

July 2016

The Path To Supersubstantivalism

Joshua D. Moulton
University of Massachusetts Amherst

Follow this and additional works at: https://scholarworks.umass.edu/dissertations_2



Part of the [Metaphysics Commons](#), and the [Philosophy of Science Commons](#)

Recommended Citation

Moulton, Joshua D., "The Path To Supersubstantivalism" (2016). *Doctoral Dissertations*. 673.
<https://doi.org/10.7275/8319743.0> https://scholarworks.umass.edu/dissertations_2/673

This Open Access Dissertation is brought to you for free and open access by the Dissertations and Theses at ScholarWorks@UMass Amherst. It has been accepted for inclusion in Doctoral Dissertations by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.

THE PATH TO SUPERSUBSTANTIVALISM

A Dissertation Presented

by

JOSHUA D. MOULTON

Submitted to the Graduate School of the
University of Massachusetts Amherst in partial fulfillment
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

May 2016

Philosophy

© Copyright by Joshua D. Moulton 2016

All Rights Reserved

THE PATH TO SUPERSUBSTANTIVALISM

A Dissertation Presented

by

JOSHUA D. MOULTON

Approved as to style and content by:

Phillip Bricker, Chair

Maya Eddon, Member

Christopher Meacham, Member

Jennie Traschen, Member

Joseph Levine, Department Chair
Philosophy

For Rebecca.

ACKNOWLEDGMENTS

Thanks to Chris Meacham for his extensive and detailed comments on numerous drafts of chapters, to Maya Eddon for her insightful comments (and especially for some encouraging words when I really needed them), and to Jennie Traschen for being willing to patiently listen as I explained this project to her. Finally, I want to thank my advisor, Phil Bricker, who read a great many version of a great many drafts of the dissertation. He, probably more than he realizes, has had a huge impact on the development of my thinking.

Finally, I truly could not have finished this without my parents' support. It would be impossible to overstate my gratitude to them. I also want to thank Gina, Ella, and Ben for tolerating the occasional evening moodiness that would follow a frustrating day of writing.

ABSTRACT

THE PATH TO SUPERSUBSTANTIVALISM

MAY 2016

JOSHUA D. MOULTON

B.A., REED COLLEGE

Ph.D., UNIVERSITY OF MASSACHUSETTS AMHERST

Directed by: Professor Phillip Bricker

This dissertation is divided into two parts. In the first part I defend substantivalism. I do this by offering, in chapter 1, a counterpart-theoretic defense of substantivalism from Leibniz' shift arguments. Then, in chapter 2, I defend substantivalism from the hole argument and argue, against the consensus, that the question of haecceitism is irrelevant to substantivalism in the context of general relativity.

In the second part of the dissertation I defend supersubstantivalism. I do this by offering, in chapter 3, an argument against dualistic substantivalism. The argument appeals to plausible principles of modal plenitude to show that the dualist is committed to a range of problematic possibilities. Then, in chapter 4, I consider a range of supersubstantivalist positions. I conclude by arguing for a version of supersubstantivalism I call *compresence supersubstantivalism*.

PREFACE

Our world is made either mostly or completely of material things – material objects, or material fields – separated by various spatiotemporal distance relations. But this leaves much open. In particular, it leaves open the question of just what sort of entity material objects are, and what it is in virtue of which they bear various distance relations. Perhaps the only fundamental, concrete objects there are are the material ones. And perhaps the way to explain spatiotemporality – the fact that material objects are separated by various spatiotemporal distances – is in terms of relations that these material things bear directly to one another. This is *relationalism*. Relationalism is typically traced back to Leibniz, and has enjoyed the support of people like Mach, and, at times, Einstein.

But perhaps relationalists have it wrong. Perhaps the relational ontology omits an entire category of entity. Perhaps material objects are not the only sort of fundamental, concrete objects; perhaps spacetime (or its basic parts) is also a fundamental, concrete object. This is *substantivalism*. On this sort of view, spatiotemporality is in the first place a feature of spacetime, and it is in virtue of their relationship to spacetime that material objects are spatiotemporal. This view also has an impressive pedigree, which includes people like Newton, and, at times, Einstein, as well as most contemporary philosophers.

Given substantivalism, there is a further question regarding the nature of the relationship between material objects and parts of spacetime. Is this relationship like the relationship between a muffin-tin and its muffins – the former being distinct from the latter, but with the two united by a relation of occupancy. Or is it more like the relationship between an eddy in a body of water – the former being something

like a property of the latter. Substantivalists who favor the muffins-in-muffin-tins analogy are *dualists*. According to dualists, material objects and spacetime are fundamentally different sorts of thing, the two are united by an occupation relation that binds them together and underwrites their collaboration in composing physical reality. Substantivalists who favor the eddies-in-water analogy are *supersubstantivalists*. Supersubstantivalists think that spacetime (or its basic parts) is the only fundamental, concrete object, and that material objects are reducible to regions of spacetime and their properties.

Each of the following chapters is more or less self-contained. But cumulatively they lead to the conclusion that supersubstantivalism is the best metaphysical theory of spacetime. In the first chapter I consider some of the history of the dispute between substantivalists and relationalists. I offer my own, preferred response on behalf of the substantivalist to the challenges posed by Leibniz in his famous correspondence with Clarke. Then in the second chapter I offer a response to the most famous contemporary challenge to substantivalism: the hole argument. I agree with many others that substantivalism escapes unscathed from the argument. But I disagree about the overall impact of the argument. Most have interpreted it as placing pressure on substantivalists who also accept haecceitism. I appeal to a novel conception of the relationship between the mathematical models of a physical theory and the worlds at which that theory is true to show that this interpretation is wrong: there is no pressure at all on substantivalists who also accept haecceitism.

The first two chapters comprise my defense of substantivalism. In the following two chapters, I argue that supersubstantivalism is the best version of substantivalism. In the third chapter I argue against dualism by showing that the conjunction of dualism and plausible principles of modal plenitude entail a range of troubling consequences. In the fourth and final chapter I explain that supersubstantivalism is not a single view, but a family of views. I then show that different members of the family have different

strengths and weaknesses. I conclude by arguing for a version of supersubstantivalism – *compresence supersubstantivalism* – that is compatible with the possibility of co-location.

