text would apply to Lewis's theory of events as *properties* of regions of spacetime.

[22]Phillip Bricker suggested defining 'involves' in terms of spatiotemporal location.

because Armstrong's argument is intended to apply to transcendent universals but not to his own theory. This last version of the argument, then, may be close to Armstrong's intention.

But I have a nagging doubt even about this final version of the argument. The argument seems to unfairly stack the deck against TU by its definition of 'involves'. Why not define the term as follows?:

(I4) event E involves $x =_{df} x$ is wholly located in E or x is instantiated by something wholly located in E.

In a clear sense, the properties of objects in a region of spacetime are involved in that region: they are instantiated there! Again, if an event, construed as a region of spacetime, causes something, it seems that the properties of the objects contained in that region are no less involved than the objects themselves. After all, the causal powers of such an event depend on the properties instantiated therein.

Thus, the argument from causal impotence that is based on (I3) seems faintly question-begging—I see no reason to accept (I3) over (I4), other than a prejudice against transcendent universals. We seem to have gone round and round in a vain search for a non question-begging version of the argument from causal impotence, and I have a diagnosis of this fact. The basic intuition behind the argument is that we have no reason to postulate non-spatiotemporal entities—the bit about causal impotence is an irrelevant detour.[23] Why not be forthright about this?:

Argument from Spatiotemporal non-Location

(1) We have no reason to postulate spatio-temporally unlocated entities

(2) Transcendent universals are spatio-temporally unlocated

(3) Therefore, we have no reason to postulate transcendent universals

Quite plainly, the discussion of *this* argument can focus directly on premise (1). As I have mentioned, I have nothing new to add to this discussion. I stand with those who appeal to the precedent of mathematics, and hold that sometimes it is correct to postulate non-spatiotemporal entities.

[23] I offer this diagnosis only in the present case of the argument applied to all transcendent universals. The notion of causal impotence may be important when the argument is offered with respect to possible worlds, numbers, and even perhaps uninstantiated properties.

6.4 Conclusion

The proposal to use universals to analyze naturalness is an attractive one, but it seems inferior in power to primitive naturalness since universals cannot be used to analyze the *more natural than* relation. My preferred view is abundant TU—an abundant conception of transcendent universals—plus a primitive *more natural than* relation. I argued that Armstrong's arguments against transcendent universals fail, and so provide no barrier to accepting this view.

Chapter 7

Dunn on Intrinsicality

In a recent paper Michael Dunn has criticized David Lewis's theory of intrinsicality and in its place put his own theory based on Relevance logic. Since I have accepted Lewis's analysis of intrinsicality in terms of duplication, Dunn's criticisms of Lewis are criticisms of me. I will argue that Dunn's objections are mistaken and that his own theory is uninteresting. Finally, I want to examine Dunn's contention that two traditional characterizations of intrinsicality come apart.

7.1 Dunn's Criticisms of Lewis

7.1.1 Dunn's Formulation of Lewis's View

As we have seen, Lewis defines 'intrinsic' as follows:

(D1) Property P is intrinsic iff for any possible objects x and y, if x and y are duplicates then x has P iff y has P

But Dunn construes Lewis's analysis of intrinsicality differently. Let ϕx be any formula with free occurrences of at most one variable x. Dunn interprets Lewis as claiming that ϕx is a formula *of a kind to determine an intrinsic property* iff the following statement is true (Dunn, 1990, p. 184):

(IPD*) $\phi a \rightarrow (x \approx a \rightarrow \phi x)$ (Indiscernibility of Perfect Duplicates)

where '$\approx$' stands for 'is a perfect duplicate of', a is a name, and ϕa is the result of substituting the name a for all free occurrences of x in ϕx. I discuss below how '→' is to be interpreted. The variable x is to be interpreted as implicitly universally quantified, whereas the name a, being a name, denotes a particular individual. Thus, (IPD*) seems awkward, since presumably the intent is that ϕx is of a kind to determine an intrinsic property if (IPD*) holds regardless of what a denotes.[1] Thus, I will reinterpret Dunn's construal of Lewis as follows: formula ϕx is of a kind to determine an intrinsic property iff

(IPD) $\forall x \forall y[\phi y \rightarrow (x \approx y \rightarrow \phi x)]$

is true, where ϕy is the result of substituting y for all free occurrences of x in ϕx. It is clear that Dunn does *not* intend the quantifiers here to be possibilist.[2]

Dunn construes Lewis as offering an account of what it is for a *formula* to be "of a kind to determine intrinsic properties". But Lewis's account is of intrinsic *properties*. How to move from the former to the latter? Say that property P is *determined by* ϕx iff P is the property denoted by the phrase: ⌜being an x such that ϕx⌝. We may suppose that Dunn is construing Lewis as follows: property P is intrinsic iff every formula ϕx that determines P is of a kind to determine intrinsic properties. This raises a host of questions (e.g. what about properties for which we have no predicates?), but let us set them aside—what I have to say will not depend on this.

So, Dunn interprets Lewis's account of intrinsicality via an account of what it is for a formula to be of a kind to determine intrinsic properties, which in turn is analyzed using (IPD). We now turn to the interpretation of the connective '→' in (IPD). Dunn says that he engages in some creative exegesis, and interprets '→' in various ways. After raising objections to these interpretations, he finally interprets it in the sense of relevance logic. Let us look at these allegedly unacceptable possibilities for '→'.

Dunn says (1990, pp. 184–5):

> ...it seems that if the "arrow" in (IPD) is the material conditional, then the definition allows that *Socrates being wise* is an intrinsic property of Reagan...Even if one employs strict implication, *Socrates being wise or not wise* ends up as an intrinsic property of Reagan.

[1] See Dunn (1987, p. 361), formula (7∀). This is Dunn's analog of Lewis's definition. Why he does not formulate Lewis's theory in a parallel fashion, I do not know.

[2] His Socrates example that I consider below makes this clear.

Let us examine Dunn's points. Dunn asks us to suppose that "*the* arrow" in (IPD) is a material conditional. Since there are two arrows, we may take him to be treating each arrow as a material conditional. Let 'F' express the property *being wise*, and let 'a' denote Socrates. Finally, let $\phi x =$ 'Fa'. Notice that this formula has no free occurrences of x, and hence $\phi v =$ 'Fa' for any variable v. This formula corresponds to the property *being such that Socrates is wise*. Clearly, this is not an intrinsic property. But under the present proposal, this property turns out intrinsic, for the sentence:

(1) $\forall x \forall y[Fa \supset (x \approx y \supset Fa)]$

(where the '$\supset$' is understood as expressing material implication) is a theorem of predicate logic. We should join Dunn in rejecting this definition of 'intrinsic'.

Next Dunn considers strict implication, or entailment. Again, I suppose he means to interpret *each* arrow as a strict conditional. That is, (IPD) is to be interpreted as:[3]

(IPD′) $\forall x \forall y[\phi y \Rightarrow (x \approx y \Rightarrow \phi x)]$

where $\ulcorner \alpha \Rightarrow \beta \urcorner$ is definitionally equivalent to $\ulcorner \Box(\alpha \supset \beta)\urcorner$. First, let us note that (IPD′) solves the problem above—that *being such that Socrates is wise* turned out intrinsic. Only if the following sentence is true does the present interpretation imply that *being such that Socrates is wise* is intrinsic:

(2) $\forall x \forall y[Fa \Rightarrow (x \approx y \Rightarrow Fa)]$

Since (2) is false, we do not get this damaging implication. Let x = me, and let y = my brother Mike. Presumably, there is a possible world w in which Socrates is wise (the actual world, for example), and another possible world w' in which i) Mike and I are perfect duplicates, but ii) Socrates is not wise. (Here I slip into talking as if Mike and I ourselves inhabit other possible worlds.) Assuming S5 modal logic, (2) is thereby falsified.

[3]Note: we must recall an issue from chapter 2. Here are using a language with quantification across the modal operator $\Box$. Given the usual possible worlds understanding of the meaning of $\Box$, this means that we must make sense of the notion of an object x having property P in a world in which x does not exist (since we are assuming the thesis of world-bound individuals). As I mentioned in chapter 2, I understand this in terms of counterpart theory: to say that x has P at w is to say that x has a counterpart at w that has P.

But (IPD′) is utterly implausible for reasons other than Dunn considers. According to (IPD′), almost *no* properties will turn out intrinsic. For example, the property *roundness* is intrinsic, on the present proposal, only if the following sentence is true (where the predicate 'R' expresses roundness):

(3) $\forall x \forall y[Ry \Rightarrow (x \approx y \Rightarrow Rx)]$

Unfortunately, (3) is false. Let y be a certain actual round tennis ball, and x be another actual tennis ball. Suppose that in world w, while neither y nor x is round, x and y are perfect duplicates. Again assuming S5 modal logic, (3) is false. Clearly, we could repeat this procedure for most properties commonly thought to be intrinsic.

This construal of Lewis's theory is a clear mistake—Lewis's own definition, (D1), has no defect of this kind. Moreover, there is no way to acceptably weaken (IPD′). The following two attempts seem to be the only possibilities. We could weaken the first or the second '$\Rightarrow$' to a '$\supset$':

(IPD$'_a$) $\forall x \forall y[\phi y \Rightarrow (x \approx y \supset \phi x)]$

(IPD$'_b$) $\forall x \forall y[\phi y \supset (x \approx y \Rightarrow \phi x)]$

The problem with (IPD$'_a$) is the same as the problem with taking each conditional in (IPD) to be material—*being such that Socrates is wise* comes out intrinsic. For the instance of (IPD$'_a$) in this case is:

$$\forall x \forall y[Fa \Rightarrow (x \approx y \supset Fa)]$$

that is,

$$\forall x \forall y \Box[Fa \supset (x \approx y \supset Fa)]$$

which is a theorem of modal predicate logic. The problem with (IPD$'_b$) is the same as the problem with (IPD′)—almost no properties turn out intrinsic. Consider, for example, the instance of (IPD$'_b$) when we let ϕy be 'Ry', with 'R' interpreted as meaning "is round":

$$\forall x \forall y[Ry \supset (x \approx y \Rightarrow Rx)]$$

This sentence turns out false for essentially the same reasons that (3) turned out false above. Even if y is in fact round, x's being a duplicate of y does not *entail* that x is round, since x could be a duplicate of y in a world in which y is not round.

I conclude, then, that Dunn's interpretation of Lewis along the lines of (IPD) using strict conditionals is a mistake—all possibilities for what Dunn could have had in mind are in trouble for reasons that Dunn never considers. What is the source of this mistake? Dunn obtains his interpretation, the schematic (IPD), from Lewis's words "if something has an intrinsic property, then so does any perfect duplicate of that thing" in Lewis's paper "Extrinsic Properties".[4] In that paper, Lewis is not clear that the quantifiers are possibilist; this seems to be the reason for the misunderstanding.

Fortunately, Dunn's major objections do not depend on his mistaken formulation of Lewis's view. The objections all apply equally well to (D1), which is what I take to be Lewis's actual theory. Let us then think of the objections as directed at (D1).

7.1.2 Dunn's First Objection

Let us return to the passage quoted above. Dunn says: "Even if one employs strict implication, [*being such that Socrates is wise or not wise*] ends up as an intrinsic property of Reagan" (1990, p. 185).

In general, if N is a necessary truth, then *being such that N* will be an intrinsic property, according to (D1). One might object to Dunn's example on the grounds that 'Socrates is wise or Socrates is not wise' does not express a necessary truth because of the possibility of Socrates failing to exist, so let us substitute the property *being such that* $2+2=4$. On Lewis's view, this is an intrinsic property.

There are two other related consequences that we might reasonably expect Dunn to find objectionable. First, on Lewis's theory, all impossible properties are intrinsic. Of course, no impossible property is an intrinsic property *of* anything, since no object can have an impossible property. Still, impossible properties can never differ between perfect duplicates, so they turn out intrinsic according to (D1).

Second, if P and Q are necessarily coextensive properties, then Lewis's

[4] Dunn quotes this on p. 184 of Dunn (1990); the quotation is from Lewis (1983*a*, p. 197).

view has the consequence that P is intrinsic iff Q is intrinsic. But, one might think, this is implausible: consider $P =$ *being green*, and $Q =$ *being green and being such that there is no largest prime number*. Let us summarize these consequences thus:

(L) Every necessary or impossible property is intrinsic, and if P and Q are necessarily coextensive properties, then either both or neither are intrinsic.

I simply do not find (L) objectionable, at least for the notion of intrinsicality with which I am concerned. I grant that there may be other notions of intrinsicality, and maybe (L) would be objectionable for those notions. But I do not grant that there is only one single notion of intrinsicality that we all have before our minds, which is such that (L) is false of it.

The way I see it, 'intrinsic' is a term of art. There are some non-negotiable constraints on how any notion deserving the name "intrinsic" must behave. For example, on any construal of intrinsicality, relational properties like *being within ten feet of a perfect sphere* should not turn out intrinsic. But I don't think that these non-negotiable constraints say anything about (L), one way or another. (L) is negotiable.

A good distinction to mark is that between those properties that can never differ between perfect duplicates, and those that can. This distinction deserves to be called a distinction between intrinsic and extrinsic properties. That (L) turns out true on this distinction seems to me to be of little import. What is more important is whether this distinction can do philosophical work. I think it can, and what is more, I don't think that (L) impedes this work in the least. So, I regard (L) as a "throwaway" consequence of the theory.

7.1.3 Dunn's Second Objection

After dismissing the interpretation of (IPD) as involving a strict conditional, Dunn suggests interpreting the arrows in (IPD) as being those of *relevant implication*. This blocks both of the problem cases from Dunn that we discussed: neither *being such that Socrates is wise* nor *being such that* $2+2=4$ turn out intrinsic on this view. This is because of special features of relevance logic that Dunn discusses informally as follows (1990, pp. 180–181):

> First, the antecedent of a relevant implication is supposed to be a really sufficient condition; it, *all by itself*, is supposed to be sufficient for the consequent. There should not be the slightest hint of background or *ceteris paribus* conditions for a true relevant conditional, unlike the case with the Lewis-Stalnaker analysis of the so-called “counterfactual conditional”. There should be no suppression of “premisses” merely because they are true…
>
> Second, the consequent of a relevant implication is supposed to depend on the antecedent in a somewhat technical sense, but one that intuitively means that the antecedent can be used in deriving the consequent.

But Dunn has an objection even for his final construal of Lewis’s view ((IPD) with the arrows construed as relevant conditionals). Let b be some actual black marble. According to Dunn, the property of *being a perfect duplicate of b* turns out to be an intrinsic property, on his final construal of Lewis’s view. For this property corresponds to formula ‘$x \approx b$’, and the following formula is derivable in relevance logic from the uncontroversial assumption that duplication is an equivalence relation:

(4) $y \approx b \rightarrow (x \approx y \rightarrow x \approx b)$

Dunn finds this consequence objectionable. He says that the answer to the question of whether a given object a has this property “Clearly…does not depend on a alone, but equally depends on b and its intrinsic properties.” (1990, p. 185).

This argument appears to apply equally well to (D1). The property *being a perfect duplicate of b* can never differ between perfect duplicates, for if a and c are duplicates of b, then they must be duplicates of each other (again, we use the fact that duplication is an equivalence relation). So I will discuss the argument as an argument against what really is Lewis’s theory: (D1). The argument seems to be:

i) If Lewis’s theory is true, *being a perfect duplicate of b* is an intrinsic property

ii) *being a perfect duplicate of b* is not an intrinsic property

iii) therefore, Lewis’s theory is not true

Fortunately, defenders of (D1) have nothing to fear from the argument, for the argument equivocates. The ambiguous phrase is '*being a perfect duplicate of b*'. Recall that I assume the thesis of worldbound individuals, so *b* exists in only one possible world. Here are two possible properties this phrase could denote:

P_1: the property had by an object *x* iff *x* and the counterpart of *b* in *x*'s world are duplicates

P_2: the property had by an object *x* iff *x* and *b* itself, the object in the actual world, are duplicates

To know whether an object in world *w* has P_1, I need to compare it to the counterpart of *b* at *w*. To know whether an object at some world has P_2, we should compare it to *b* itself (back in the actual world).

It might be thought that it is implausible that P_1 is the property Dunn had in mind. But this is not implausible at all. In fact, I think that P_1 is the most natural interpretation of the phrase '*being a perfect duplicate of b*'. This is just another example of the counterpart theorist's interpretation of *de re* modal claims. For example, when we discuss what would have happened if Dukakis had won the election, we will discuss what would have happened to *Dukakis*, despite the fact that it is only a counterpart of the Duke in the relevant world. If we discuss the property of *losing an election to Dukakis*, we will attribute that property to whoever "Dukakis beats" in this possible world, despite the fact that it isn't the Duke himself that does the beating.

In everyday English, we do not mention the counterparts of Dukakis. Rather, we discuss what happens to Dukakis himself in counterfactual situations. We say: Dukakis might have won. In this counterfactual situation, Dukakis did win. The counterpart theorist does not deny the truth of these claims, but he does give an analysis of them on which their truth is consistent with Dukakis himself being present in only one world. This analysis is not given in English; it is in a language that, we might say, is more literal. In this language, we mention Dukakis's counterparts. In this language, it is false to say that Dukakis himself wins in other possible worlds—rather, only his counterparts do. Let's call this language the "possibilist" language.

Whether '*being a duplicate of b*' refers to P_1 or P_2 depends on whether this phrase is taken as a phrase of English, or as a phrase in the possibilist language. If we take it as a phrase of English, then it refers to P_1, just as "*losing an election to Dukakis*" refers to a property had by people beaten by *counterparts* of Dukakis. On the other hand, if it is interpreted as a phrase of the possibilist language, it denotes P_2. Let's consider how the argument fares on each reading.

If the phrase denotes P_1, premise i) is false. Object b, recall, is a black marble in the actual world. Consider a world w that contains a marble, c, that is a duplicate of b, and also b', the counterpart of b at w. Suppose that b' is white. First note that b has P_1, for b is a duplicate and counterpart of itself.[5] But c does not have P_1, for b's counterpart at w, b', is white whereas c is black. Thus, P_1 differs between duplicates (b and c) and hence is not intrinsic according to (D1).

On the other hand, if '*being a duplicate of b*' denotes P_2, then premise i) seems true. *Duplication* is symmetric and transitive, so if a and c are duplicates of b, then they are duplicates of each other. But premise ii) can now be rejected.

Let P be the conjunction of all of b's intrinsic properties. P is an intrinsic property (see section 4.2.3, principle A3). I will show that P is necessarily coextensive to P_2; (L) then implies that P_2 is intrinsic. So, any objection to this reasoning would need to be an objection to (L), and we have already considered such objections in the last section.

The argument here employs principles about *intrinsic profiles* from sections 4.2.2 and 4.2.3. Suppose Px. P is an intrinsic profile (principle (C3)); since b and x share intrinsic profiles they are duplicates (principle (C2)), hence P_2x. On the other hand, suppose P_2x. Thus, x and b are duplicates (since *duplication* is an equivalence relation); since P is intrinsic, by (D1), we have Px. Thus, P and P_2 are necessarily coextensive.

I think that whatever initial plausibility premise ii) enjoys is the result of thinking of '*being a perfect duplicate of b*' as denoting P_1, for this reading implies that there is something special about b that makes a given object a have the property. Look again at what Dunn says in support of premise ii) (1990, p. 185):

> Consider the question of whether a given object a is a "perfect duplicate" of an object b… Clearly the answer to this question does not *depend* on a alone, but equally depends on b and its intrinsic properties. [my emphasis]

[5] We may stipulate that b has no other counterparts in the actual world.

Dunn does not say what sense of 'depend' he intends here, but on a natural reading, what he says is true only if '*being a perfect duplicate of b*' denotes P_1 rather than P_2. This sense is that of counterfactual dependence. For suppose that *a* and *b* exist in the actual world, and are duplicates. Object *a* has the property *being a perfect duplicate of b* regardless of whether we read this phrase as denoting P_1 or P_2. But now let us ask whether (5) is true:

(5) if *b* were to change color but *a* remained unchanged, *a* would no longer have the property *being a perfect duplicate of b*

(5) expresses the claim that *a*'s having the property *being a perfect duplicate of b* counterfactually depends on the color of *b*. In fact, (5) is true if the property in question is P_1, false if the property in question is P_2. So if Dunn has in mind counterfactual dependence when he says 'depends', his defense of premise (ii) rests on reading '*being a perfect duplicate of b*' as denoting P_1. I grant the truth of premise (ii) on that reading, but on that reading premise (i) is false, as I have already shown.

Premise ii) sounds true, for we usually speak English, not the possibilist language. But if we fix on the possibilist language, our intuitions must be distrusted. Consider the property *having the same color as b*. In the possibilist language, the property we express is necessarily coextensive to *blackness*, since *b* is in fact black. But of course *usually* when we use the phrase '*having the same color as b*', we usually have in mind the property had by object *x* in world *w* iff *x* has the color of the counterpart of *b* at *w*. To get the other reading in English, we would have to say '*having the same color as b has in fact*'. This phrase expresses in English what '*having the same color as b*' expresses in the possibilist language.

7.2 Dunn's Theory of Intrinsicality

Next I turn to Dunn's own definition of 'intrinsic'. Dunn offers two definitions:[6]

[6]For the first, see Dunn (1990, p. 180). Dunn says for the left hand side of this definition "*a* has ϕ intrinsically" rather than "*a* has the property determined by ϕ intrinsically", but ϕ is assigned a formula, not a property. For the second definition see Dunn (1990, p.185) formula (II). As before, I substitute a universally quantified variable instead of the name '*a*' that appears in Dunn's paper.

a has the property determined by ϕ intrinsically $=_{\text{df}} \forall x(x = a \to \phi x)$

ϕ is of a kind to determine relevant properties $=_{\text{df}} \forall x \forall y[\phi y \to (x = y \to \phi x)]$

where the "arrows" are those of relevant implication.

Say that property P is a *relevant* property iff every formula that determines that property is of a kind to determine relevant properties. Dunn's theory of intrinsicality seems to be this: a property is intrinsic iff it is relevant. The notion of a property had intrinsically by an object a is slightly different, for he allows that, if *wisdom* is an intrinsic property of Socrates, then he has *wisdom or being such that* $2+2=4$ intrinsically, since he has the latter property purely in virtue of himself. But the latter property would not be a relevant property (1987, p. 363, 1990, p. 183).

(Notice that we can offer a parallel analysis of the notion of "object-relative" intrinsicality within the Lewis framework:

(RI) a has P intrinsically $=_{\text{df}}$ every possible object that is a perfect duplicate of a has P

Notice also that a property is intrinsic *simpliciter* (as defined by (D1)) iff it is intrinsic to every possible object that has it.[7])

We should be clear that Dunn's theory is a theory of a different sense of 'intrinsic' than the one in which I have been interested in this dissertation. One of the consequences of Dunn's analysis that he favors is that haecceities are intrinsic (Dunn, 1990, p. 186). He defends this position with a quotation from G.E. Moore (1922, p. 262):

> It is obvious that there is a sense in which when two things are exactly alike, they must be 'intrinsically different' and have different intrinsic properties, merely because they are two...the mere fact that they are *numerically* different does in a sense constitute an intrinsic difference between them, and each will have at least one intrinsic property which the other has not got - namely that of being identical with itself.

Now, in that paper, Moore is merely pointing out that there is a *sense* of 'intrinsic' according to which the property *being identical to Ted* is intrinsic.

[7]In conversation Lynne Rudder Baker has suggested to me that some uses of 'intrinsic' in everyday language involve an object-relative notion of intrinsicality.

I am perfectly willing to grant that there is such a sense. Perhaps Dunn has given a correct theory of it. The sense of 'intrinsic' in this dissertation, Lewis's sense of 'intrinsic', might be called a *qualitative* sense of 'intrinsic', and Dunn's account fails to capture *this* sense of 'intrinsic'. Of course, this is no objection to Dunn's theory as a theory of some other sense of 'intrinsic'.

Whatever sense of 'intrinsic' Dunn is interested in, his theory is in the end of little importance to the project of this dissertation—it is of no help to the metaphysician who seeks a reductive definition of 'intrinsic'. His definitions are stated for a formal language; the definitions involve formulas of that language. But we need to choose primitive predicates of that language. If the primitive predicates are chosen to express extrinsic properties, then the definition will yield the result that these are relevant properties. So his definitions are of no help in distinguishing the intrinsic properties unless we already understand the distinction.

Dunn knows this. Consider the following quotations:

> These observations seem finally to constitute a definition of intrinsic property, at least for an ideal language where complex relational ideas are not expressed deceptively by monadic predicates (1990, p. 202).
>
> But it is not the business of logic, but rather of metaphysics (or perhaps of whatever field whose subject matter is being formalized, e.g., physics) to determine what formulas "really" determine properties... logic should tell us only that *if* certain formulas are postulated to "really" determine properties, *then it follows* that certain other formulas "really" determine properties" (1987, p. 355).

But Lewis is, and I am, of course, engaged in metaphysics, not logic. Why Dunn offers his theory as a competitor to Lewis's theory I cannot fathom.

7.3 Two Conceptions of Intrinsicality

There are several traditional ways of distinguishing intrinsic from extrinsic properties. Dunn claims that these do not always yield mutually consistent results. In particular, he mentions the "metaphysical" and "syntactical" criteria (1990, p. 178):

> Metaphysically, an intrinsic property of an object is a property that the object has by virtue of itself, depending on no other thing... An-

> other common way of characterizing the intrinsic properties of an object (let us call it "the syntactical criterion") is to say that they are non-relational.

According to Dunn, these two criteria do not always agree because of the possibility of "a relation that an item *a* has to an item *b*, but which depends in some sense on only *a* itself..." Dunn's example is a non-Humean notion of causality; he quotes Kripke as follows (Dunn, 1990, p. 179):

> Indeed to say that *a* by itself is a sufficient cause of *b* is to say that had the rest of the universe been removed *a* still would have produced *b*.

Consider the property *causing b*. It seems like a relational property, and hence an extrinsic property according to the syntactical criterion. But, if we accept this view of causation, we seem to want to say that *a* has the property of *causing b* "purely by virtue of itself". The metaphysical criteria then would call this property intrinsic.

I do not accept these traditional conceptions as conceptions of the notion of intrinsicality of this dissertation. The syntactic criterion seems to me to count *haecceities*, properties like *being Ted*, as intrinsic, since these do not seem relational (in a sense of 'relational' that might be elaborated as 'involving relations to distinct things'). The metaphysical criterion is close to the criterion I accept, but I dislike it for two reasons. The first is that it, too, seems to make haecceities intrinsic. The second is that some properties are extrinsic and yet are sometimes "had by an object by virtue of itself...". Consider a green object that has the property *being green or being 10 feet from some red thing*. The related criterion I accept is that an intrinsic property is one such that *whenever* it is had by an object, it is had just in virtue of the *qualitative way* that object is in itself. But I do want to consider and reject Dunn's example, for I think it illustrates a misunderstanding of the notion of an intrinsic property.

I find the example unconvincing for two reasons. The first is that Dunn seems to have misunderstood the intent of the metaphysical criterion. When we say that an object has an intrinsic property "by virtue of itself", this is intended to have strong modal force. Suppose that my father is extremely dignified. Because of his stern demeanor, he has the property of *being respected by me*. In a sense, this is in virtue of himself, since he is so dignified. But of course, it is only because of certain facts about me as well that his dignity inspires my respect. It would be *possible* for him to remain as dignified as he in fact is, and yet for me to disrespect him.

In the case where *a* causes *b* there is an analogy. The non-Humean will claim that, in *some* sense, this occurs purely by virtue of *a*'s nature. But she will presumably not claim that it would be metaphysically *impossible* for *a* to occur without causing *b* (if she does, then I find that notion of causation implausible). In rough form, since *a* might have occurred without causing *b*, *causing b* isn't an intrinsic property.

Of course, this isn't right. If this were the argument, I would be appealing to the principle that an *accidental* property of an object can't be intrinsic, and this is clearly false. One of my intrinsic properties is my mass, but I could have had some other mass. What I mean to argue is that it is possible that *a* occur, *with the same intrinsic nature as it has in fact*, and yet not cause *b*, and that this fact implies that *causing b* isn't an intrinsic property of *a*. That is, there is a possible world in which (a counterpart of) *a* is a duplicate of its actual self but does not cause *b*.

Now of course, this looks question begging. For *a* to have the same "intrinsic nature" as it has in fact is for it to have the same intrinsic properties as it has in fact. So I seem to be directly asserting as a premise that *causing b* isn't an intrinsic property.

I am. The claim that *causing b* is intrinsic, as I see it, is so clearly and basically wrong that arguments against it are bound to beg the question. But I am hoping that the way I am denying it will jar the reader into fixing on the concept of an intrinsic property. Think about my claim that *a* might been a perfect duplicate of its actual self, without causing *b*. Isn't that clearly true? Imagine a world where *a* causes *b*. Now remove *b*. You didn't *have* to change *a* did you?

Perhaps *a* causes *b* in any world with the same laws as our world—this may be what is indicated in the Kripke quotation by the counterfactual locution "…*had* the rest of the universe been removed *a* still *would have* produced *b*." [my emphasis] This may indicate a belief that particular instances of causation are independent of matters of particular fact, so long as the causal laws are held constant. Or, perhaps the non-Humean believes that it is possible for *a* to cause *b* even when there are no other objects or laws. But surely it is not the case that in every metaphysically possible world in which *a* occurs and has its actual nature, *b* occurs as well.

My second objection to Dunn's example is similar to the first. Surely the non-Humean would allow that even though *a* in fact causes *b*, it is not necessary that it cause *b* in particular. Couldn't *a* have had its actual nature, but caused some other event instead? Perhaps even a duplicate event. For example, Bob's swinging his fist caused Rob's pain. Now imagine a world where the person Bob hit was not Rob, but someone else. Then, Bob's swinging his fist didn't cause Rob's pain; it caused, say, Nob's pain.

In the end, one with a radically non-Humean view of causation can consistently maintain that *causing b* is intrinsic. But she would be committed to theses that, I say, are rather unintuitive. In particular, there would be metaphysically necessary connections between states of affairs like

> *a's having such and such an intrinsic nature*

and

> *b's occurring*

So, the radical non-Humean would be committed to severe restrictions on possibility, for we normally think that such distinct states of affairs are metaphysically compossible, if not physically compossible. Since I find this implausible, I find Dunn's example implausible.

Chapter 8

Analysis of the Notions

In this chapter I discuss two related questions. The first is the question of whether certain attempts to analyze naturalness succeed, and the second is the question of whether certain attempts to analyze intrinsicality succeed. My answer in each case is *no*. In chapter 6 I considered and rejected the proposal that naturalness be analyzed using a sparse theory of universals. Obviously, it is impossible to review all proposals for analyzing the notions of intrinsicality, naturalness, and duplication. Still, I hope that the results of these two chapters, together with my claim that naturalness is an important and fruitful notion, justify my taking naturalness as a primitive.[1]

8.1 Can We Define 'Natural'?

The notion of naturalness that is up for analysis is that characterized in chapter 3 as "Conception 1". First, let's focus on the *perfectly* natural properties and relations. I characterized these as the "most fundamental properties". The general idea underlying the hope that naturalness can be analyzed is this: once we fix the facts about perfectly natural properties and relations, we thereby fix all other facts *of a certain sort*.

[1]Lewis (1986*c*, p. 63) discusses attempts to define 'natural' in terms of "robust" notions such as laws of nature and resemblance. His objection is that naturalness should be used to analyze these robust notions, rather than the other way around. Quinton (1958, pp. 53-58) discusses what he calls "formalistic" analyses of naturalness—those that do not appeal to any such robust notions. Since these proposals have already been discussed, I will focus on proposals of different kinds.

It must be emphasized that the qualification "of a certain sort" is necessary. The perfectly natural properties and relations do not form a supervenience base for all properties whatsoever. In section 3.2.1 I showed that haecceities do not supervene on the set of perfectly natural properties and relations. Moreover, perhaps there are possible worlds that are alike in their distribution of perfectly natural properties and relations, but differ with respect to the laws of nature that hold there, with respect to the causal relations that hold there, or with respect to the objective chances of various events there.[2] If so, then properties and relations involving laws, causation, or objective chance will be instantiated differently in the worlds in question, and so won't supervene on the perfectly natural properties and relations.