...

TABLE OF CONTENTS

	Page
ACKNOWLEDGMENTS	v
ABSTRACT	vi
PREFACE	vii
 CHAPTER	
 1. SPATIOTEMPORAL STRUCTURES AND LEIBNIZ AGAINST SUBSTANTIVALISM	 1
1.1 Absolute Space and Time and Leibniz' Attack	1
1.1.1 Symmetry and Covariance Groups	2
1.1.2 Absolute Space and Time	3
1.1.2.1 Absolute Space	4
1.1.2.2 Absolute Time	8
1.1.2.3 Symmetries of Absolute Space and Time	12
1.1.2.4 Mechanical Symmetries In Newtonian Physics	13
1.2 Leibniz' Shift Arguments	15
1.2.1 Qualitative and Non-Qualitative Properties	15
1.2.2 The Static Shift	16
1.2.2.1 The PII Argument	17
1.2.2.2 The PSR Argument	17
1.2.3 The Kinematic Shift	17
1.2.3.1 Leibniz' Argument	18
1.3 Responding to the Shifts	18
1.3.1 Neo-Newtonian Spacetime and the Kinematic Shift	18

1.3.2	The Static Shift and The Modal Commitments of Substantivalism	20
1.3.2.1	Possible Worlds, Haeccetism, and Antihaeccetism	20
1.3.3	The Static Shift	27
1.3.3.1	Haeccetism: Biting The Bullet	27
1.3.3.2	Counterparts of Parts of Spacetime	28
1.4	Conclusion	33
2.	CLIMBING OUT OF THE HOLE	34
2.1	Introduction	34
2.2	Preliminaries	36
2.2.1	Worlds, Models, and Points	36
2.2.2	Properties and Relations	37
2.2.3	Tensors and What They Represent	39
2.2.4	Individualistic and Non-Individualistic Differences	42
2.2.5	Notation	43
2.2.6	Diffeomorphism	44
2.3	What do GR Models do?	46
2.4	The Hole Argument	48
2.4.1	The Hole Argument Against Substantivalism	51
2.4.2	The Argument Against Haeccetistic Substantivalism	53
2.4.3	The Argument Against Haeccetistic Determinism	54
2.5	The Robust Conception	55
2.5.1	Determination	55
2.6	Consequences of The Robust Conception	57
2.6.1	Determinism	58
2.6.2	Haeccetism and Antihaeccetism	59
2.7	Conclusion	60
3.	DUALISM AND THE PRICE OF HARMONY	62
3.1	Introduction	62
3.2	Preliminaries	63

3.3	Occupation	65
3.4	Harmony	66
3.4.1	The Shapes of Objects and their Regions	68
3.5	Modal Plenitude	69
3.5.1	A Plenitude of Structures	70
3.5.2	A Plenitude of Recombinations	72
3.5.3	The Principle of Solitude	73
3.6	Problems for Exact Occupation	73
3.6.1	Against Imposition	73
3.6.2	Against Inheritance	76
3.6.3	Against Independence	78
3.6.3.1	Independence, Harmony and PR are Inconsistent	79
3.6.3.2	Restricting PR	80
3.6.3.3	Preserving Harmony?	81
3.7	Conclusion	83
4.	PROSPECTS FOR SUPERSUBSTANTIVALISM	84
4.1	Introduction	84
4.2	Preliminaries	86
4.2.1	Spacetime	86
4.2.2	Natural Properties and Relations	87
4.2.3	Modality	88
4.2.4	The Occupation Relation	89
4.3	Supersubstantialisms	90
4.3.1	Modest Supersubstantialisms	92
4.3.1.1	Identity Supersubstantialism	92
4.3.1.2	Constitution Supersubstantialism	96
4.3.2	Radical Supersubstantialism	97
4.4	Advantages of Supersubstantialism	99
4.4.1	Field Theories	100
4.4.2	Parsimony	103

4.5	Problems for Supersubstantivalism	106
4.5.1	Counterintuitiveness	106
4.5.2	The Argument From Common Sense	108
4.5.2.1	Modest Supersubstantivalism and Neutral Monism	109
4.5.2.2	Radical Supersubstantivalism and the <i>A</i> <i>Posteriori</i>	110
4.5.3	The Argument From <i>De Re</i> Modal Difference	112
4.5.3.1	The Constitution View	113
4.5.3.2	The Identity View	114
4.5.4	Co-Location	116
4.5.4.1	The Argument from Causal Powers	119
4.5.4.2	The Argument from Quantum Mechanics	121
4.5.4.3	Co-Locating Bosons	123
4.5.4.4	Co-Located Bosons and Modal Plenitude	125
4.5.5	Supersubstantivalism and Co-Location	127
4.5.5.1	Tollensing The Argument	128
4.5.5.2	Supersubstantivalism and The Compresence Relation	129
4.6	Conclusion	130

APPENDICES

A. MANIFOLDS	132
-------------------------------	------------

B. MODEL-DETERMINISM AND WORLD-DETERMINISM	135
---	------------

BIBLIOGRAPHY	143
-------------------------------	------------

CHAPTER 1

SPATIOTEMPORAL STRUCTURES AND LEIBNIZ AGAINST SUBSTANTIVALISM

Modern substantivalism is typically traced back to the conceptions of space and time Newton appealed to in developing his physics. In this chapter I will discuss these conceptions, and examine two of Leibniz' famous arguments against them: the so-called *static shift* and *kinematic shift* arguments. We will find that these arguments prompt different replies from the substantivalist. The kinematic shift will show a substantivalist that he must give up Newton's Absolute space and time and accept a different sort of spatial and temporal structure instead. The static shift, on the other hand, requires a substantivalist to clarify his modal commitments.¹

1.1 Absolute Space and Time and Leibniz' Attack

The Ancient Greeks believed that different sorts of material obeyed different sorts of natural laws. Thus, it was thought to be of the nature of Earth and Water to move in a straight line towards the center of the universe, of the nature of Fire and Air to move in a straight line away from it, and of the nature of *aether* to uniformly rotate about it. Different things have different sorts of natural motion.