Thus, we have the requirement that all facts *of a certain sort* supervene on the perfectly natural properties and relations. What sort of facts? *Qualitative* facts. For the kinds of facts just mentioned that fail to supervene on the set of perfectly natural properties and relations seem non-qualitative.

Suppose we were to succeed in defining 'natural' in terms of 'qualitative'. This would not be an achievement quite so momentous as that of defining 'natural' purely in terms of supervenience and other "quasi-logical" notions. The notion of a qualitative property or relation seems to be of a piece with that of naturalness. Still, such a definition would be important. But I will argue that 'natural' cannot be so defined.

The working idea here is that the set N of perfectly natural properties and relations at any world w forms a supervenience base for the set Q of qualitative properties at that world—N is a "Q-base" for w.[3] But not just any Q-base for w is the set of perfectly natural properties at w. Any set is a supervenience base for itself. Hence, the following definition:

(N1) Property or relation P is perfectly natural iff it is a member of a Q-base for some world

would err in making every qualitative property perfectly natural. In fact, (N1) makes every property and relation perfectly natural since the set of *all* properties and relations is a supervenience base for any set of properties whatsoever.

[2] See Lewis's discussion of "Humean Supervenience" in the introduction to Lewis (1986*d*).

[3] In fact, because of the possibility of a "world of endless complexity", these remarks need to be complicated, but I will ignore this twist for simplicity. See section 3.2.2. The argument I give here applies equally well to the more complex versions.

(N1) makes too many properties natural. We should therefore limit our attention to certain special Q-bases—the properties in these Q-bases will be the natural properties. Intuitively, these will be the "non-redundant" Q-bases. But there are problems in carrying this out.

First, there is a problem with stating an appropriate necessary condition for perfect naturalness along these lines. The following necessary condition was suggested in section 3.2.1:

(1) for any world w, the set of perfectly natural properties and relations at w is a minimal Q-base for w.

Let us consider this extended into a definition:

(N2) Property or relation P is perfectly natural iff it is a member of a minimal Q-base for some world w.

where a *minimal* Q-base for w is defined to be a Q-base for w that has no other Q-base for w as a subset. The difficulty with the necessary condition for perfect naturalness contained in (N2) involves the "problem of minimality" I discussed in section 3.2.1. What if one perfectly natural property or relation can be analyzed in terms of others? Then, there may be a world at which the set of perfectly natural properties is not a *minimal* Q-base.

I suggested a solution to a special case of this problem—the problem of "permutations"—by revising the definition of 'minimal' to allow that a relation R and a permutation of R might both be present in a minimal supervenience base. However, it still seems an open possibility that some other instance of the problem of minimality might falsify (1). To the extent that this possibility is a genuine possibility, (1) is to be distrusted.

Whether or not the problem of minimality can be solved, I think it is clear that the *sufficient* condition laid out by (N2) is unacceptable.[4]

Let N be the set of perfectly natural properties and relations at some world w; let A be the set of qualitative properties and relations at w; suppose that N contains some property P. In section 3.2.1 it was argued that the negation of a perfectly natural property is not perfectly natural, so $\sim P \notin N$. Let N' be the result of replacing P by $\sim P$ in N. (N2) implies that N is a minimal Q-base for w, and we will show that it then follows that N' is as well. (N2) then implies that $\sim P$ is perfectly natural, and is therefore refuted.

We first note the following lemma:

[4] Phillip Bricker made an important suggestion here.

(*) for any sets A and B and property P, B supervenes on $A \cup \{P\}$ iff B supervenes on $A \cup \{\sim P\}$

which follows from principle (S5) of section 3.2.1. Now, since A supervenes on N, by (*) it also supervenes on N'. Moreover, N' is a *minimal* supervenience base for A. For suppose $N'' \subset N'$ and A supervenes on N''. If $\sim P \notin N''$ then $N'' \subset N$, and we contradict the fact that N was a minimal supervenience base for A. On the other hand, if $\sim P \in N''$, then by (*), A supervenes on $(N'' - \{\sim P\}) \cup \{P\}$. Since this latter set is a proper subset of N, we again have a violation of N's minimality.

It should be noticed that this example refutes a strengthened version of (N2):

(N3) Property or relation P is perfectly natural iff it is intrinsic and a member of a minimal Q-base for some world w.

for the property in question, $\sim P$, is intrinsic ($\sim P$ is intrinsic since P is perfectly natural and hence intrinsic—see section 4.2.3 where it is shown that Boolean combinations of intrinsic properties are intrinsic and that perfectly natural properties are intrinsic). It should also be noticed that we cannot block this example by banning "negations" from the minimal Q-bases we consider, for *every* property P is the negation of $\sim P$.

So, the definition we have been considering cannot account for the fact that the negation of a perfectly natural property is not perfectly natural. The trouble is caused by the fact that the negation of a property is, to put it colorfully, as good as the property itself as far as supervenience is concerned. I do not see how any definition of 'natural' along the lines we have been considering can get around this problem. The prospects, then, for defining 'natural' in terms of 'qualitative' and 'supervenience' look dim.

8.2 Can We Define 'Intrinsic?

A property is *intrinsic* iff it never differs between any two possible duplicates. Objects are *duplicates* iff they share all their intrinsic properties. We have a circle of interdefinability between 'intrinsic' and 'duplicate'. Given either, we may define the other (and given naturalness, we may define them both). But this may seem like defining the obscure in terms of the obscure. Who would understand 'duplicate' if she did not understand 'intrinsic'? I myself think

I understand both notions. I also think that the equivalences relating these concepts increase my understanding of each. But some philosophers have desired more than mere equivalences relating these two notions. They have sought reductive definitions of 'intrinsic' and 'duplicate' in terms of modal concepts, property exemplification, the part-whole relation, etc.

I believe that all such attempts must fail. But my goal in this section is more modest. I will review two proposed sets of definitions, one by Jaegwon Kim and one by Michael Slote. Each proposal will be shown to be flawed. In Slote's case I will discuss extensively the possibilities for revising the analysis; it will be seen that the prospects are not good.

8.2.1 Preliminaries

In this section I will use both of the languages I mentioned in chapter 2. I will engage in some use-mention sloppiness in the interest of smooth exposition.

The *actualist language* is a standard modal language with the ordinary apparatus of the predicate calculus, plus the modal operators '$\Box$' and '$\Diamond$'. Quantifiers that range over objects range over objects in "the world of evaluation". For example, "$\exists xFx$" is true at a world iff some object is F at that world. "$\Diamond\forall xFx$" is true iff it would be possible for every object that would then be actual to be F; that is, iff there is some possible world such that every object that exists *at that possible world* is F. This language has a temporal existence predicate—"Exist(x,t)" means that x exists at time t.

The *possibilist language* is my usual language. It contains no modal operators; in their place we have quantifiers that range over possible worlds (variables: 'w', 'w'', etc.), and quantifiers that range over all possible objects (variables: 'x', 'y', etc.). The sentence "$\forall xFx$" says that *every* object, *both actual and possible*, is F.

I will also employ "lambda abstraction". For example, where ϕ is some formula with no occurrences of free variables other than 'x' and 't', the expression $\ulcorner\lambda x\lambda t\phi\urcorner$ shall be understood as denoting the property had by possible individual z at time t' iff ϕ is true when z is assigned to 'x' and t' is assigned to 't'.

In this chapter only I will speak of properties of objects *at* times rather than the properties had simpliciter by the temporal slices of those objects. I do this to match the language of the philosophers I discuss. To this end, each language contains variables that range over *times*: 't', 't'', etc. The sentence "$t > t'$" means that time t is *after* time t'. Where 'P' names a property, I will take the liberty of using "$P(x,t)$" to mean that object x has property P at time t; "$R(x,y,t)$" means that x and y stand in relation R at t. I will also allow quantifiers over *places*, with corresponding variables 'p', 'p'', etc.

8.2.2 Kim's Definition

In discussing Kim's definition we may be brief, since it has already been adequately discussed by David Lewis (1983*a*). Kim (1982, pp. 59–60) offers a sequence of definitions building on a suggestion by Chisholm (1976, p. 127).

I follow Kim in using an actualist modal language. We need the concept of two objects being "wholly distinct". An object is not wholly distinct from other objects that it overlaps; neither is it wholly distinct from, say, its unit set. I will say x and y are wholly distinct at time t ("Dist(x,y,t)") iff a) both are contingent objects and b) x and y have no parts in common at t, and c) neither is such that its existence entails the other's existence.

Here are the definitions; the first is Chisholm's; the others are Kim's:

D1: G is *rooted outside times at which it is had* $=_{\text{df}} \Box\forall x\forall t[G(x,t)\rightarrow\exists t'(t\neq t'\wedge \text{Exist}(x,t'))]$

D2: G is *rooted outside the objects that have it* $=_{\text{df}} \Box\forall x\forall t[G(x,t)\rightarrow\exists y\exists t'\ \text{Dist}(x,y,t')]$

D3: G is *internal* (i.e. intrinsic) $=_{\text{df}}$ G is neither rooted outside times at which it is had nor outside the objects that have it

David Lewis, in 'Extrinsic Properties', notes that Kim's analysis is unsuccessful. The property of *loneliness*, had by x at t iff at t, there exists no contingent object that is wholly distinct from x, satisfies D3. But loneliness is not intrinsic. Imagine two possible worlds w and w' containing duplicate black balls at some time t. In w the ball is entirely isolated, whereas the ball has plenty of company in w'. Only the first ball is lonely, so *loneliness* can differ between perfect duplicates and hence is not intrinsic.

Similarly, many disjunctive properties incorrectly turn out intrinsic according to Kim's definitions. As Lewis notes, the property that is the disjunction of loneliness and the property of coexisting with exactly six pigs (wholly distinct from oneself) also satisfies those definitions.

Let us leave Kim's definition, and discuss instead some definitions proposed by Michael Slote.

8.2.3 Slote's Project

Slote does not directly concern himself with the project of defining either 'duplicate' or 'intrinsic'. But he does attempt to define certain related notions in chapter 8 of *Metaphysics and Essence*. I shall critically examine these definitions, as well as the question of whether they can be extended to define 'duplicate' and 'intrinsic'.

Slote's goal is to define the locution 'objects x and y are exactly alike at time t'', where x and y are understood to be objects in the same possible world. The definition of this locution employs the concept of an *alteration*; this also is analyzed by Slote.

The definitions have an idiosyncratic form, for they do not mention properties or related entities. I will take the liberty of changing their form to one more standard. Most of what I say about the definitions I discuss applies straightforwardly to Slote's original definitions.

8.2.4 Slote's Analysis of Alteration

Slote begins with an analysis of what it is for a single object x to be at some time intrinsically unlike the way it was at some earlier time. This Slote calls "alteration". This term is somewhat unfortunate, for on at least one natural use of this word, a ball that changed from being green at t to being blue and then back to being green by t' could be said to have "altered" between t and t', even if it was exactly similar at t to how it was at t'. It should be kept in mind, therefore, that such a case would *not* count as alteration in Slote's sense. In other terminology, an object x *alters* in Slote's sense between t and t' iff it is not the case that the stage of x at t is a perfect duplicate of the stage of x at t'.[5]

I take it that we have a pre-theoretical grasp of the concept of alteration against which proposed philosophical analyses may be evaluated; likewise for the concept of *alikeness* in the next section. To fix on that notion, I will sometimes appeal to the familiar notions of intrinsicality and duplication in my discussion of Slote's definition. Of course, these notions will not be used in the definitions. That would be circular.

Slote's definition runs as follows (1975, p. 138).

[5] I will take this to entail that x's passing into or out of existence counts as an alteration in x.

> ⌜ a alters between t_0 and t_1⌝ means about the same as ⌜there are an x identical with a, a t identical with t_0, and a t' identical with t_1 such that there are a y at t and a z at t' that are incompatible and both dependent on x; y and z are both not temporally bound; for all places p and times t'''', x, y, and z can each exist at p and t'''' and y and z can exist even if x never moves; and neither y nor z depends on anything c such that of necessity if one of y or z exists at some time t'' and the other at some other t''' and c exists at both t'' and t''' then either x or c is somewhere at t''' where it is not at t''⌝

When Slote says that object x is *temporally bound*, he means that x is either (a) essentially unceasing, or (b) such that if it exists at some time t then it essentially exists only at t (1975, p. 101). Object x *depends on* object y iff necessarily, if x exists at some time then y exists at that time. Objects are *incompatible* iff it would be impossible for them to exist at the same time.

Slote's idea is that this definition will be satisfied in virtue of the variables 'y' and 'z' being assigned certain **states of affairs**. Suppose that x undergoes a genuine change—say a change in color from red to blue. Then the definition will be satisfied in virtue of the states of affairs *x's being red* and *x's being blue*, respectively, being assigned to the variables 'y' and 'z'. On the other hand, suppose x merely changes its position without changing any of its intrinsic properties. Then the assignment of the states of affairs *x's being at position* p_1 and *x's being in position* p_2 to 'y' and 'z', respectively, will not satisfy the definition, nor will any other assignment to 'y' and 'z'. So the first, but not the second, will count as a genuine case of alteration.

Slote's definition does not mention states of affairs. This is not because he does not believe in these entities. Indeed, his intention is that the definition will be satisfied in virtue of certain states of affairs being assigned to the variables 'y' and 'z'. But he believes that the definition works adequately as it is—phrased only in terms of modal notions, spatiotemporal notions, and logical notions. The notion of a state of affairs is not a primitive notion for Slote.

But it is not part of my project to eliminate states of affairs or related entities from primitive ideology, so I will develop a more intuitive definition that is based on Slote's definition but makes reference to properties. First, some terminology. I continue to use an actualist modal language. Say that a relation (property) is *existence entailing* iff necessarily, if $x_1 \dots x_n$ stand in that relation (property) at a time, then $x_1 \dots x_n$ must all exist at that time. Properties P and Q are *incompatible* iff it would be impossible for any object to have both P and Q at the same time. Property P is *temporally bound* iff $\Box\forall x\forall t[P(x,t)\rightarrow\forall t'>t\ P(x,t')]$.[6] I will denote the location of x at time t by "loc(x,t)".[7] I will write "$P(x,t,p)$" iff loc$(x,t)=p$ and $P(x,t)$. I represent Slote's definition thus:

> *Analysis of Alteration 1*
>
> $Alter(x,t_0,t_1) =_{\text{df}} \exists P\exists Q$ such that
>
> a $P(x,t_0)\wedge Q(x,t_1)$
>
> b P and Q are existence-entailing[8] and incompatible and neither is temporally bound
>
> c $\forall p\forall t[\Diamond P(x,t,p)\wedge\Diamond Q(x,t,p)]$
>
> d $\Diamond[\exists t P(x,t)\wedge\exists t Q(x,t)\wedge x$ never moves$]$
>
> e $\sim\exists z\Box\forall t\forall t'([P(x,t)\ \wedge\ Q(x,t')\ \wedge\ \text{Exist}(z,t)\ \wedge\ \text{Exist}(z,t')]\rightarrow[\text{loc}(x,t)\neq\text{loc}(x,t')\vee\text{loc}(z,t)\neq\text{loc}(z,t')])$

As an example of the application of this definition, let us return to the case where x changes position between times t_0 and t_1 without altering. Let $P =$ *having position* p_1 and let $Q =$ *having position* p_2. Conditions a, b, d, and e seem satisfied, but c is not. If other assignments to the variables also fail to satisfy the definition, then it yields the correct result: x does not alter between t_0 and t_1.

As stated, however, the definition is inadequate. Suppose that x never alters throughout its existence, but at t_1 it is moving while at t_2 it is stationary. Suppose x's color (all over) is blue. Consider the following assignments:

$P =$ *moving or being non-blue*

[6]The second half of Slote's definition of 'temporally bound' can be left out because of condition c.

[7]Let us introduce some object that is not a place to be the value of loc(x,t) when x does not exist at time t. Thus, if x exists at t but not at t', then loc$(x,t)\neq$ loc(x,t').

[8]I include the assertion that P and Q are existence-entailing to mirror Slote's assertion in his definition that y and z *depend on* x. In Slote's original definition, the intended values for 'y' and 'z' are states of affairs involving x.

$Q =$ *being stationary and being blue*

As near as I can tell, the conditions above are satisfied, and hence the theory gives the incorrect result that x alters between t_1 and t_2.

The difficulty can be fixed by adding condition d′:

d′ $\Diamond[\exists t P(x,t) \wedge \exists t Q(x,t) \wedge x$ moves whenever it exists$]$

Intuitively, d′ is in the same spirit as d: they each rule out properties that entail certain states of motion. Condition d′ prohibits Q in my example.

There is another difficulty with this definition. Suppose x remains in place for its entire life and never changes any of its intrinsic properties. But suppose at t_0 there is a red ball (called "A") ten feet from x, whereas at t_1 there is nothing ten feet from x. This is the sort of situation for which condition e is designed: when $P =$ *being ten feet from A* and $Q =$ *not being ten feet from A*, condition e is not satisfied. But now consider the following assignment to 'P' and 'Q'

$P =$ the property of having something or other within ten feet

$Q =$ the property of having nothing within ten feet

Under this assignment, condition e in the definition is satisfied, for there is no one object such that, necessarily, if x has P and Q at two times t and t', and that object exists at t and t', then either x or *it* must have moved. As far as I can see, the other clauses in the definition are satisfied as well, and so the definition again gives the incorrect result that x alters between t_0 and t_1.

Let us then add the following two conditions that are in the spirit of Kim's definition (recall that $\text{Dist}(x,y,t)$ iff x and y are contingent objects that share no parts at t, and are such that neither's existence entails the other's existence):

f $\Diamond[\exists t P(x,t) \wedge \exists t Q(x,t) \wedge \sim\exists y \exists t\, \text{Dist}(x,y,t)]$

g $\Diamond[\exists t P(x,t) \wedge \exists t Q(x,t) \wedge \exists y \exists t\, \text{Dist}(x,y,t)]$

(It is necessary to add both conditions for reasons analogous to those that necessitated our adding d′ to supplement d). These conditions will rule out properties like *having something within ten feet*. In fact, these conditions seem to rule out the cases that the original condition e was intended to rule out,[9] so I propose that we drop condition e. Here is the modified analysis.

Analysis of Alteration 2

$Alter(x, t_0, t_1) =_{\text{df}} \exists P \exists Q$ such that

a $P(x, t_0) \wedge Q(x, t_1)$

b P and Q are incompatible and existence-entailing

c $\forall p \forall t [\Diamond P(x, t, p) \wedge \Diamond Q(x, t, p)]$

d $\Diamond[\exists t P(x, t) \wedge \exists t Q(x, t) \wedge x$ never moves$]$

d′ $\Diamond[\exists t P(x, t) \wedge \exists t Q(x, t) \wedge x$ moves throughout its existence$]$

f $\Diamond[\exists t P(x, t) \wedge \exists t Q(x, t) \wedge \sim\exists y\, \text{Dist}(y, x)]$

g $\Diamond[\exists t P(x, t) \wedge \exists t Q(x, t) \wedge \exists y\, \text{Dist}(y, x)]$

It should be noticed that by adding f and g, we have employed a new primitive: the part-whole relation.

From our discussion of Kim's definition above, we should notice that d and d′ do not rule out the property $P =$ *either being stationary or moving with velocity 5 meters per second*. However, this causes no problem in the present context. If Q is to satisfy condition b, it must be incompatible with P, and hence must entail the property *neither being stationary nor moving at 5 meters per second*. But then Q will fail condition d. Similarly, the conditions f and g do not rule out the case where P is the disjunction of *loneliness* and *coexisting with six pigs*. But again, if Q is to satisfy condition b then it must entail the property *neither being lonely nor coexisting with six pigs*, and hence will fail to satisfy condition f. Indeed, a major reason for examining this definition is that it avoids the problems Lewis presents for Kim.

[9]The example Slote gives on p. 137 is essentially the example I gave above involving the properties *being ten feet from A* and *not being ten feet from A*. There is another reason to drop condition e. Let $z =$ any object that *cannot* exist at two different times. (e.g. an instantaneous time stage). In this case, the conditional in e has an impossible antecedent, the statement beginning with '$\Box$' is true, and the condition is failed, for any properties P and Q.

Our new conditions f and g also rule out counterexamples to Slote's original definition that were noticed by Frederick Schmitt (1978, pp. 406–407). Here is the general form of Schmitt's counterexamples. Let x and y be objects, and P' and Q' be incompatible intrinsic properties. The properties

$$P = \lambda x(y \text{ is } P')\ (\textit{being such that } y \textit{ is } P')$$
$$Q = \lambda x(y \text{ is } Q')\ (\textit{being such that } y \textit{ is } Q')$$

seem to satisfy the original definition. So suppose that x never changes any of its intrinsic properties during its lifetime, but y changes from having P' to having Q'. This means that x goes from having P to having Q, and the original analysis counts this as an alteration to x.

The Analysis of Alteration 2 is a significant improvement over Slote's original definition. Unfortunately, it is still inadequate. It cannot rule out cases based on Nelson Goodman's grue/bleen examples.[10] Let t_0 be some time, and define P and Q as follows:

$$P = \lambda x \lambda t[(\text{Green}(x,t) \wedge t < t_0) \vee (\text{Blue}(x,t) \wedge t \geq t_0)]\ (\textit{Grueness})$$
$$Q = \lambda x \lambda t[(\text{Blue}(x,t) \wedge t < t_0) \vee (\text{Green}(x,t) \wedge t \geq t_0)]\ (\textit{Bleenness})$$

Suppose that object x does not alter at all between times t and t' where $t < t_0 < t'$. Further suppose that x is green during this interval. Properties P and Q fit the definition of alteration, and hence that definition yields the incorrect result that x alters between t and t'.

We can generalize further. Let ϕ be any function from objects and times to propositions, and F be any intrinsic property. Further suppose that object x does not alter between times t_0 and t_1 and has F during that period, but that $\phi(x,t_0)$ is true while $\phi(x,t_1)$ is false. It seems to me that the following P and Q will satisfy the Analysis of Alteration 2:

$$P = \lambda x \lambda t[(F(x,t) \wedge \phi(x,t) \text{ is true}) \vee (\sim F(x,t) \wedge \phi(x,t) \text{ is false})$$
$$Q = \lambda x \lambda t[(\sim F(x,t) \wedge \phi(x,t) \text{ is true}) \vee (F(x,t) \wedge \phi(x,t) \text{ is false})$$

provided that, for any x and t, $\phi(x,t)$ and the proposition *that x is F at t* are logically independent, and provided that neither F nor its negation is essential to x.

[10] Goodman (1955, p. 74). Phillip Bricker suggested looking at grue/bleen.

I see no way around this problem, and so I believe that Slote's analysis of alteration is a failure. It should be noticed that the property grueness is also a counterexample to Kim's definition of 'intrinsic' above. A formidable task for any reductive definition of words like 'intrinsic' and 'duplicate', or others of their ilk, is to solve this problem.

8.2.5 Slote's Analysis of Alikeness

Slote uses his analysis of alteration to analyze the notion of distinct objects being exactly alike at some given time. In more familiar terminology, x and y are exactly alike at time t iff x's stage at t is a perfect duplicate of y's stage at t. We found reason to reject Slote's analysis of alteration, but I will consider his analysis of alikeness anyway. Perhaps an adequate analysis of alteration is possible after all. Also, it is interesting to see whether duplication may be analyzed in terms of alteration, for the concept of alteration could be taken as a primitive notion. If analyses of duplication and intrinsicality were then forthcoming, these would not be entirely reductive, since the notion of alteration seems in the same boat with duplication and intrinsicality. Still, there may be some merit to taking alteration as our primitive.

A new notion employed in Slote's definition of 'alike' is that of two objects being "incongruous". The idea here is a spatial one: left hand/right hand mirror images are incongruous despite perhaps being otherwise exactly similar.

His definition is as follows (1975, pp. 142–3):

> ⌜a and b are non-identical mutable unbound entities exactly alike at t_0⌝ means about the same as ⌜there are an x identical with a, a y identical with b, and a t identical with t_0 such that x and y are non-identical mutable unbound entities and there is at t a temporally unbound z dependent on x and y but not depended on by either x or y; z cannot exist at all times when either x or y exists if x alters between some two times and y does not (or vice versa); z can exist at some time, or during some period, even if both x and y are altering at that time, or during that period, and even if both x and y are not altering at that time or during that period; z can exist between some t' and t'' only if for every two times t''' and t'''' between t' and t'' or identical with t' or t'', it is not the case that x at t'''' is (to some degree) unlike x at t''' but y at t'''' is not (to any degree) unlike y at t'''; z does not depend on any

> temporally unbound immutable entity w that depends on x and y; and at t, x and y are not incongruent or incongruous⌝

The intent is that the definition will be satisfied when 'z' is assigned the state of affairs: *x and y's being exactly alike*.

I represent this definition as follows. An object is *mutable* iff it is possible that it alters. For any objects z and z', let $t_i(z,z')$ be the earliest time when either z or z' exists; let $t_f(z,z')$ be the latest time when either z or z' exists. When we have $R(x,y,t'')$ for every t'' between t and t' (inclusive), I will write "$R(x,y,[t,t'])$".

Furthermore, I extend the notion of alteration from the previous section as follows: when $t = t'$, I interpret "Alter(x,t,t')" as meaning that x is *altering* at t. Slote does not define the locution "x is altering at t''. He alludes to the possibility of such a definition in terms of "Alters(x,t,t')" and spatiotemporal terms on p. 141 footnote. 14. One presumably would go about it thus. *x is altering "from the past" at t* iff there is some $t' < t$ such that for any $t'' \in [t',t)$, Alter(x,t'',t). The definition of "x is altering at t from the future" is analogous. Presumably, x is altering at t iff x is altering at t from either the past or the future.

Analysis of Alikeness 1:[11]

Let x and y be non-identical mutable temporally unbound entities.

$Alike(x,y,t_0) =_{\text{df}} \exists R$ such that R is existence entailing and:

a $R(x,y,t_0)$

b $\forall t \Diamond R(x,y,t)$

c $\forall t \Diamond [\text{Exist}(x,t) \wedge \text{Exist}(y,t) \wedge {\sim}R(x,y,t)]$

d $\Box\{R(x,y,[t_i(x,y),t_f(x,y)]) \rightarrow \forall t \forall t'(t \neq t' \rightarrow [\text{Alter}(x,t,t') \leftrightarrow \text{Alter}(y,t,t')])\}$

e $\Diamond \exists t,t'[R(x,y,[t,t']) \wedge \text{Alter}(x,t,t') \wedge \text{Alter}(y,t,t')]$

f $\Diamond \exists t,t'[R(x,y,[t,t']) \wedge {\sim}\text{Alter}(x,t,t') \wedge {\sim}\text{Alter}(y,t,t')]$

g $\Box \forall t \forall t'[R(x,y,[t,t']) \rightarrow \forall t'' \forall t''' \in [t,t']\{t'' \neq t''' \rightarrow (\text{Alter}(x,t'',t''') \leftrightarrow \text{Alter}(x,t'',t'''))\}]$

h The state of affairs *R's holding between x and y* does not entail (include) any "temporally unbound immutable state of affairs"

i x and y are not incongruous at t

The idea here is that R will be the relation *being exactly similar to*.

Condition h is obscure. It is intended to rule out the following case which I will discuss below:

> $R = \lambda x \lambda y$(x and y are exactly alike except that x has the determinate shade green$_n$ and y has the determinate shade red$_n$)

R is ruled out by condition h because the state of affairs *x and y's standing in R* entails the state of affairs *x's having the determinate shade green$_n$ and y's having the determinate shade red$_n$*, which, according to Slote, is a "temporally unbound immutable state of affairs". It is immutable, presumably,

[11] Clauses a, b, and c are intended to capture Slote's condition a (p. 139). c is slightly stronger than what Slote says, however.

because there are two "determinate" intrinsic properties, namely, greenness$_n$ and redness$_n$, such that it entails that x has the first, while y has the second. I do not know how to further clarify this notion, and I will not even attempt to, since, as I will argue below, condition h should be dropped anyhow.

This definition is somewhat daunting. But we can make some simplifications. First, g together with the fact that R is existence-entailing entails d;[12] hence, the latter can be eliminated. Moreover, g is is equivalent to the simpler g′:

g′ $\Box\forall t\forall t'[(t \neq t' \wedge R(x,y,[t,t'])) \rightarrow (\text{Alter}(x,t,t') \leftrightarrow \text{Alter}(y,t,t'))]$

So we have:

Analysis of Alikeness 2:

Let x and y be non-identical mutable temporally unbound entities.

$Alike(x,y,t_0) =_{\text{df}} \exists R$ such that R is existence-entailing and:

- a $R(x,y,t_0)$
- b $\forall t\Diamond R(x,y,t)$
- c $\forall t\Diamond[\text{Exist}(x,t) \wedge \text{Exist}(y,t) \wedge \sim R(x,y,t)]$
- e $\Diamond\exists t,t'[R(x,y,[t,t']) \wedge \text{Alter}(x,t,t') \wedge \text{Alter}(y,t,t')]$
- f $\Diamond\exists t,t'[R(x,y,[t,t']) \wedge \sim\text{Alter}(x,t,t') \wedge \sim\text{Alter}(y,t,t')]$
- g′ $\Box\forall t\forall t'[(t \neq t' \wedge R(x,y,[t,t'])) \rightarrow (\text{Alter}(x,t,t') \leftrightarrow \text{Alter}(y,t,t'))]$
- h The state of affairs *R's holding between x and y* does not entail (include) any "temporally unbound immutable state of affairs"
- i x and y are not incongruous.

[12] Suppose $R(x,y,[t_i(x,y),t_f(x,y)])$. Since R is existence-entailing, x and y come into and go out of existence at the same times. Let $t \neq t'$. If both t and t' are not in $[t_i(x,y),t_f(x,y)]$ then clearly neither x nor y alters between t and t'. If exactly one of t and t' is in this interval, then both x and y alter between t and t' (in note 5 it is mentioned that coming into existence shall be taken to be an alteration). Finally, if t and t' are each in $[t_i(x,y),t_f(x,y)]$ then it follows from g that Alter(x,t,t') iff Alter(y,t,t').

This definition will not do in its present form. Condition i is present to rule out the case of two perfect mirror images of each other—say a left hand and a right hand that are exact duplicates except for being mirror images. For let $R =$ *being a perfect mirror image of*; R satisfies all of the conditions before condition i. In particular, R satisfies condition g′, since any alteration in one of the hands would need to be accompanied by an alteration in the other, if they were to continue being perfect mirror images.[13]

Similarly, condition h is intended to rule out the case where

> $R = \lambda x \lambda y$(x and y are exactly alike except that x has the determinate shade green$_n$ and y has the determinate shade red$_n$)

since, again, this relation satisfies the conditions before condition h. In particular, notice that it satisfies g′. For suppose it holds between objects x and y during some time period. Any alteration in the color of exactly one of x and y would make R cease to hold, and any other alteration in exactly one of x and y would also make R cease to hold.

But these cases can be generalized in ways that get around conditions h and i. For the sake of definiteness, let us develop a specific objection that generalizes the mirror image case. Suppose that there is some fundamental intrinsic magnitude called "parity" that comes in two degrees: on or off. Further suppose that an object's parity is independent of its other intrinsic properties (except those like *being on and being green*), and let:

> $R = \lambda x \lambda y$(x and y are intrinsically exactly alike except that x is on iff y is off)

Now suppose that x is on and y is off, but they are otherwise exactly alike. Our definition gives the wrong answer: that x and y are perfect duplicates. In particular, notice why R satisfies g′. If x but not y altered by changing its parity then R would cease holding between them. On the other hand, if x but not y altered in some other way, then again R would cease to hold between them, since R requires that its *relata* be exactly alike in respects other than parity. Also notice that R seems to satisfy h. The state of affairs *x and y's standing in R* does not entail any determinate parity for either x or y.