Part of what was revolutionary about the physics Newton offered as an alternative to that of the ancient Greeks was that its central claims did not discriminate between

¹Note that it is usually more convenient to speak of material objects as though they occupy, rather than are identical, various regions of spacetime. In the second part of the dissertation, I will defend an identity account over an occupancy account. In the interim, however, I'll help myself to the language of the occupancy account.

the various sorts of matter. Matter – *all* matter – obeys certain laws of motion. One of the best known of these – and the only one that need detain us – is Newton’s First Law.

First Law: Every body perseveres in its state of either rest or of uniform motion in a straight line, except insofar as it is compelled to change its state by impressed forces.²

The philosophical question is this: How must physical reality be composed so as to make the First Law possible? What sorts of structures does it presuppose? In particular, how are we to cash-out the notions of *rest* and *uniform motion in a straight line*?

Newton believed that in order to underwrite the First Law (and also subsequent Laws of motion), reality must include both *Absolute space* and *Absolute time*. Leibniz thought that these structures were metaphysically objectionable, and so rejected the physics on which they were based (although he failed to ever provide a satisfactory alternative). In this section we will examine the structures of Absolute space and time and consider Leibniz’s argument against them. In the following section, we will consider replies.

1.1.1 Symmetry and Covariance Groups

Distinguish a *geometric symmetry* from a *mechanical symmetry*. To do this, first note that an *isometry* is a distance-preserving map between geometric structures. Now let us say that an isometry from a geometric structure \mathcal{G} to itself is a *geometric symmetry*, and that the set of all geometric symmetries of \mathcal{G} is the *symmetry group* of \mathcal{G} . Roughly, the symmetry group associated with a geometric structure of a certain type – for example, a two dimensional sphere – is constituted by the range of actions

²This is Maudlin’s (2012, p. 4) translation.

that may be performed on that structure without changing the properties that make it count as a thing of that type – that is, all the stuff you can do to a sphere without making it something other than a sphere. A mechanical symmetry, on the other hand, has to do with the forms equations governing a physical system may take. For example, an equation governing a physical system involving coordinates displays a mechanical symmetry when there are alternative choices of coordinates which preserve the form of the equation. When such alternatives exist, the equation is said to be *covariant* with respect to this class of alternatives, and the *covariance group* of the equation is the set of all alternatives that leave the form of the equation unchanged.

A useful way to look at the dispute between Newton and Leibniz – and certainly a good way to characterize the general dispute between substantivalists and relationalists – is as one that concerns the relationship between the class of geometric symmetries and the class of mechanical symmetries associated with a particular physical theory. For example, a typical anti-substantivalist line of attack proceeds by charging the substantivalist with accepting an interpretation of a physical theory wherein the group of geometric symmetries associated with that theory form a proper subgroup of the group of mechanical symmetries associated with that theory. This is thought to be problematic because it is only mechanical phenomena and not the geometric backdrop of a theory, that can be associated with observable quantities. So, if there is more geometric structure than required to ground the mechanics, we seem to have an ontology that is needlessly bloated. And, the thought continues, since a lean ontology is *ceteris paribus* better than a bloated one, we ought to reject the substantivalist ontology. Now to the details.

1.1.2 Absolute Space and Time

In his Scholium to the Definitions in *Principia*, Newton claims that “Absolute, true, and mathematical time, of itself, and from its own nature, flows equally without

relation to anything external [and] . . . Absolute space, in its own nature, without relation to anything external, remains always similar and immovable.” These claims are representative of a picture that is standardly interpreted³ as Newton’s endorsement of the view that Absolute space has the structure of an instantaneous, 3-dimensional, infinite Euclidean space, and that Absolute space persists through Absolute time. The following discussion will be an explication of these concepts.

1.1.2.1 Absolute Space

An n -dimensional Euclidean space – \mathbb{E}^n for short – is a special sort of set of points. It is a set of points that possesses *topological*, *affine*, and *metrical* structure.

We begin with topology. Consider a circle drawn on a piece of paper. Now consider two sets: the set of points that lie strictly *within* the interior of the circle, and the set of points that are within or *on* the circle. The first is an example of an **open set**, the second is an example of a **closed set**. The open set does not include any boundary points, the closed set does. Given the notion of an open set, we may characterize topological spaces in the following terms.

Definition 1.1.1. When X is a non-empty set, a set τ of open subsets of X is a **topology** on X if and only if:

1. $\{\emptyset\} \in \tau$, and $\{X\} \in \tau$.
2. τ is closed under arbitrary unions.
3. τ is closed under finite intersections.

When these conditions are satisfied, subsets of X whose complements are in τ are the **closed sets**, and the pair (X, τ) is a **topological space**.

³*C.f.*, Earman (1989), and Pooley (2012)

Defining a topology on X allows us to ground certain important distinctions. Chief among these is the distinction between *continuous* and *discontinuous* paths through a topological space. To illustrate, imagine a lump of clay. Suppose I were to stretch the lump into something resembling a snake, and then join the ends together to form a ring. When I do this I change the open set structure of the points that compose the lump of clay: whereas prior to the joining there were boundaries at both the tail and the head, and so certain closed sets at the tail and the head, after the joining, these boundaries, and so the corresponding closed sets, were removed, leaving open sets in their place. Likewise, if we imagine that I first reform the clay into a lump, and then tear a hole in the middle, we can see that I change the open set structure of the clay, and so its topology, by tearing: where certain open sets populated the area around the tear, there are now certain closed set – sets with boundaries – in that area.⁴

We can extend the notion of open and closed set structure to the notion of continuity by thinking of the boundaries included in closed sets as “barriers” that paths cannot cross. In the case of the snake and the ring, it is clear to see that in the process of changing the open set structure by joining I also change the set of continuous paths. Consider a point a near the head of the snake, and a point b near the tail. Prior to the joining, the only way to get from a to b is to traverse the length of the snake. After the joining, there is a path between a and b that exploits the newly-joined head and tail. Similarly, we can see that in the case of the lump with the hole, certain continuous paths within the lump are rendered discontinuous after creating the hole.