[13]It is not clear that mirror images are not duplicates. In fact, the analysis of duplication I accept in section 4.2.1 is consistent with mirror images being duplicates. And this consequence seems acceptable.

We may generalize the color case if we help ourselves to a set of determinate shades of red$_i$ where the index i ranges over, say, the real numbers between 0 and 1. The following relation satisfies the conditions in Slote's definition:

> $R = \lambda x \lambda y(x$ and y are exactly alike except that $\forall i \in (0, 1)$, x is red$_i$ iff y is red$_j$, where $j = \frac{i+1}{2}$

The cases can be generalized further. We need the concept of an *intrinsic profile* from section 4.2.2. An intrinsic profile is a maximally specific intrinsic property. I note two facts about intrinsic profiles:

(F1) every (possible) object has exactly one intrinsic profile

(F2) a possible object alters iff it changes intrinsic profiles

(F1) follows from principle (C2) from section 4.2.3, which says that a property P is an intrinsic profile iff the set of P's (possible) instances is a maximal set of possible duplicates (provided we assume, as I do, that properties are individuated by necessary coextension). For the set of maximal sets of possible duplicates is a partition of the set of possibilia, and thus every possible object falls into exactly one of these sets. (F2) also follows from (C2), since an object alters between t_1 and t_2 iff its instantaneous stage at t_1 is not a duplicate of its stage at t_2.

Call the set of intrinsic profiles "IP". A function f is a *permutation* of IP iff f is a bijection from IP onto itself. For any such function f, if the relation:

(*) $R = \lambda x \lambda y \lambda t [\forall P \in \text{IP}\; P(x,t) \text{ iff } f(P)(y,t)]$

holds between objects x and y at some time, then all of the conditions before condition i will be satisfied.[14] In particular, condition g′ will be satisfied. For suppose that x, but not y, were to alter between t_1 and t_2. We show that it follows that R cannot hold between x and y throughout the closed interval $[t_1, t_2]$. *Suppose otherwise.* Call x's intrinsic profile at t_1 "P", its profile at t_2 "Q", and call y's profile at both t_1 and t_2 "S" (we here appeal to (F1)). Since

[14]Actually, in some extreme cases condition e might fail. Suppose R is defined as in (*) via some function f such that, necessarily, if R holds between objects x and y, then for *any* intrinsic profile F that x could then alter to have, it would be *impossible* for y to alter to have $f(F)$ (perhaps because the complement of $f(F)$ is an essential property of y).

y does not alter during this interval, we know that y has the same profile at t_1 and t_2 because of (F2). From (F2), we can infer that $P \neq Q$, since x does alter during the interval. Since R holds between x and y at t_1, we have $S = f(P)$, by (*) and (F1). Similarly, since R holds between x and y at t_2, we have $S = f(Q)$. But since $P \neq Q$ and f is a bijection, we have a contradiction.

So let x and y be any objects that exist at time t and are not incongruous pairs (say, my head and the Eiffel Tower). Let f be any permutation of IP such that f(x's intrinsic profile at t) = y's intrinsic profile at t, and define R as in (*). Slote's definition yields the result that x and y are exactly alike at t.

It seems that conditions h and i must be replaced by some more general condition that will rule out all these counterexamples. Notice that there is one permutation of IP such that, when relation R is defined in as in (*), then R *must* be the relation *being a perfect duplicate of*. This is the identity permutation: $f(F) = F$, for all $F \in \text{IP}$. For on that choice, R holds between x and y iff they have the same intrinsic profile—i.e. iff they are duplicates, whereas for other permutations, the resulting R can sometimes hold between objects when they have different profiles—i.e. when they are not duplicates. So I suggest that we add the following condition:

j R is (necessarily) reflexive

Intuitively, the situation is this. Condition g′ insures that, if the relation R holds between two objects through some period of time, any alteration in one of the objects must be exactly accompanied by simultaneous alteration in the other object. This is satisfied when R = *is a perfect duplicate of*. But, as we saw in (*), this can be achieved by any other choice of R that sets up a "determinate correspondence" between the intrinsic profiles of its *relata*. Requiring that R be reflexive rules out all relations of the form specified by (*) except the the one corresponding to the choice of f as the identity function, and this is the one choice that yields the correct answers.

It is unclear whether or not Slote could make use of condition j in his original project. Recall that he does not use the notion of a relation in his definition. Whether or not condition j could be adapted to suit Slote's purposes, I do not know.

Given condition j, we can simplify the definition. Condition e is intended to rule out the case where $R = \lambda x \lambda y \lambda t$(neither x nor y is altering at t). Condition f is intended to rule out the case where $R = \lambda x \lambda y \lambda t$(both x and

y are altering at t). Since there are objects that can exist without altering, and objects that can exist and alter, condition j rules these cases out. So I propose we drop conditions e and f.

Furthermore, I recommend that we strengthen g′ to a stronger claim that is every bit as intuitive as the original:

$$\text{g}'' \quad \Box\forall x'\forall y'\forall t\forall t'([t\neq t'\wedge R(x',y',[t,t'])]\rightarrow[\text{Alter}(x',t,t')\leftrightarrow\text{Alter}(y',t,t')])$$

So, I propose the following adaptation of Slote's definition:

Analysis of Alikeness 3:

Let x and y be non-identical mutable temporally unbound entities.

$Alike(x,y,t_0) =_{\text{df}} \exists R$ that is existence-entailing and such that

a $R(x,y,t_0)$

b $\forall t\Diamond R(x,y,t)$

c $\forall t\Diamond[\text{Exist}(x,t)\wedge\text{Exist}(y,t)\wedge\sim R(x,y,t)]$

g″ $\Box\forall x'\forall y'\forall t\forall t'([t\neq t'\wedge R(x',y',[t,t'])]\rightarrow[\text{Alter}(x',t,t')\leftrightarrow\text{Alter}(y',t,t')])$

j R is (necessarily) reflexive

I have not been able to create problems for this analysis of alikeness. But the concept of alikeness is not as general as the concept of *duplication*. The actualist modal language, the language used by Slote to state his definitions, has quantifiers that range, at any given time, over objects in a single world. In this language, when one says "Necessarily, x and y are alike iff …", x and y are assumed to be in the same possible world. So the concept of alikeness that I have been analyzing applies only to objects in the same possible world. Moreover, the concept only applies to objects at a single given time. But we might want to ask whether objects that exist at different times, or in different possible worlds, are duplicates—i.e. exactly alike. Furthermore, we use the concept of duplication to analyze the concept of an intrinsic property. Can the concept of alikeness do this job just as well?

One might think to generalize from Slote's definition as follows. First use 'alike' to define 'intrinsic':

P is intrinsic iff $\Box\forall x\forall y\forall t[\text{Alike}(x,y,t)\rightarrow(Px\leftrightarrow Py)]$

(That is, *intrinsic* properties are those that can never differ between worldmates that are exactly alike at a time.) Then go on to define 'duplicate' (here I use possibilist quantifiers). For any times t_1 and t_2, and possible objects x and y (perhaps in different worlds), let us use "$\text{Dup}(x,y,t,t')$" to mean "x as it is at t is a duplicate of y as it is at t'". (The cumbersome locution must be used since we are following Slote in thinking of objects as temporal continuants, and thinking of properties as being had relative to times.)

$\text{Dup}(x,y,t,t')$ iff for any intrinsic property P, $P(x,t)$ iff $P(y,t')$

Roughly, this says that objects are duplicates (relative to two times) iff they have the same intrinsic properties (at those times).

This attempt is a failure. Take a clearly intrinsic property—the property *roundness*. Now consider any possible world w such that for every time t, there are no distinct x and y such that $\text{Alike}(x,y,t)$. We define a new property as follows:

x has P at time t iff either (x's world is not w and x is round at t) or (x's world is w and x is not round at t)

Surely *roundness* never differs between pairs of alike objects (since alike objects are, in my terminology, duplicates that inhabit the same world and time). But then P never differs between alike objects either. For P could never differ between alike objects in any world other than w since it has the same extension as *roundness* in those worlds. But it can never differ between alike objects in w either, since in w, at no time is there a pair of distinct alike objects. Thus, P satisfies the current definition of 'intrinsic'. This seems incorrect. And there is more trouble. Let x be an object in w that is round at time t, and let y be a round object in some other possible world that is, intuitively, at time t exactly like x is at t. By the definition of 'P', y does not have property P whereas x does have property P, and thus x and y turn out *not* to be duplicates at t.

A better strategy would be to rework Slote's definition from the beginning to apply to objects in different possible worlds. To do this, I will use the possibilist language with possibilist quantifiers. I will also assume that we can speak meaningfully of crossworld and crosstemporal relations (e.g. *x at t_0 is redder than y at t_1*, where x and y are in different worlds). Let x and

y be any two possible objects. When relation R holds between object x at time t and object y at time t', I will write "$R(x,y,t,t')$".

Finally, I need notation for the crossworld and crosstime analog of $R(x,y,[t,t'])$. Intervals $[t,t']$ and $[t'',t''']$ are *congruent* iff they have the same (temporal) length. Let $[t,t']$ and $[t'',t''']$ be congruent intervals; for any $t^* \in [t,t']$, let "$t''+t^*-t$" denote the time in the interval $[t'',t''']$ that is as far from t'' as t^* is from t. I will write "$R(x,y,[t,t'],[t'',t'''])$" when i) the intervals $[t,t']$ and $[t'',t''']$ are congruent, and ii) for every $t^* \in [t,t']$, we have $R(x,y,t^*,t''+t^*-t)$.

Analysis of generalized alikeness:[15]

Let x and y be non-identical mutable temporally unbound entities.

$GenAlike(x,y,t_0,t_1) =_{\text{df}} \exists R$ that is existence-entailing and such that:

a $R(x,y,t_0,t_1)$

b $\forall t \forall t' \exists x' \exists y' R(x',y',t,t')$

c $\forall t \forall t' \exists x' \exists y' {\sim} R(x',y',t,t')$

g'' $\forall x' \forall y' \forall t \forall t' \forall t'' \forall t''' [t \neq t' \wedge R(x',y',[t,t'],[t'',t'''])) \rightarrow (\text{Alter}(x',t,t') \leftrightarrow \text{Alter}(y',t'',t'''))]$

j R is (necessarily) reflexive (i.e. $\forall x \forall t R(x,x,t,t)$)

(The quantifiers here are possibilist.) The intention is that the definition will be satisfied just when R is the relation born by x at t to y at t' iff x's stage at t is a perfect duplicate of y's stage at t'.

But there is a problem even with the new definition. The notion of generalized alikeness is still not general enough. It is restricted so as not to apply to immutable objects, objects that cannot alter. This is no accident. Let R' be *any* binary relation that satisfies conditions b and c, and holds between two immutable objects a and b. Now consider a new relation R, defined as follows:

[15] I ignore the fact that the original definition involved *de re* modality—this was irrelevant to the substance of the definition.

> $R(x, y, t, t')$ iff EITHER (i) either x or y is mutable, and x at t and y at t' are duplicates, OR ii) both x and y are immutable and $R'(x, y, t, t')$)]

This relation R satisfies the conditions in the definition. Moreover, a and b stand in R; thus, a and b would be said to be GenAlike by that definition, despite the fact that R' was an arbitrarily selected relation and a and b were arbitrarily chosen immutable objects. In particular, notice why R satisfies g''. Consider any two possible objects x' and y', and suppose we have $R(x', y', [t, t'], [t'', t'''])$. If either x' or y' are mutable, then clearly Alter(x', t, t') iff Alter(y', t'', t''') in virtue of clause i) of the definition of R above. But if x' and y' are each immutable, then this will also hold trivially, because of the fact that neither can alter.

Our concept of generalized alikeness, then, is restricted to mutable objects, and therefore is less general than the notion of duplication, which applies equally to mutable and immutable objects.

One might follow the generalizing strategy outlined above. First define 'intrinsic' using generalized alikeness, and then analyze the more general notion of duplication in terms of intrinsicality. The proposal is this:

> P is intrinsic iff $\forall t \forall t' \forall x \forall y[(x$ and y are mutable and GenAlike$(x, y, t, t')) \rightarrow (P(x, t) \leftrightarrow P(y, t'))$
>
> Dup(x, y, t, t') iff for any intrinsic property P, $P(x, t) \leftrightarrow P(y, t')$

(The quantifiers here, remember, are possibilist. The definitions say the following: a property is intrinsic iff it can never differ between mutable GenAlike objects; two objects, even immutable objects, are duplicates iff they share all intrinsic properties.)

This approach will fall into error if there are some properties that are *not* intrinsic, but which are particular to immutable objects. For let P be such a property. P will never differ between mutable generalized alike objects, since no such object ever has P, so we get the incorrect result that P is intrinsic. Moreover, this will create problems for the definition of 'duplicate', since this intuitively extrinsic property will be required to be shared by duplicates.

The question, then, is: are there any extrinsic properties particular to immutable objects? I think there are. The property *being an instantaneous stage of Ted* is extrinsic, for I might have had an exactly similar identical twin whose stages would be duplicates of my stages, but would not have this property. But no mutable object ever has this property, because of the fact that instantaneous stages cannot change, being unable to persist through time.[16]

I think, then, that we cannot analyze intrinsicality and duplication along Slotean lines. I argued above that Slote's analysis of alteration is unsuccessful. Even if we grant ourselves that concept, I have failed to use this concept to analyze intrinsicality and duplication. I suggest that Lewis's view on this matter is correct: 'intrinsic' and 'duplicate' cannot be defined in terms of "quasi-logical" notions.

[16]Another sort of example appeals to "purely" abstract entities, if there are such things. *Impure sets* are perhaps *not* immutable; perhaps when I get a haircut my unit set alters. But numbers, on the other hand, seem different. Surely, orthodoxy says that the number 9 is immutable. So, if the property *numbering the planets* is an extrinsic property, we have another example. However, I do not urge this example since I am wary of applying the notion of intrinsicality to purely abstract entities.

Chapter 9

Naturalness and Metrical Realism

The usefulness of the natural/nonnatural distinction is not limited to its relations to intrinsicality, duplication, and related notions. In "New Work for a Theory of Universals", David Lewis finds an abundance of applications for the distinction. Indeed, the philosophical utility of the concept of naturalness is a primary reason for its importance. In this chapter I discuss yet another of its applications.

9.1 Preliminaries

First, let me review some of the assumptions that I laid out in chapter 2. I assume that necessarily coextensive properties and relations are identical; this means that I may conveniently single out a property or relation by specifying its extension in every possible world. In fact, for convenience I will often talk as if I *identify* properties with the sets of *possibilia* that instantiate them; similarly for relations. Unless otherwise indicated, my quantifiers are possibilist.

The natural properties and relations are the most fundamental properties and relations—chapter 3 lays out my theory of naturalness. In this chapter I will ignore the fact that the natural/nonnatural distinction is best construed as a matter of degree, so when I say 'natural', I mean 'perfectly natural'. Some take naturalness to be an unanalyzable primitive; I call these theorists *primitive naturalists*.

We must be clear on the distinction between *mathematical* and *physical* spaces. Mathematical spaces are abstract mathematical objects. Physical spaces may be *represented* by mathematical objects, but they are themselves "concrete".[1]

Let us look at an example (the mathematics is simplified for clarity). The mathematical *Euclidean 3-space* may be thought of as the pair $\langle \mathbb{R}^3, d_E \rangle$. $\mathbb{R}$ is the set of real numbers, so $\mathbb{R}^3$ (i.e. $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$) is the set of three-tuples of real numbers. A member of $\mathbb{R}^3$, $\langle x, y, z \rangle$, is a "point" in the space. d_E is the "Euclidean distance function": the real-valued two-place function over $\mathbb{R}^3$ defined as follows:

$$d_E(\langle x,y,z\rangle, \langle x',y',z'\rangle) = \sqrt{(x-x')^2 + (y-y')^2 + (z-z')^2}$$

Once, people thought that the physical space of the actual world was Euclidean. Let's suppose for the sake of the example that they were right. We may take the *physical* space of the actual world to be the pair $\langle S, D \rangle$, where S is the set of the actual world's *physical points*, and D is the actual two-place *physical distance function* defined over the points in S.[2] S and D are *not* thought of as being purely mathematical entities. The members of S are real points in space, not representations. The value of the function D for arguments p_1 and p_2 is the actual distance between points p_1 and p_2. I consider in section 9.1 below just what it is for a function to be the real distance function for a physical space.