A further important feature of topological spaces regards criteria for equivalence. Topological spaces are topologically equivalent when there is a mapping between them – called a *homeomorphism* – that preserves open and closed set structure. A

⁴This example assumes a connection between the metric and topological structures of spaces. Namely, that the topology of the space is the topology induced by the metric. See fn. 5 of this chapter.

homeomorphism is a mapping ϕ between topological spaces (X, τ_X) and (Y, τ_Y) where ϕ assigns every open subset of (X, τ_X) to an open subset of (Y, τ_Y) and every closed subset of (X, τ_X) to a closed subset of (Y, τ_Y) , in such a way that every subset of (Y, τ_Y) is mapped to by exactly one subset of (X, τ_X) . What is interesting about these criteria for equivalence is that they count things which seem intuitively to be very different – like spheres and cubes – as topologically equivalent.

Next we need to characterize the notion of affine structure over our set X . Here a rough, intuitive gloss will do. Affine structure over a set X grounds the distinction between straight and bent lines through X . Notice that such a distinction does not fall out of the definition of a topological space in terms of open sets – if a sphere and a cube are topologically equivalent (they are), then topology must not recognize a distinction between the straight lines of the cube and the bent lines of the sphere. But we all know that it is central to Euclid’s geometry that there be a notion of straightness – the first two postulates make explicit reference to straight lines. So if this notion is not supplied by the topology, we’ll need to either posit the affine structure of a space as *sui generis* structure or we’ll have to show that it falls out of some other structures associated with our space.

The third geometric structure we require is that given by a *metric* on \mathbb{E}^n . We approach this idea via the following definition.

Definition 1.1.2. Let X be a non-empty set and let d be a function $d: X \times X \rightarrow \mathbb{R}$. If d is such that for all $(u, v) \in X$:

1. $d(u, v) \geq 0$, with $d(u, v) = 0$ iff $u = v$ *Positive definiteness*
2. $d(u, v) = d(v, u)$ *Symmetry*
3. $d(u, v) + d(v, w) \geq d(u, w)$ *Triangle Inequality*

then d is a **metric** on the set X , and the pair (X, d) is a **metric space**.⁵

Now, intuitively, a metric on a set is like physical distance between places as measured by a ruler. But a bit of care is required in unfolding the potential physical significance of a metric. To begin, note that each element of the codomain of d is simply some $n \in \mathbb{R}$, where n has no units attached to it: n could represent miles, meters, minutes, or moles. The metric is silent about units. What it allows us to do, however, is calculate relative magnitudes. If $d(u, v) = 2$ and $d(s, t) = 8$ then we know that $d(s, t)$ is four times $d(u, v)$. The metric is also silent about scale in the following sense. If $d(u, v) = 2$ and $d'(u, v) = 10$, then d and d' disagree with respect to (u, v) by a scale factor of 5. If d and d' differ by a scale factor of 5 for all elements in their domain, then d and d' represent the same metrical structure. The consequence of this is that differences in scale factor of a metric function have no real significance for the spaces they describe.

So far we have been discussing various geometric properties of \mathbb{E}^n . These geometric properties have fairly straightforward physical significance: topological structure grounds the distinction between continuous and discontinuous motion in physical space; affine structure grounds the distinction between inertial and accelerated motion in a physical space; metrical structure grounds distance relations in physical space. But there is a great deal missing here from the perspective of physics. To actually do physics, it is convenient to be able to calculate quantities, to add and subtract things, multiply and divide things, differentiate and integrate things, etc. To do this, one typically relies on a systematic link between numbers – more specifi-

⁵A quick note about metric spaces. Given a metric d on a space X , we may define a topology on X in terms of d by thinking of the open set structure as induced by d . For example, we may specify a topology over X by appeal to **open balls**, where an open ball centered on $y \in X$ is the set of all points $x \in X$ such that, for some specified $n \in \mathbb{R}$, $d(x, y) < n$. This is perhaps a more intuitive avenue for characterizing the notion of open set structure, but it obscures somewhat the fact that the set of metric spaces forms a proper subset of the set of topological spaces.

cally an algebraic structure – and geometric structure. We can get this by introducing coordinates.

Although there are many ways to coordinatize a space, a particularly convenient way is to appeal to the real numbers. We do this by putting the elements of \mathbb{E}^n into one-to-one correspondence with elements of \mathbb{R}^n , and we require that this correspondence satisfy several constraints. First, we want the topology of the physical space and the topology of the coordinate space to be related in such a way that continuous paths through physical space correspond to continuous changes in coordinates. Second, we want our coordinate space to reflect the affine structure of physical space. We can get this by requiring that the coordinates of straight lines in physical space satisfy linear equations. Finally, we want the metric structure of our coordinate space to reflect the metric structure of our physical space. One way to do this is by using the familiar Euclidean metric function. Where x , y , and z are variables that range over Cartesian coordinates of points p and q under a given coordinatization in a space with dimension 3, the distance between points p and q is given as follows:

$$d(p, q) = \sqrt{(x(p) - x(q))^2 + (y(p) - y(q))^2 + (z(p) - z(q))^2}. \quad (1.1)$$

So, to sum up: Newton’s Absolute space is a set of points endowed with the topological, affine, and metrical structure of \mathbb{E}^3 . Moreover, the standard coordinatization of \mathbb{E}^3 involves using \mathbb{R}^3 in a way that satisfies the constraints discussed above. Now what about time?

1.1.2.2 Absolute Time

Absolute time, like Absolute space, is a set of points. We call these points ‘instants’, and we call paths through this set ‘intervals’. Also like space, the set of instants may be endowed with various geometric structures. In particular, the set of instants has a topology, which grounds facts about continuous and discontinuous

intervals, and it is equipped with a metric, which determines the relative magnitudes of intervals. Finally, there is an important relationship between Absolute space and time: Absolute space *persists* through Absolute time. That is, somehow or other Absolute space exists at multiple distinct instants of Absolute time. Let's begin by clarifying how this might happen.

Let us say that an object *perdures* if and only if it persists by possessing distinct parts at distinct times. A perduring object, then, is spread out temporally in much the way that objects are spread out spatially; just as I have feet and hands for spatial parts, I have infant, teenage, and retiree temporal parts. Let us juxtapose perdurance with *endurance*. We'll say that an object endures if and only if it persists by being multiply located in time, and "wholly present" at every time at which it is located. To clarify this a bit more, consider the difference between perdurantist and endurantist analyses of change.

According to a perdurantist, objects can gain and lose parts and undergo change with respect their intrinsic properties by having temporal parts that differ with respect to which parts they possess and which intrinsic properties they instantiate. So the change I undergo, for example, from being a 3 foot tall toddler to a 6 foot tall adult is to be explained by me – the temporally extended continuant – having some temporal parts that are 3 feet tall and other temporal parts that are 6 feet tall. According to the endurantist, on the other hand, change in parts and with respect to intrinsic properties is to be explained by taking parthood and certain properties to be relations born between objects and times. So, for example, the endurantist analysis of the change I undergo from being a 3 foot tall toddler to a 6 foot tall adult will involve recasting my height properties as relations I bear to times: I bear the three-feet-at relation to some times, and I bear that six-feet-at relation to some other times.

(Perdurantism is sometimes called a "crazy metaphysics" since, for example, it entails that when I hold a piece of chalk in my clenched hand from 12:05 to 12:06,

there is something that I have in my hand at 12:06 that wasn't there at 12:05 despite my hand remaining clenched – to wit, some temporal part of the chalk. There is a similarly intuition-based complaint about endurantism: it too is crazy since, for example, it entails that an apple's being red involves something other than just the apple and the color red – to wit, the time relative to which the apple is red. These sorts of complaints should be given no weight at all. For surely the “intuitive” position to take about persistence entails, for example, that my change from a short toddler to a tall adult is one according to which it is *me* that has various height properties (not temporal parts of me), and the height properties I have are had by me *intimately* (they are not relations between me and anything else). Insofar as this is the folk intuition, both perdurance and endurance are strange views. But this doesn't bear against either of these views, rather it shows that the folk view is incoherent!

Back to Newton. In order to ground his claims about Absolute motion, Newton requires that the points of space persist through time. Newton's view that Absolute space “remains always similar and immovable” suggests that he himself thought of points as enduring.⁶ Be that as it may, Newton's conception of Absolute space and time is compatible with either modern view about the nature of persistence. But the difference is worth mentioning. If one views Newton's picture through the lens of the endurance view, then the arena in which the mechanics of the universe unfold is an infinite, three dimensional euclidean space that is multiply located in time. This means that the very same points of Absolute space exist across distinct instants of Absolute time. This numerical identity between points at distinct times gives the endurantist a tidy way to characterize Absolute motion. First, characterize Absolute rest in terms of the identities of points:

⁶*C.f.*, Maudlin (2012).

an object O experiences **Absolute rest** between times t and t' iff the set of points S of Absolute space that O occupies at t is identical to the set of points of Absolute space that O occupies at t' , and there is no time t'' between t and t' such that O fails to occupy S at t'' .

Then define Absolute motion in terms of Absolute rest:

an object O experiences **Absolute motion** between times t and t' iff O does not experience Absolute rest between t and t' .

Under the assumption of the perdurance view, points at distinct times are not numerically identical. This means that a perduring object cannot satisfy the endurantist characterization of Absolute rest, and that means that the perdurantist requires some other means of characterizing Absolute motion. So, to begin, if one views Newton's picture through the lens of the perdurance view, then the arena in which the mechanics of the universe unfold is a set of copies of an infinite, three dimensional euclidean space – one copy at each instant of time. So while the endurance view is one of a single space “sweeping through” time, the perdurance view is of successive spaces “stacked up” time-wise. Another way to put the point is that while the endurance view is compatible with thinking of Absolute space and time as rather distinct things, the perdurantist is committed to viewing Absolute space and time as a product space: a four-dimensional spacetime that is the product of Absolute space and Absolute time. There is a consequence here for the perdurantist. In order to ground Absolute motion in the absence of numerical identity between points at distinct times, the perdurantist needs to claim that there are certain brute facts concerning which non-simultaneous points are the “same” points. Once these facts are fixed, the perdurantist can proceed to offer a definition of Absolute motion analogous to that offered by the endurantist, save for the substitution of the perdurantist's notion of sameness of points for the endurantist's identity between points.

Note: while the endurantist account of Absolute motion can ultimately be cashed out in terms of the identity of points, the perdurantist account requires certain brute facts about the sameness of non-simultaneous points. This, because it is a sort of theoretical economy, looks like an advantage for the endurance view. On the other hand, endurantists require a fundamental distinction between spatial and temporal points in order to cash out their account of Absolute motion. Perdurantists require no such distinct; they treat all points simply as spacetime points. This, because it is a sort of theoretical economy, looks like an advantage for the perdurance view. We seem to have a tie between endurantism and perdurantism in the context of Absolute space and time.

1.1.2.3 Symmetries of Absolute Space and Time

Because Absolute space has the geometry of \mathbb{E}^3 , it inherits both the translational and rotational symmetries that \mathbb{E}^3 possesses. The translational symmetry of \mathbb{E}^3 ensures that the geometric properties of objects do not depend on which regions of \mathbb{E}^3 they are embedded in, and the rotational symmetry of \mathbb{E}^3 ensures that the geometric properties of objects do not depend on their orientation in \mathbb{E}^3 . A space that possesses translational symmetry is *homogenous*, one that possesses rotational symmetry is *isotropic*. So, because Absolute space is homogenous, the intrinsic properties and relations of material objects are not affected by their placement in regions of Absolute space: as far as the intrinsic properties and relations of material objects are concerned, one region of Absolute space is just like any other. Because Absolute space is isotropic, the intrinsic properties and relations of material objects are also unaffected by direction – as far as the intrinsic properties and relations of material objects are concerned, orientation in one direction of Absolute space is just like orientation in any other. Absolute time also possesses translational symmetry, and it is this symmetry that ensures that the intrinsic properties and relations of material

objects do not depend upon their location in Absolute time. As we will see, these geometric symmetries will be exploited by Newton's opponents.