The claimed relation between the mathematical Euclidean 3-space and the physical space was one of *representation*; to say that physical space is "Euclidean" is to say that the mathematical Euclidean 3-space correctly represents physical space. "Representation" here means isomorphism, so the mathematical space $\langle \mathbb{R}^3, d_E \rangle$ correctly represents physical space $\langle S, D \rangle$ iff there is a one-one map f from S onto $\mathbb{R}^3$ such that for any points p_1 and p_2 in S, $D(p_1, p_2) = d_E(f(p_1), f(p_2))$.

Modern differential geometry teaches us of other mathematical spaces besides Euclidean 3-space. For example, if the Euclidean distance function d_E in $\langle \mathbb{R}^3, d_E \rangle$ is replaced by an appropriate distance function d; the resulting space $\langle \mathbb{R}^3, d \rangle$ will be "non-Euclidean". The claim that the geometry of physical space might be non-Euclidean is the claim that a non- Euclidean

[1] The word 'concrete' is a bit of a weasel-word. See Lewis's tirade in Lewis (1986*c*, section 1.7). All I mean by calling physical spaces "concrete" is to distinguish points of physical space from the mathematical objects that represent them in a mathematical space.

[2] Actually, there will be a family of such functions corresponding to arbitrary choices of units.

mathematical space might correctly represent (be isomorphic to) physical space. There are constraints d must satisfy in order to qualify as a distance function—these constraints are laid out by mathematicians. For example, d must always assign positive numbers. Another example: if point p_2 is (linearly) *between* points p_1 and p_3, then $d(p_1, p_2) + d(p_2, p_3)$ must be equal to $d(p_1, p_3)$.[3] But many distance functions are possible that satisfy the constraints. The distance function that is paired with the set of points in a mathematical space is that space's *metric*. The physical distance function, as well, is called the "metric" for that physical space.

I will discuss only space, ignoring time and spacetime. I intend what I say to carry over to time and spacetime. I generalize this simple discussion of mathematical and physical spaces in section 9.2.3.

I now want to explore the application of naturalness to a problem in the philosophy of space and time: the statement of metrical realism.[4] Along the way we will learn a good deal about primitive naturalism and naturalness. Section 9.2 introduces the problem of the statement of metrical realism and gives a solution involving naturalness. In sections 9.3 and 9.4 I reply to objections to the proposal of section 9.2.

9.2 Metrical Realism

9.2.1 Introduction to Metrical Realism

The *absolutist* believes that points of space exist. The *relationalist* denies this—there are spatial relations between *material objects*, but the only "points of space" are mathematical surrogates we introduce for convenience in science. Let us be absolutists, realists with respect to the existence of points of space.

We could go further. We could be realists with respect to the metric structure of space as well, countenancing objective facts about distances.

There are the various Euclidean and non-Euclidean mathematical spaces. The natural expectation is that, just as there are various mathematical spaces, so there are various possible physical spaces. But some have denied this. According to Hans Reichenbach, for example, there is no non- conventional

[3] Betweenness in $\mathbb{R}^3$ may be defined in terms of the usual structure of $\mathbb{R}^3$.

[4] Phillip Bricker suggested to me the idea of using naturalness and the Ramsey-Lewis method for defining theoretical terms.

answer to the question: Is space curved or flat? Suppose that I say that space is flat, whereas you say that it is curved. We still may not disagree on any prediction of the outcome of any possible observation if I invoke "universal forces" that systematically distort physical objects. Where you say that an object retains its geometrical properties but moves through curved space, I say that it moves through flat space but is distorted by universal forces. According to Reichenbach, we do not really disagree. Our theories are the same theory expressed in different vocabulary.[5] Since geometrical theories can always be protected from refutation in this way by appropriate adjustments in physical theory, Reichenbach's positivistic view is that the various geometrical theories do not correspond to different physical possibilities.

Others say that there is a fact of the matter as to the metric structure of space. Some possible worlds have Euclidean spaces, others have Riemanian spaces. At some possible worlds, physical space is curved; at others it is flat. If we agree then we are metrical realists.

Let us be realists in this way as well. It is surely the natural view. But how shall we understand this realist thesis? The claim we need to make is that the different mathematical metric functions correspond to different possibilities for physical space. More specifically, let W be the class of possible worlds, and S be the class of mathematical spaces. In a strong and crude[6] form, the claim would be:

Metrical Realism 1 there is a function f from W onto S, such that $\forall w \in W$, $f(w)$ is the unique member of S that correctly represents the physical space of w.

9.2.2 Physical Distance Functions

On the face of it, this claim of the metrical realist seems a little mysterious. In the introduction I said that mathematical spaces "represent" physical spaces by isomorphism. I called the "physical space of the actual world" a pair $\langle S, D \rangle$, where S is the set of the actual physical points, and D is the "real distance function". But this raises two related issues. First, we need to know what makes a function the "real distance function" of a given possible world.

[5] See Reichenbach (1958, chapter 1), and Sklar (1974, pp. 88–146).

[6] This is crude firstly because some worlds might have no spacetimes at all, secondly because of the problem of the arbitrary choice of unit, and thirdly because the notion of a mathematical space must be generalized. These issues are addressed below.

Second, we need some account of how we single out this correspondence with our language. The issues might be stated in question form: i) what facts about w's physical space make D its distance function? and ii) how do *we* single out D with words like 'distance'? The second question becomes pressing when it is recalled that distance between spatial points only indirectly connects up with observable states of affairs. We cannot sense points in space. Distances between points in space cannot be identified with, say, results of measurements using standard apparatus, for it is possible for that apparatus to be physically distorted and give results that deviate from true distances. We seem to have no direct "access" to metric facts about space, so how can we link up our words with spatial relations?

Here is our predicament. Consider a possible world w. As metrical realists, we believe in a distance function D for the physical space of w. What makes D the real physical distance function of w, and how do we single it out with our word 'distance'? Here is a first stab at an answer to both questions. Maybe there is exactly one function from pairs of physical points of w to real numbers with the characteristics appropriate to distance functions. This, the realist could claim, is the real distance function. We achieve reference to this function by stipulating that 'the distance function' refers to the one and only function with such and such characteristics.

This of course would be a mistake. Suppose S is the set of w's physical points and D is such a function. Then the image of D under any one-one map of S onto S will be such a function as well. But the correction is close at hand. I will complicate my proposal in later sections, but the basic idea is that the realist's claim should be that there is exactly one *natural* function from pairs of w's physical points to real numbers with the appropriate structure, and this is the real physical distance function. We achieve reference to this function by stipulating that our word 'distance' is to refer to the unique natural function with the appropriate characteristics. We should treat claims about physical distance using the Ramsey-Lewis method for defining theoretical terms (Lewis, 1970). The sentence

> a and b are three feet apart

would be analyzed as

> There is exactly one natural function f from pairs of w's physical points to numbers with such and such characteristics, and $f(a,b) = 3$.

9.2.3 The Proposal Generalized

Some comments should be made about geometrical concepts other than metrical concepts, such as the concept of *betweenness*. First, I would treat references to such concepts in the same way as references to metrical concepts: using the Ramsey-Lewis method. For example, as a first approximation,

> physical point p_1 is between physical points p_2 and p_3

could be interpreted as meaning:

> there is exactly one natural relation R over triples of w's physical points with such and such characteristics, and $R(p_1, p_2, p_3)$

This is a "first approximation" since I must immediately refine the claim. *All* geometrical concepts should be "Ramsified" at once. Suppose there are m geometrical concepts $C_1 \ldots C_m$. Let sentence ϕ be a complete geometrical characterization of these concepts. For each geometrical concept C (distance, betweenness, etc.) ϕ will have a conjunct ϕ_C that characterizes C, and a free predicate variable V_C that stands for C. ϕ_C characterizes C in part by having free occurrences of variable V_C; for example, ϕ_C could require that C be a transitive binary relation by containing the following conjunct: $\ulcorner \forall x \forall y \forall z[(V_C xy \wedge V_C yz) \rightarrow V_C xz] \urcorner$ specifying the characteristics of C. But ϕ_C may contain free variables corresponding to other geometrical concepts as well, since constraints made on one concept may involve its relations to another concept. (For example, I noted above that a distance function D is constrained by the requirement that, if point p_1 is *between* points p_2 and p_3, then $D(p_2, p_1) + D(p_1, p_3) = D(p_2, p_3)$.)

The geometrical concepts $C_1 \ldots C_m$ will be defined to be the unique m natural relations/functions satisfying ϕ. Any sentence ψ containing references to some geometrical concepts will be interpreted as meaning:

> there are exactly m natural relations/functions $V_{C_1} \ldots V_{C_m}$ over w's physical points such that ϕ, and ψ'

where ψ' is the result of replacing every term purporting to refer to a geometrical concept C in ψ by the variable V_C.

We may now generalize our analysis of the claim that a given mathematical space represents a given physical space. The mathematician explores affine structure, topological structure, etc. in addition to metric structure.[7] But when I characterized mathematical spaces as pairs of the form $\langle \mathbb{R}^3, d \rangle$ for various distance functions d, I was focusing on a special case. A more general way of characterizing a mathematical space would be to take it to be a pair $\langle S, Q \rangle$ of a set S and a set Q. S could be any set whatsoever, thought of as the the set of "points" of the space, and Q would be a set of mathematical entities defined over the members of S that specifies the structure of the space. We have focused on one sort of member of Q: distance functions. But Q could also include a 3-place betweenness relation, or other, more complicated, mathematical entities.[8] There is a corresponding generalization of the concept of physical space: take the physical space of world w to be the pair $\langle S_w, Q_w \rangle$, where S_w is the set of w's physical points and Q_w is the set of physical geometrical relations/functions restricted to the points of S_w. For example, Q_w might include the physical distance function over points of w, the physical betweenness relation over points of w, etc. The members of Q_w are defined all at once to be the unique natural relations satisfying such and such restrictions, as I outlined above. And to say that mathematical space $\langle S, Q \rangle$ *represents* $\langle S_w, Q_w \rangle$ is, as before, to say that the two are isomorphic.

A goal for metrical realism was to interpret the claim that world w has such and such metrical structure. Notice that the view I am proposing lets us interpret analogous claims for non-metrical structure. For example, to claim that the physical space of a certain world has certain topological features would be to claim that its physical space is represented by a mathematical space with those topological features. This claim, on the present view, is the claim that there are certain natural relations over physical points with certain properties.

Notice another feature of the view: some claims about physical space will be claims that there are *no* natural relations over physical points with such and such properties. For example, suppose we want to claim that a world w has a neo-Newtonian space. Along with the claim that certain entities are natural, this will involve the claim that no entities are natural that would correspond to absolute velocities or positions.[9]

---

[7]See Sklar (1974, pp. 46–54) for a brief nontechnical survey.

[8]In my simplified example where $\mathbb{R}^3$ was the set of points, I was assuming its usual structure. In the general case, this structure would need to be put into Q, for in the general case the set S of points is merely a set with no tacitly assumed structure.

[9]See Sklar (1974, pp. 202–209) for a brief introduction to neo- Newtonian spacetime.

9.2.4 The Proposal Refined

For the rest of the chapter, I will mostly discuss distance functions, but it should be kept in mind that I envision all geometrical concepts being Ramsified at once, as in the previous section. For simplicity, I will pretend that geometrical facts about a world w may be described completely by a two place real-valued distance function over the physical points of w's space. Hence, the simple theory I am working with at present is the theory *Metrical Realism 1* from section 9.2.1, conjoined with the claim that the physical distance function at a world w is the unique two-place real-valued natural function D defined over w's physical points with characteristics appropriate to distance functions.

There are several refinements to our proposal that we must make. The first deals with units. Surely there is nothing special about using *feet* as our unit of distance. Hence it is implausible that there is a natural distance function in terms of feet, but not one in terms of meters. The realist's claim might be modified: there is exactly one *family* of natural functions from pairs of points to real numbers (with the appropriate structure), such that the functions in the family are suitable transformations of each other. In our simple case of three-dimensional space, the suitable transformations would be exactly the *scalar transformations*: function D is a scalar transformation of function D' iff there is a real number k such that $\forall x \forall y[D(x,y) = kD'(x,y)]$.

But there is another problem. I argued in chapter 5 that natural properties and relations cannot involve numbers if numbers are entities that can be constructed in various equivalent ways from sets. The brute force solution would be to accept real numbers as primitive entities. There is, however, a more moderate solution.

In measurement theory, questions of the following form are answered: given a certain class of entities and relations among those entities, what numerical measurement functions can be defined over those entities corresponding to those relations, and to what extent are those functions unique? For example, given the relation of *greater than or equal length* and the operation of *concatenation*, and their formal properties, what corresponding real valued *length* function can we construct for a given set of measuring rods?[10] The length function to be constructed must "correspond" to the original relations in the following sense. Let "Rxy" mean that rod x is at least as long

[10] See Krantz et al. (1971) for specifics on the theory of measurement.

as rod y, and let “$x = y \circ z$” mean that rod x is a concatenation of rods y and z. Length function l corresponds to the original relations iff for all rods x, y, and z, i) if Rxy then $l(x) \geq l(y)$, and ii) if $x = y \circ z$ then $l(x) = l(y) + l(z)$.

I had imagined the realist claiming that the distance function, a function from pairs of physical points to numbers, is the entity to which naturalness accrues. But the realist could avoid this by using measurement theory. In the measuring rod example, relations among the set of measuring rods enabled definition of a corresponding length function. The metrical realist would pursue an analogous strategy based on natural relations among physical points of space. If certain natural relations among physical points satisfy certain formal constraints, then distance functions will be definable. In fact, there will be a family of definable functions corresponding to different arbitrary choices of units. The realist's claim, then, should be modified: the distance functions of a world are those definable from certain natural relations over that world's physical points using measurement theory.

I will illustrate this proposal with an example: the definition of a metric for a three dimensional Euclidean space, taken from Suppes, Luce, Krantz, and Tverski's *Foundations of Measurement Volume* 2, pp. 84–111 (“FM2”). Suppose at world w there are just two natural relations over the set S_w of w's space points: a ternary relation B and a quaternary relation $\sim$. Intuitively, B is betweenness and $\sim$ is congruence; we have $Bxyz$ iff y is between x and z, and we have $xy{\sim}zw$ iff the distance between x and y is equal to that between z and w. I say ‘intuitively’ for I assume no primitive grasp of these relations, nor any operational definitions in terms of measurements. These relations are singled out as being a betweenness relation and a congruence relation solely by the fact that they are natural relations satisfying certain formal constraints. FM2 lists these constraints; here are a few examples:

if $Bxyx$ then $x = y$

$xy{\sim}yx$

if $Bxyz$, $Bx'y'z'$, $xy{\sim}x'y'$, and $yz{\sim}y'z'$, then $xz{\sim}x'z'$

If natural relations B and $\sim$ satisfy these constraints then, I say, they are a betweenness relation and a congruence relation, respectively. Now suppose further that that these relations satisfy certain additional formal constraints. These further constraints have the result of determining whether or not the space is Euclidean. Here is one example, which we may suppose holds of the physical points from w:

(*) if $Bxvw$, $Byvz$, and $x \neq v$, then there are u and t such that $Bxyu$, $Bxzt$, and $Buwt$

(*), in fact, corresponds to Euclid's famous parallel postulate, which requires that for any point p not on line l, there is exactly one parallel to l through p.

Say that D is an *acceptable distance function* for w iff D is defined over all pairs of points of S_w, and $D(x,y) = D(z,w)$ iff $xy{\sim}zw$. FM2 theorem 13.11 implies that for any acceptable distance function D, the physical space $\langle S_w, D\rangle$ is isomorphic to the mathematical Euclidean 3-space $\langle \mathbb{R}^3, d_E\rangle$ defined above. That is, D is a *Euclidean* distance function. Were (*) above replaced by something different, the acceptable distance functions for w might be non-Euclidean.

This example showed how we may start with a pair of relations (B and $\sim$), assumptions about those relations' formal properties, and prove that the resulting distance functions have a certain structure (in the example, they were Euclidean). Rather than claiming that there are natural functions relating physical points and numbers, I construe metrical realism as accepting natural relations over physical points, and constructing corresponding functions involving numbers using measurement theory.