1.1.2.4 Mechanical Symmetries In Newtonian Physics

It will be easiest to approach the idea of mechanical symmetry by introducing the notion of a *frame of reference*. To do this, we take a short detour. Recall Galileo's famous thought experiment concerning the relativity of inertial motion. We are to suppose that we are in a windowless cabin below deck on a ship. In the cabin with us are some things like balls and butterflies. Now we consider two scenarios. The first scenario involves the boat docked near shore, the second involves the boat moving with constant velocity – that is, with change in neither speed nor direction. Apparently, nothing we could do below deck would reveal to us our state of motion. For, if the boat is moving with uniform velocity, then the butterflies fly about freely – they are not forced back to the rear of the cabin – but this is just as it would be if you were docked near shore. If the boat is moving with uniform velocity, then if you hold one of the balls and drop it to the floor, it falls straight down rather than in front or behind you, but this is just as it would be if you were docked near shore. What does this show? It shows us that inertial states of motion enjoy a sort of equivalence. In particular it shows us that, like the outcomes of the experiments in the above scenarios, the laws of nature themselves should not vary between states of inertial motion. Indeed, Newton's Second Law, $\mathbf{F} = m\mathbf{a}$, is invariant in just this sense.

Suppose that you are on the deck of a ship and that I stand on the shore and watch you sail away with constant velocity. Now suppose that I want to lay down a Cartesian coordinate system with me as the origin, and that you want to lay down one with you as the origin, but we otherwise agree on the orientation of axes and the passage of time. This is a roundabout way of saying that while we agree about global things like directions and the passage of time, I think of myself as a stationary

point watching as you move away, while you think of yourself as a stationary point watching as I move away. Let us suppose that this perceived motion occurs along what we've agreed is the x -axis of our Cartesian coordinate systems.

Now consider the coordinate system (x, y, z, t) that I lay down with me as the origin. (x, y, z, t) is a *frame of reference*, it is *my* frame of reference, the one according to which I and my immediate surroundings are at rest. Next consider the coordinate system (x', y', z', t') that you lay down with you as the origin. (x', y', z', t') is also a frame of reference, its is *your* frame of reference, the one according to which you and your immediate surroundings are at rest.

Now it turns out that there is a special relationship between our frames of reference. This relationship, where \mathbf{v} is a velocity vector in the x direction, is given by (Gal).

$$x' = x - vt, \quad y' = y, \quad z' = z, \quad t' = t \quad (\text{Gal})$$

When the frame whose inertial coordinate system is (x', y', z', t') (this is your frame) is related to the original frame whose coordinates are (x, y, z, t) (this is my frame) by (Gal), the primed frame is said to be related to the unprimed frame by a *Galilean transformation*.

Now, it turns out that Newtonian mechanics is invariant under the group composed by the Galilean transformations. That is, calculations employing Newton's laws do not differ between coordinate systems related by (Gal). But notice that Absolute space does not possess a geometric symmetry that corresponds to (Gal). That is, while it follows from the homogeneity and isotropy of Absolute space and time that the state of motion of an object may be described by a coordinate system whose point of origin is anywhere and anywhen, and whose axes are oriented in whichever way one chooses, it does not follow from these symmetries that the state of motion of a particle can be described equally well by two different coordinate systems related by

(Gal). If it could be, then there would be no facts about Absolute motion. What this means is that the symmetries of Absolute space and time form a proper subgroup of the symmetries of Newtonian mechanics which include not just the symmetries that arise from the homogeneity and isotropy of Absolute space and time, but also the symmetry group that arises from Galilean transformations. In short, there is more geometric structure to the Newtonian regime than is licensed by or strictly required by its mechanics. As we'll see in the following section, this idea was famously exploited by Leibniz.

1.2 Leibniz' Shift Arguments

Now we'll consider Leibniz' famous attacks on the Newtonian regime. We begin by noting that there are two principles that form the background of Leibniz' attacks. The first is a theological principle called the Principle of Sufficient Reason (PSR), which we may take to amount to the following:

PSR: for every act A , if God performs A , then there is a sufficient reason for God's having performed A , rather than some alternative act B .

The second principle is the Principle of The Identity of Indiscernibles (PII). To properly formulate this principle, we first need to consider a distinction between *qualitative* and *non-qualitative* properties.

1.2.1 Qualitative and Non-Qualitative Properties

It is more or less standard to characterize non-qualitative properties⁷ as properties that concern particular individuals, and to characterize qualitative properties as those that do not.⁸ The property *being Obama*, for example, is a non-qualitative property

⁷For simplicity I'll confine my talk to qualitative and non-qualitative properties, although similar points apply to relations as well.

⁸*C.f.*, Hawthorne (2006), Skow (2005, 2008).

since it concerns a particular individual, whereas *being president* is a qualitative property since it doesn't concern any particular individual. Another way to characterize the distinction between qualitative and non-qualitative properties is in terms of the distinction between the natural properties and relations and all the others. We can say that qualitative properties are those that are built from the natural properties and relations, and the non-qualitative ones are not.⁹ More examples of qualitative properties include properties like *being green*, *being round*, *being negatively charged*, and so on. Examples of non-qualitative properties include properties like *being in the same room as Ben* and *living just outside of Tucson*.

Given this notion of qualitative properties, we may formulate the PII as follows.

PII: for any objects a and b , if, for any qualitative property P , a has P if and only if b has P , then a is identical to b .

Leibniz' assault on Absolute space proceeds along two routes. The first involves showing that Absolute space entails certain possibilities that jointly run afoul of the PII, the second involves the charge that Absolute space in conjunction with the PSR makes a Buridan's Ass of God.

1.2.2 The Static Shift

Suppose that the material universe has a specific location in some region of Absolute space. Call this region ' R .' Now consider the possibility of the material universe having been placed in some *other* region R^* of Absolute space. Call the possibility of the material universe having been located in R 'possibility #1', and call the possibility of the material universe having been located in R^* , 'possibility #2'. Now assume that possibilities #1 and #2 are related by either a rigid rotation or a rigid translation.

⁹See Bricker (1996) for a very clear explanation.

1.2.2.1 The PII Argument

According to friends of Absolute space, possibilities #1 and #2 are distinct because, by hypothesis, they involve different distributions of matter across regions of Absolute space. However, possibilities #1 and #2 do not differ qualitatively, they differ *only* in which regions of Absolute space are occupied by matter. Because there is no qualitative difference between possibilities #1 and #2, it follows from PII that possibilities #1 and #2 are identical. So, it follows from the assumption of Absolute space and the PII that possibilities #1 and #2 are both identical and distinct: *reductio!*

1.2.2.2 The PSR Argument

A second line of attack employs the PSR. First, because the only differences between possibilities #1 and #2 are non-qualitative differences related to the distributions of matter over regions of Absolute space, God could have no reason to choose one possibility over another as the one to make actual. But, by PSR *every* act is such that if God performed it, then He had a sufficient reason for so doing. Thus, the fact of the actual world's existence shows that God had a sufficient reason for creating the actual world as it is. So the existence of the actual world undermines Absolute space, since if space were Absolute, God would have been frozen by indecision between indistinguishable possibilities at the time of creation, and the world would never have been created.

1.2.3 The Kinematic Shift

Suppose that the material universe has some state of Absolute motion through Absolute space and time. Call this state '*S*.' Now consider the possibility that it have some *other* state of motion, '*S**.' Now suppose that the velocity change in this case is uniformly distributed across the entire material universe so that the relative velocities and relative distances amongst the various material parts of the universe remain the

same. Call the possibility that the material universe be in state S , ‘possibility #1,’ and call the possibility that the material universe be in state S^* , ‘possibility #2.’

1.2.3.1 Leibniz’ Argument

Leibniz’ argument proceeds, as in the case of the static shift, by pointing out that by PII possibilities #1 and #2 collapse, and by PSR the fact that the universe exists at all is evidence that God was not forced to choose between indiscernible possibilities. In either case, there is no Absolute space.

1.3 Responding to the Shifts

I do not think that many contemporary philosophers are committed to the PSR. So I’ll assume that the arguments against Absolute space that appeal to the PSR as a premise can be rejected. (Although, perhaps it’s worth noting this much: if we suppose that there are infinitely many possible velocities for the material universe for God to choose from, it doesn’t necessarily follow from the PSR that God would be unable to pick one. What about a velocity of 0? 0 seems to be, at least in a way, unique among possible velocities.) But this still leaves the arguments that appeal to the PII. In the following sections these will be our focus.

1.3.1 Neo-Newtonian Spacetime and the Kinematic Shift

The real force of the kinematic shift argument comes from the methodological principle according to which, when it comes to unobservable physical quantities – like absolute motion – its not that we could never be justified in making do with them, it’s rather that “[c]eteris paribus, one would prefer a theory without them” (Maudlin, 1993, p. 193). So how can we do without them? One way, presumably, is to be a relationalist. Another way is to concede that, although the kinematic shift shows absolute velocity to be a suspect quantity, this does not undermine substantivalism. Instead, it shows that certain of the structures posited as part of the substantival-

ist picture need to be rejected. Namely, those structures that underwrite absolute motion.

Recall that Newton's picture – packaged with absolute motion – involves a three-dimensional set of spatial points endowed with various intrinsic structures persisting identically through time. As I mentioned, we could alternatively view the arena of Newtonian physics the way that a perdurantist does – as a single, four-dimensional set of points endowed with brute facts about cross-temporal spatial distance relations. In either case, the trouble-causing piece of structure is the one that grounds claims about spatial intervals between non-simultaneous points of space. To put this point more simply: the problem comes from having a notion of being at the same place at different times. Now, for the spatial endurantist, a notion of being at the same place at different times is built directly into their metaphysics – it falls directly out of the claim that identically the same points exist at different times. So getting rid of this piece of structure means letting go of spatial endurantism.

Once we've taken spatial endurantism off the table, the solution becomes apparent. We accept spatial perdurantism, but reject the existence of brute facts about spatial distances between non-simultaneous points. On this picture, space and time are a four-dimensional set of points within which there are well-defined temporal intervals between all points, and also well-defined spatial intervals between simultaneous points, but where there are no well-defined spatial intervals between non-simultaneous points. The resulting spacetime is called *Neo-Newtonian spacetime* (or, alternately, *Galilean spacetime*). The move from Newtonian Absolute space and time to Neo-Newtonian spacetime brings the geometric symmetries of spacetime into alignment with the mechanical symmetries of Newtonian physics and is perfectly compatible

with substantivalism since it still involves treating spatiotemporal points as fundamental constituents of reality. This is an easy fix.¹⁰

Before moving on I want to make note of a shift in terminology. Throughout the remainder of the chapter I'll drop the terms 'Absolute space' and 'Absolute time' and instead use 'substantival spacetime' with the understanding in place that substantival spacetime has Neo-Newtonian structure. I think the kinematic shift argument does undermine Absolute space and time, but, as I've just shown, this doesn't undermine substantivalism since one can be a substantivalist about Neo-Newtonian spacetime. So in what follows we'll consider what sort of impact the static shift has on the plausibility of substantival spacetime.

1.3.2 The Static Shift and The Modal Commitments of Substantivalism

In this section I'll cover some relevant background material concerning modality and the question of haecceitism. I'll then consider the sorts of things a substantivalist might say in response to the Static Shift argument.

1.3.2.1 Possible Worlds, Haecceitism, and Antihaecceitism

Consider the following two propositions:

1. Josh is 6 feet and 2 inches tall.
2. $2 + 2 = 4$.

Both (1) and (2) are true, but there is something different about the *way* they are true. It seems like (1) could have very easily been false; things could have gone just slightly different and I would be an inch or two shorter or taller. But it seems like (2) would be true no matter how things had gone. This suggests that propositions don't just come in two flavors: true and false; their truth or falsity also comes in different

¹⁰*C.f.*, Earman 1970, Friedman 1983, Sklar 1974.

modes. Some true propositions are necessarily true – like (2) – and others are only contingently true – like (1). Others are necessarily false – like ‘2+2=7’ – and others are only contingently false – like ‘Josh has blue eyes’. The study of these sorts of modes of truth or falsity is called the study of *modality*.