An important fact from measurement theory is that that there are often different ways to construct a given measurement function. A certain system of relations may suffice, but there may also be some other system of relations that will do the trick as well. The metrical realist should not focus on any one system of relations. If one set of relations among a given world's physical points is natural and has the appropriate formal properties to enable construction of distance function D, then this would be enough for the realist to say that D is a physical distance function for w. And if more than one set of natural relations can do the job, then this would be fine as well, so long as the resulting distance functions agree.[11] These distance functions are only required to be unique up to certain transformations (recall the arbitrary choice of units).

Presumably, if at a given possible world there is no such set of natural relations, or if there are more than one that yield incompatible metrics, then realism should have the consequence that there is no fact of the matter at that world about the metric structure of space. Our statement of metrical

[11] Recall the problem of minimality from section 3.2.1.

realism, then, is as follows. Let W be the set of possible worlds, and S be the set of mathematical spaces.

Metrical Realism 2 there is a function f such that for every $w \in W$, f assigns to w the set of members of S that represent (are isomorphic to) $\langle S_w, D\rangle$, where S_w is the set of w's physical points of space, and D is some physical distance function for w.

(R1) D is a physical distance function for world w iff i) there is a family of natural relations R over w's physical points of space sufficient to define D (unique up to an appropriate transformation) and ii) no other such family of natural relations R' allows definition of a function D' (unique up to an appropriate transformation) that is *not* an appropriate transformation of D.

Notice that $f(w)$ will be empty if no mathematical spaces represent the physical space of w. This might be if there are systems of relations at w that allow definition of conflicting metrics, or if w has no physical space at all. Remember that this is all done for a special case—physical spaces of the form $\langle S, D\rangle$. I indicated in section 9.2.3 how this may be generalized.

My proposal is schematic at certain points. For example, I say that the distance function definable from some set of natural relations must be unique up to an "appropriate transformation", but I never say what transformations are appropriate. I expect that this and similar questions will be answered by the mathematicians. Also, I take as unanalyzed the notion of a point of space. (It would be interesting to investigate how this could be avoided.)

9.3 The Problem of Extra Relations

In the next two sections I discuss problems for this approach to metrical realism. The first is the problem of "extra relations". My approach to metrical realism will fall into error if there is some spurious natural relation that, intuitively, has nothing to do with distance or metric, but gets mistaken for a real distance relation because it just happens to have formal properties that enable definition of a metric. Let R be such a relation; suppose it has formal properties adequate for the definition of a bogus metric M. If no other natural relations enable definitions of conflicting metrics, then (R1) entails that M is the metric of the actual world, despite the stipulated fact that R has nothing to do with distance or metric. On the other hand, if there are other natural relations that enable definition of a conflicting metric (intuitively, the real metric), then (R1) says that our space has no metric structure. Either way, we have bad news. This is the problem of extra relations.

There are several responses to this problem; I remain undecided. One would be to claim that we have some way of singling out, from all the natural relations among points with the appropriate structure, the ones that correspond to real facts about distance and metric. This would be an admission that (R1) is unacceptable as stated, and needs to be supplemented. But how could we single out the "real" metrical natural relations from the bogus ones? Not perceptually. Not by means of formal structure, since the problem is that these bogus relations have the same formal structure as real metrical relations.

Another response to the problem of extra relations is to deny the possibility of extra relations. Perhaps it is a reasonable conjecture that the only natural relations holding between space points with formal properties adequate to define metrics are the right kind of relations, and not spurious relations that happen to have the relevant formal properties. The only other natural relations among space points that one can think of seem not to have the right formal properties for definition of metrics (consider *betweenness*, *simultaneity*, etc.) And, it may be a reasonable conjecture that all natural relations are spatiotemporal. We seem to have no clear counterexamples.

Finally, there is biting the bullet—standing by the consequences of (R1), come what may. Dialectically, this response is in fairly good shape. "You who maintain that there could be an extra, "bogus" natural relation: what makes that relation *bogus*? What distinguishes it from a *real* metrical relation? If some relation over spatial points is natural, and has formal properties that enable definition of a distance function, then I am prepared to stand by (R1) and call it a genuine metrical relation; I am perfectly happy to call the resulting function a genuine distance function. My concepts of metricality and distance extend no further."[12]

[12]Phillip Bricker has convincingly defended the third option (which, perhaps, is a dramatic version of the second option).

9.4 Primitive Naturalism and Contingency of the Metric

9.4.1 An Argument

In this section I want to consider the following objection. The metric structure of space is surely a contingent fact. But (R1) does not seem, on the face of it, to allow for the metric to vary properly from world to world. Suppose that the physical space of our world is Euclidean. According to (R1), this is because a certain system of natural relations among our world's physical points enables definition of a Euclidean distance function. But on my conception of naturalness, naturalness is a *property* of properties and relations; not a *relation* between these entities and possible worlds. So, since the relations are natural not just at our world, but natural simpliciter—natural at every world—it seems that they will make the spaces of those worlds Euclidean as well.

I believe, however, that there is no problem here. Furthermore, showing that there is no problem will, I believe, bring to light some facts about naturalness.

First, we must get clearer about the objection. Let us accept (R1), for the present. Further suppose that the metric of the space of @, the actual world, is Euclidean. To simplify matters, suppose that this is because there is exactly one natural relation *R* relating @'s points with the formal structure necessary to enable the definition of a distance function, and that the function is Euclidean. Where w is any world, the conclusion of the following valid argument implies that w cannot have a non-Euclidean space. But w was an arbitrarily selected world. So if I am committed to each premise, then this would show that I cannot account for the contingency of the metric.

Argument Against Contingency

(1) *R* relates the physical points of w.

(2) If (1), then (a): *R* enables the definition of a Euclidean distance function over the physical points of w.

(3) If (a) then the metrical structure (if any) of w's physical space is Euclidean

(4) Therefore, the metrical structure (if any) of w's physical space is Euclidean

The idea behind the argument is as follows. (1) is true because R, being a natural relation, cannot have "gaps" in its spread over logical space. (2) is also defended by appeal to R's naturalness: a natural relation ought to have the same formal properties from world to world. Finally, I am committed to premise (3) since i) natural relations are natural *simpliciter*, so ii) R is a natural relation allowing definition of a Euclidean distance function over w's points, and hence iii) (R1) implies that w's distance functions (if there are any) will be Euclidean. (The "if any" proviso is present because if there are some *other* natural relations that relate the points of w and enable definition of a distance function that is *not* an appropriate transformation of the Euclidean distance function definable from R, then (R1) has the result that w has *no* physical distance function).

As I see it, the first two premises are open to question. The justifications for the premises will be explained in more detail when I consider attacks on those premises. Suppose that w is a world that we would want to say has a non-Euclidean space; I consider premises (1) and (2) in turn.

9.4.2 Rejecting Premise (1)—Relation splitting

I see no reason why I cannot hold that R contains no 'tuples of points from w. R only relates points from Euclidean worlds like @. There may be other relations that relate w's points, and have over w's points the same formal structure as R has over @'s points. Indeed, if w has as many physical points as @, there must be some such relations.[13] But I can claim that R is not among these relations, and moreover that none of these relations are natural—hence, I am not forced by (R1) to hold that w has a Euclidean metric.

Let us focus on one of these relations that relates w's points and is a formal analog of R. Indeed, let us focus on a relation that is like this not only in w, but in every world that is non-Euclidean—call it R'. We can think of R' as the twin of R. Like R, R' enables definition of Euclidean metrics in worlds whose points it relates. R' relates only points in non-Euclidean worlds, while R covers the Euclidean worlds; together, R and R' span all of the worlds with physical spaces that have metrics.[14] Indeed, one might have

[13] Let $R^{@}$ be R restricted to points of @. $R^{@}$ has formal properties that enable definition of a Euclidean metric. Now, let f be any one-one correspondence between the points of @ and the points of w. The "mirror image" of $R^{@}$ under f will have the same formal properties as $R^{@}$.

[14] I ignore worlds with two physical spaces, one Euclidean, the other non- Euclidean.

expected the union of R and R' to be a natural relation. But the primitive naturalist can deny this, and claim that R is a natural relation, but neither R' nor $R \cup R'$ is natural. Premise (1) is false, for R only relates points from Euclidean worlds.

I call this the "relation-splitting" approach. The natural expectation is that a single natural relation relates points from both w and @. But the relation-splitter denies this, claims that the relation $R \cup R'$ is *not* natural, splits it into R and R', and claims that only the first is natural.

One might have doubts about the claim that R, but not $R \cup R'$, is natural. Natural relations, I said above, should not have "gaps". Fragments of more inclusive relations might seem, intuitively, to be unnatural. The subrelation gotten by restricting *being ten feet from* to objects in one possible world would be highly unnatural. Convinced by such a case, one might go on to argue that since R is a fragment of $R \cup R'$, it can't be natural.

In this general form, however, the objection is clearly mistaken. The principle:

(P1) if a property or relation P is a subset of some other property or relation P', then P is not natural

is obviously false. For any natural binary relation we can construct a superset by arbitrarily adding extra pairs.

We might try to modify the objection, employing instead the principle:

(P2) A proper subset of a natural property or relation is unnatural.

Unlike (P1), (P2) has some plausibility, at least if the natural properties are construed as the most fundamental properties.[15] Unfortunately, it does not help the argument. We would need to claim that $R \cup R'$ is natural, and then employ (P2). This, however, would blatantly beg the question against the relation splitter.

One might instead try to modify (P1). That principle is clearly false, but there is a sound intuition behind it. Natural relations should not be *arbitrarily* restricted, split, or fragmented. There should be no *arbitrary* gaps in

[15] Section 3.2.1. However, notice that I do not defend (P2) in that chapter; instead, I defend the weaker principle that conjunctions of distinct perfectly natural properties are not perfectly natural. Keep in mind that by 'natural' here I mean *perfectly* natural. Some fairly natural properties have even more natural subsets. For example, determinate shades of red are subsets of redness, and yet the determinate shades seem more natural than redness, their disjunction.

their span across logical space. The difficulty here is to say what constitutes an "arbitrary" restriction, split, or gap. The following terminology will let us make an attempt.

Consider any relation R, formal property P of relations, and world w. Say that R *has* P *at* w iff the subrelation of R restricted to w's objects has property P. Say that R has P *uniformly* iff R has P at every world that contains some objects related by R. Say that P is *metrical* iff there is some distance function D such that every relation among space points that has P suffices to define D. An example of a metrical property is the property of *being Euclidean*, which is had by a relation iff its formal properties suffice for the definition of d_E over the set of objects it relates.[16]

Let us recall what the relation splitter has said so far. R is uniformly Euclidean, as is R', although the class of worlds with points related by R is completely disjoint from the class of worlds with points related by R'. But R is natural, while R' is not; this is why @, but not w, has a Euclidean space.

The criticism I want to consider invokes the following principle:

(P3) If the relevant formal properties of a relation are uniform, then any proper subset of that relation is unnatural.

Since R and R' are both uniformly Euclidean and never relate points that are worldmates, $R \cup R'$ is also uniformly Euclidean. If we grant that *being Euclidean* is a "relevant" property, then (P3) implies that R is unnatural, contradicting the claim of the primitive naturalist.

Above we noted the intuition that natural properties and relations should not be arbitrarily split. (P3) is an attempt to say when a split is arbitrary. It says it is arbitrary to split up a relation when that relation's relevant formal properties do not vary from world to world. The relations that result from such a split, says (P3), are always more unnatural than the unsplit relation. Of course, we need an account of what formal properties of relations are "relevant". I will not attempt this (among other reasons, because I think (P3) is false, provided that 'relevant' is not so defined as to trivialize the principle). For present purposes I will assume that, for relations involving space points, metrical properties are relevant properties.

[16] These definitions can and should be extended to apply to *sets* of relations. Recall that I have simplified things by imagining that a single relation, R, enables definition of a Euclidean metric.

Though (P3) may represent a gallant attempt to say when a split of a relation is arbitrary, we can argue against (P3). Imagine a "mush world" w_m. No natural relation or system of relations that enables definition of a distance function relates the points of w_m's space. The realist who adopts (R1) would describe the mush world by saying that there are no facts of the matter about the metric of its space. There are no distances at the mush world; just "point soup". Now, suppose there is some natural relation R that is uniformly Euclidean.[17] Of course, R relates none of w_m's points. But we can form a new relation R' by adding 'tuples of points from w_m, taking care to only add 'tuples that will keep R' uniformly Euclidean. R' cannot be natural, since no natural relation enables definition of a metric for w_m. But it has a natural proper subset, namely R, and R has the same "relevant" (in this case, metrical) formal properties as R'. This is a counterexample to (P3).

It must be admitted that this counterexample depends on a rather exotic possibility. The defender of (P3) may reject the possibility of the mush world. Even in the absence of a counterexample to (P3), however, I think the relation splitter can simply reject (P3). It is no part of common sense! Moreover, (P3) seems to be loosely based on the intuition that what makes a relation natural is something about its formal properties. I discuss this intuition in the next section. It is mistaken.

9.4.3 Rejecting Premise (2)

So far I have been exploring the relation-splitting method of allowing for contingency of the metric. On this approach, the Argument Against Contingency is rebutted by rejecting premise (1). Throughout, you may have noticed, it was assumed that metrical properties of natural relations are uniform. For example, in the counterexample to (P3), I assumed that there was a uniformly Euclidean natural relation, and in the overall discussion I assumed that the relation R is uniformly Euclidean. Another approach, one that I favor over relation-splitting, questions this assumption. I turn next to this approach, which rejects premise (2) instead.

Let us forsake relation splitting, and grant premise (1). Thus R, the natural relation that is Euclidean at @, relates points of w. Indeed, let us assume

[17] If you deny that any natural relations among points have uniform metrical properties, then be patient. You will agree with the rejection of premise (2) of the Argument Against Contingency instead.

that *R* relates the points of every world with a physical space.[18] Premise (2) of the Argument Against Contingency claims that *R* is Euclidean at *w*. This claim might be defended by appeal to the following principle:

(P4) A natural relation has its relevant formal properties uniformly

'Relevant' is intended here in the same sense as in (P3). In particular, I assume that *being Euclidean* is relevant in the case of *R*. Hence, (P4) implies (2).

I see no reason to accept (P4). Moreover, I think we can provide arguments against it. I will consider two examples.

It is plausible that the topological relation of *betweenness* among spatial points is a natural relation. In our world, that relation presumably has a certain property of *denseness*: between any two points there lies a third. But at a world with discrete space, *betweenness* lacks this property. Assuming that *denseness* is "relevant", (P4) is false.

Secondly, a more abstract counterexample. Imagine the following situation that a universals theorist would describe as follows. There is an alien natural dyadic universal *U*, not instantiated in our world, which has no necessary (relevant) formal constraints. Think of *U* as being an ethereal rubber band linking its *relata*; the bands can be present in any combination. In terms of primitive naturalism, every (relevant) formal property of a certain alien natural relation is non-uniform. The universals view can clearly admit the possibility of this example. Primitive naturalism is intended to be the equal of the universals view in power; surely we do not want it to preclude this possibility. But if (P4) were true, then such an alien relation would be impossible.

Someone might object to the first counterexample using the distinction between rigid and nonrigid designators.[19] Consider any term *t* and object *o* in world *w*. Intuitively, we say that *t* rigidly designates *o* iff *t* picks out *o* and only *o* in every possible world. But this is too loose, for various reasons. Supposing *o* to be some concrete possible object, say that *t* "rigidly designates *o*" iff i) in any possible world containing exactly one counterpart of *o*, *t* denotes only that object, and ii) *t* denotes no other object in any other world. On the other hand, if *o* is a property or relation that exists "from the

[18] Perhaps: of every world whose physical space has appropriate non- metrical (e.g. topological) structure. Let us ignore relativistic considerations.

[19] Phillip Bricker pressed this objection to the first counterexample.

point of view of" every world, we may simply say that t rigidly designates o iff at every world, t designates o and only o.

Consider the counterexample involving the relation *betweenness*. The objector would claim that all I have established is:

(a) In the actual world, the betweenness relation is dense, but, at w, the betweenness relation is not dense

But what I need for a counterexample to (P4) is

(b) The relation that *in fact* is the betweenness relation is in fact dense, but *it* is not dense at w.

(a) implies (b) iff 'the betweenness relation' is a rigid designator, and a defender of (P4) might claim that 'betweenness' indeed is not a rigid designator. At the actual world, a certain relation is the betweenness relation, but since relations have their relevant properties uniformly, the betweenness relation at the discrete world is a different relation.

First, I think that 'the betweenness relation' is a rigid designator. Consider how one argues that a certain *singular term*, say, the proper name 'Bush', is a rigid designator. One says: "in any possible world, if *that guy* [pointing to Bush] exists, then he would be Bush; moreover, show me a possible object properly called 'Bush'—it must be *that guy*." I think 'the betweenness relation' passes the analogous test for relations. Suppose x is between y and z, and consider any world in which three points x', y', and z' are related *in that same way*. I think that x' would be between y' and z'. Similarly, it seems convincing to me that if x' is between y' and z', then x', y', and z' have to be related *in the same way* that x, y, and z are in the actual world.

A different argument may be given against (P4) involving betweenness, but we need some machinery from section 4.2.2. An *internal* relation is one whose holding supervenes on the intrinsic nature of its *relata*: if it holds between x and y, then it must hold between any duplicates of x and y, respectively. An *external* relation is a non- internal relation that supervenes on the intrinsic nature of the mereological sum of its relata. That is, if an external relation holds between x and y, then it must hold between the corresponding parts of any duplicate of the *fusion* of x and y. Hence, we have (Lewis, 1986*c*, p. 62):

(E) If R is external or internal and holds between x and y, then R must hold between the corresponding parts of any duplicate of the mereological sum of x and y

Notice that (E) quantifies directly over relations, *not* names of those relations.

We will need a recombination principle for the argument. *Recombination principles* underlie the inference that, for example, if a cat and a dog are possible objects, then it would be possible for there to be twenty duplicates of each, for there to be a duplicate of the cat's head attached to a duplicate of the dog's body, etc. One recombination principle, called *the principle of isolation*, says that, given any possible object c, there is a possible world containing a duplicate of c and that object's parts, but no other objects besides those entailed by the existence and nature of those objects already mentioned (this is principle (I2) from section 6.3.2).[20]

Finally, we need a principle from section 4.2.3: every (perfectly) natural relation is either internal or external. This principle is a straightforward consequence of the definitions of 'internal', 'external', and 'duplicate' in terms of naturalness.

Let us rigidly designate the relation that, at the actual world, is the betweenness relation by 'betweenness$^{@}$'. Presumably, betweenness$^{@}$ is dense in the actual world. And surely, it is a natural relation—the spatiotemporal relations are our only clear examples of natural relations. (Moreover, if such relations are not natural, then the idea in this chapter for stating metrical realism, embodied in (R1), will not work! In section 9.2.3 I indicated how I intend my proposal to be generalized. On this approach, the betweenness relation is defined to be the **natural** relation with such and such formal properties.) Thus, betweenness$^{@}$ is either external or internal. Suppose that in the actual world x is between and distinct from y and z, and let XYZ be the fusion of x, y, and z. The principle of isolation applied to the case of XYZ implies the existence of a possible duplicate of XYZ existing in isolation; call that duplicate XYZ'; call the possible world at which it exists "w". Since XYZ' is a duplicate of XYZ, and betweenness$^{@}$ is either external or internal, XYZ' is made up of three distinct parts, x', y', and z' where, by (E), x' is between$^{@}$ y' and z'. We are entitled to infer that x' is between$^{@}$ y' and z', rather than x' is between y' and z', because of the fact that the variable 'R' in

[20] The related *principle of isolation* in Paull and Sider (1992) was not intended to apply to points of space.

(E) ranges over relations, not relation names—betweenness$^{@}$ is an allowable substitution for '*R*'. Since XYZ' exists in isolation, nothing is between$^{@}$ y' and x'; hence, betweenness$^{@}$ at w is not dense. (P4), therefore, is false.

I say that I may fairly reject (P4), and also premise (2) of the Argument Against Contingency. I may claim that *R* is a natural relation, and yet varies in its metrical properties from world to world. At @ it is Euclidean, but perhaps at w it is *Riemanian* instead.

I imagine a complaint: "*R* seems unnatural. It seems a conglomeration of different relations all with different formal properties. For example, it has a Euclidean subrelation of pairs of points from @, and a Riemanian subrelation of pairs of points from w. What makes a relation natural, if not its formal properties?"

Well, whatever makes a relation natural, formal properties cannot be the whole story. Let *S* be the set of spacetime points from some possible world. There are many bijections from *S* onto itself. Each such bijection induces, for any relation defined over the members of *S*, an image of that relation, a "scrambled" version of that relation. The scrambled version will, in general, be radically different from the original. For example, if one began with the betweenness relation, the resulting relation would in general not be the betweenness relation. However, the scrambled relation will have the same formal properties as the original. For example, if the original was transitive, connected, and asymmetric, then the scrambled version will be transitive, connected, and asymmetric. If the original enabled definition of a Euclidean metric, then so will the scrambled version. So, take any natural relation *R*. Arrive at a new relation R' by arbitrarily scrambling *R* within each possible world whose points it relates. R' will have the same formal properties as *R*, but clearly will not be a natural relation. Hence, formal properties are not the sole determiners of naturalness.

What does make a natural relation natural? There are two answers, depending on how this question is interpreted. It may be interpreted as a request for an intuitive picture of why relations are natural. The answer, then, is that the *facts* make relations natural. *R* is natural because it corresponds to real metrical facts involving points. This of course is no *analysis*; according to (R1) metrical facts are analyzed in terms of naturalness, so an analysis of naturalness in terms of metrical facts would be circular. But I am a primitive naturalist, and no analysis of naturalness is *ever* offered by a primitive naturalist. The second way to interpret the question "What makes a relation natural" is as a request for analysis. The reply here is that there is none to be had. Naturalness is a primitive.

Consider the analogous situation with respect to the naturalness of the property *unit positive charge*. The primitive naturalist gives no analysis of why this property is natural. Someone might object: consider the set S that is the property *unit positive charge*. Exchange one member of S with some electron, to form a new set S'. Why isn't the property corresponding to S' natural? One answer is that the members of the set share no common feature; the set is heterogeneous. This is of course no analysis, since heterogeneity would presumably be analyzed in terms of naturalness. But do not demand analysis from the primitive naturalist.

If we consider S and S' in abstraction from facts about charge, then nothing will tell us which is a natural property. We can point out that the members of S share a common feature—charge—while the members of S' do not, but this is no analysis. In the same way, if we consider R and another relation in abstraction from facts about metric and distance, then nothing will tell us which is natural. For the primitive naturalist, there is a brute fact that R is natural, and this fact is not reducible to any claim about R's formal properties. Again, we can point out that the 'tuples of points in R all share a common feature, a feature involving metric and distance. But again, this is no analysis.

The primitive naturalist, then, can escape the Argument Against Contingency and account for the contingency of the metric, either by relation-splitting, thereby rejecting premise (1), or by rejecting premise (2) (I prefer the latter approach). Obstacles to each of these paths were presented, most formidably in the form of (P3) and (P4), but each, I think, is false. Finally, I argued that the primitive naturalist can feel free to claim that the naturalness of relations is not a mere matter of those relations' formal properties.

A note on the type

The text is set in Sabon, designed by Jan Tschichold in 1964.

Bibliography

Armstrong, David M. (1978*a*). *Nominalism and Realism*, volume 1 of *Universals and Scientific Realism*. Cambridge: Cambridge University Press.

— (1978*b*). *A Theory of Universals*, volume 2 of *Universals and Scientific Realism*. Cambridge: Cambridge University Press.

— (1986). "In Defence of Structural Universals." *Australasian Journal of Philosophy* 64: 85–88.

— (1989*a*). *A Combinatorial Theory of Possibility*. New York: Cambridge University Press.

— (1989*b*). *Universals: An Opinionated Introduction*. Boulder: Westview Press.

Bealer, George (1982). *Quality and Concept*. Oxford: Clarendon.

Benacerraf, Paul (1965). "What Numbers Could Not Be." *Philosophical Review* 74: 47–73.

Bricker, Phillip (1991). "Plenitude of Possible Structures." *Journal of Philosophy* 88: 607–19.

— (1992). "Realism without Parochialism." Read to the Pacific Division of the APA.

Campbell, Keith (1981). "The Metaphysic of Abstract Particulars." In Peter French, Theodore E. Uehling, Jr. and Howard K. Wettstein (eds.), *Midwest Studies in Philosophy VI: The Foundations of Analytical Philosophy*, 477–88. Minneapolis, MN: University of Minnesota Press. Reprinted in Mellor and Oliver 1997: 125–39.

Carnap, Rudolf (1929). *Der Logisches Aufbau der Welt*. Berlin: Schlachtensee Weltkreis-Verlag.

Chisholm, Roderick (1976). *Person and Object: A Metaphysical Study*. La Salle, Illinois: Open Court Publishing Co.

Demopoulos, William and Michael Friedman (1985). "Critical Notice: Bertrand Russell's The Analysis of Matter: Its Historical Context and Contemporary Interest." *Philosophy of Science* 52: 621–39.

Dunn, J. Michael (1987). "Relevant Predication 1: The Formal Theory." *Journal of Philosophical Logic* 16: 347–81.

— (1990). "Relevant Predication 2: Intrinsic Properties and Internal Relations." *Philosophical Studies* 60: 177–206.

Forrest, Peter (1986). "Neither Magic Nor Mereology: A Reply to Lewis." *Australasian Journal of Philosophy* 64: 89–91.

Goodman, Nelson (1955). *Fact, Fiction, and Forecast*. Cambridge, MA: Harvard University Press.

Halmos, Paul (1963). *Lectures on Boolean Algebras*. Princeton: D. Van Nostrand.

Kim, Jaegwon (1973). "Causation, Nomic Subsumption, and the Concept of Event." *Journal of Philosophy* 70: 217–36. Reprinted in Kim 1993: 3–21.

— (1982). "Psychophysical Supervenience." *Philosophical Studies* 41: 51–70. Reprinted in Kim 1993: 175–93.

— (1984). "Supervenience and Supervenient Causation." *Southern Journal of Philosophy* 22: 45–56. Spindel Conference Supplement.

— (1987). "'Strong' and 'Global' Supervenience Revisited." *Philosophy and Phenomenological Research* 48: 315–26. Reprinted in Kim 1993: 79–91.

— (1993). *Supervenience and Mind*. Cambridge: Cambridge University Press.

Kitcher, Phillip (1978). "The Plight of the Platonist." *Noûs* 12: 119–36.

Krantz, David H., R. Duncan Luce, Patrick Suppes and Amos Tversky (1971). *Foundations of Measurement*, volume 1 and 2. New York: Academic Press.

Kripke, Saul (1982). *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.

Lemmon, E. J. (1967). "Comments on D. Davidson's 'The Logical Form of Action Sentences'." In Nicholas Rescher (ed.), *The Logic of Decision and Action*, 96–103. Pittsburgh, PA: University of Pittsburgh Press.

Leonard, Henry S. and Nelson Goodman (1940). "The Calculus of Individuals and its Uses." *Journal of Symbolic Logic* 5: 45–55.

Lewis, David (1968). "Counterpart Theory and Quantified Modal Logic." *Journal of Philosophy* 65: 113–26. Reprinted in Lewis 1983*c*: 26–46.

— (1970). "How to Define Theoretical Terms." *Journal of Philosophy* 67: 427–46. Reprinted in Lewis 1983*c*: 78–96.

— (1980). "Index, Context, and Content." In Stig Kanger and Sven Ohman (eds.), *Philosophy and Grammar*, 79–100. Dordrecht: D. Reidel. Reprinted in Lewis 1998: 21–44.

— (1983*a*). "Extrinsic Properties." *Philosophical Studies* 44: 197–200. Reprinted in Lewis 1999: 111–15.

— (1983*b*). "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–77. Reprinted in Lewis 1999: 8–55.

— (1983*c*). *Philosophical Papers, Volume 1*. Oxford: Oxford University Press.

— (1984). "Putnam's Paradox." *Australasian Journal of Philosophy* 62: 221–36. Reprinted in Lewis 1999: 56–77.

— (1986*a*). "Against Structural Universals." *Australasian Journal of Philosophy* 64: 25–46. Reprinted in Lewis 1999: 78–107.

— (1986*b*). "Events." In Lewis (1986*d*), 241–69.

— (1986*c*). *On the Plurality of Worlds*. Oxford: Basil Blackwell.

— (1986*d*). *Philosophical Papers, Volume 2*. Oxford: Oxford University Press.

— (1991). *Parts of Classes*. Oxford: Basil Blackwell.

— (1993). "Mathematics is Megethology." *Philosophia Mathematica* 3: 3–23. Reprinted in Lewis 1998: 203–29.

— (1998). *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.

— (1999). *Papers in Metaphysics and Epistemology*. Cambridge: Cambridge University Press.

Loux, Michael J. (ed.) (1979). *The Possible and the Actual*. Ithaca, NY: Cornell University Press.

Mellor, D.H. and Alex Oliver (eds.) (1997). *Properties*. Oxford: Oxford University Press.

Mendelson, Elliott (1987). *Introduction to Mathematical Logic*. Belmont, CA: Wadsworth & Brooks.

Moore, G. E. (1922). "The Conception of Intrinsic Value." In *Philosophical Studies*, 253–75. London: Routledge & Kegan Paul.

Paull, R. Cranston (1994). *Supervenience Relations and Their Significance*. Ph.D. thesis, University of Massachusetts.

Paull, R. Cranston and Theodore Sider (1992). "In Defense of Global Supervenience." *Philosophy and Phenomenological Research* 52: 833–54.

Plantinga, Alvin (1974). *The Nature of Necessity*. Oxford: Oxford University Press.

— (1976). "Actualism and Possible Worlds." *Theoria* 42: 139–60. Reprinted in Loux 1979: 253–73.

Putnam, Hilary (1971). *Philosophy of Logic*. New York: Harper & Row.

Quine, W. V. O. (1951). "Ontology and Ideology." *Philosophical Studies* 2: 11–15.

— (1963). *Set Theory and Its Logic*. Cambridge, MA: Belknap Press.

— (1969). "Natural Kinds." In *Ontological Relativity and other Essays*, 114–38. New York: Columbia University Press.

Quinton, Anthony (1958). "Properties and Classes." *Proceedings of the Aristotelian Society* 58: 33–58.

Reichenbach, Hans (1958). *The Philosophy of Space and Time*. New York: Dover.

Resnik, Michael (1980). *Frege and the Philosophy of Mathematics*. Ithaca, NY: Cornell University Press.

Ryle, Gilbert (1939). "Plato's Parmenides." *Mind* 48.

Schmitt, Frederick F. (1978). "Change." *Philosophical Studies* 34: 401–16.

Shoemaker, Sydney (1980). "Causality and Properties." In Peter van Inwagen (ed.), *Time and Cause*, 228–54. Dordrecht: D. Reidel. Reprinted in Mellor and Oliver 1997: 228–54.

Sider, Theodore (1991). "A Hierarchy of Supervenience Relations." Unpublished manuscript.

Sklar, Lawrence (1974). *Space, Time, and Spacetime*. Berkeley: University of California Press.

Slote, Michael (1975). *Metaphysics and Essence*. Oxford: Basil Blackwell.

Stout, G. F. (1923). "Are the Characteristics of Particular Things Universal or Particular?" *Aristotelian Society, Supplementary Volume* 3: 114–22.

van Inwagen, Peter (1990). *Material Beings*. Ithaca, NY: Cornell University Press.

Wetzel, Linda (1989). "That Numbers Could Be Objects." *Philosophical Studies* 56: 273–92.

Williams, Donald C. (1952). "The Elements Of Being." *Review of Metaphysics* 7: 3–18, 171–92. Part I reprinted in Mellor and Oliver (1997, 112–24).

Williamson, Timothy (1985). "Converse Relations." *Philosophical Review* 94: 249–62.