Some philosophers think that modally qualified claims are nonsense,¹¹ others think that modality is primitive,¹² but most philosophers prefer *reductionism* about modality: modally qualified expressions may be analyzed in non-modal terms.¹³ Most reductionists about modality have found it helpful to begin their analysis of modality by introducing possible worlds. Thus we have the familiar biconditionals.

A proposition p is necessary iff p is true according to every possible world.

A proposition p is possible iff p is true according to some possible world.

A proposition p is not necessary iff p is false according to some possible world.

A proposition p is not possible iff p is false according to every possible world.

So, according to a possible worlds analysis of modality, (1) is contingent iff there are some worlds according to which it is true, and others according to which it is false. Similarly, (2) will be necessary iff it is true according to every world.

But of course these biconditionals don’t help much in explaining modality; they merely shift the question from “what makes statements like (1) and (2) true or false?” to “what is a possible world?” and “what does it mean for something to be true *according to* a possible world?” It turns out that there are many theories about the

¹¹Quine (1951) may be among these.

¹²*C.f.* deRosset (2009).

¹³Strictly speaking, most philosophers working in the metaphysics of modality take themselves to give only *partial* reductions of modality. Possible worlds may be used to analyze necessity and possibility, but typically, somewhere in the account of possible worlds, unreduced modal notions find their way in.

nature of possible worlds – that is, many attempts to answer the question “what is a possible world?” For our purposes, it will be sufficient to distinguish *ersatzism* from *realism* about possible worlds. There are many versions of ersatzism, but what they all have in common is that they seek to construe possible worlds as something decidedly different in kind from our actual world. The ersatzist may take possible worlds to be sets of sentences, propositions, or pictures, or abstract redistributions of properties across regions of spacetime. The realist, on the other hand, takes merely possible worlds to differ from our world only in their local distributions of properties. Possible worlds, for the realist, do not differ in kind from our world; whatever ontological category our world belongs to, all the other possible worlds belong to it, too.¹⁴

Just as there are several ways to understand what sort of thing a possible world is, there are several ways to understand what it means for something to be true “according to” a possible world. Let’s focus our attention on individuals. How might it be true, according to some possible world, that some individual is different from the way that it actually is? To begin, consider (3):

3. Josh might have had blue eyes.

I actually have green eyes, but (3) seems plausible. According to the possible worlds framework we are working within, that means that there is some world according to which I have blue eyes – in other words, there is something in some other possible world that represents me here in the actual world as possibly having blue eyes. But how does that work? Consider some possible accounts.

¹⁴I have in mind here the realism of Lewis (1986), but there are other sorts of realism. Bricker (2006) is an example. He argues that Lewis’ realism is incoherent because it counts all worlds as equally actual. He offers an amendment to Lewis’ view that he takes to solve this problem. The fix involves supplementing Lewisian realism with primitive absolute actuality. The result of this move is a rift in the ontological status of possible worlds: some possible world – namely, *ours* – have the primitive property of being absolutely actual, while the others lack it. With minimal changes, everything I will say about realism can be made compatible with Bricker’s brand of realism.

To begin, let's call the sorts of accounts we're about to consider accounts of *transworld identity*. These are accounts that explain the relationship between things in distinct worlds in virtue of which certain possibilities obtain. Now first suppose that I manage – literally, without any sort of paraphrasing – to exist in multiple worlds. That is, in addition to being here in the actual world, I also manage to be in various other merely possible worlds as well – all of the possible worlds in which I exist overlap by sharing me as a common part. According to this view, transworld identity involves genuine identity. On this sort of view (3) is true because something in a different possible world represents me as having blue eyes, and this something is me, as I am in that other possible world. But there are problems here. Begin with the ersatzist. If you are an ersatzist, then you think that the actual world is different in kind from other merely possible worlds – the actual world is concrete (in some sense of 'concrete') and merely possible worlds are abstract (in some sense of 'abstract'). But if I am a part of the actual world, which is concrete, how could I also be part of a merely possible world that is abstract? That would seem to require that I be somehow unsure about my ontological status: am I abstract, or am I concrete?

Now consider the realist. If you are a realist, then you think that all worlds are ontologically on a par – in whatever sense the actual world is concrete, merely possible worlds are concrete in that sense, too. But now compare the actual green-eyed me, and the possible blue-eyed me. If we are one and the same – literally, numerically identical – then my eyes are both blue and green. Or, if that doesn't seem so incredible, consider the possibilities of me being completely bald, and me having knee-length hair. How could one and the same thing be both bald and have knee-length hair? Indeed, given an overlap account and realism, many objects will turn out to have many different and incompatible properties.

These are not decisive objections to the view that representation of counterfactual possibilities for individuals involves those individuals existing in distinct worlds,¹⁵ but they are reasons to prefer an alternative. The alternative is to deny that objects from different possible worlds are literally, numerically identical, and instead claim that they are bound by some sort of relation. On this sort of view what makes a claim like (3) true will be the holding of some relation between me and some other-worldly, blue-eyed individual. There are two main variant accounts of such a relation. According to one, the relation between individuals in distinct possible worlds in virtue of which they count as “the same” is primitive, according to another, the relation is governed by qualitative similarity.

David Lewis’ *counterpart theory* is a paradigm of qualitative similarity accounts.¹⁶ According to Lewis view, what makes a claim like (3) true is that there is a counterpart of me in some other possible world and this counterpart has blue eyes. Likewise, what makes it the case that I am possibly shorter, taller, balder, fatter, etc., is that I bear counterpart relations to individuals that are shorter, taller, balder, fatter, etc. But what determines whether something is my counterpart?

Whether something is one of my counterparts, and thereby represents a possibility for me, will depend upon qualitative similarities between me and that thing. But qualitative similarity is a somewhat loose notion, one that seems to be sensitive to context. The way to factor this in is to take the possibilities for me (and other individuals) to be sensitive to context: different contexts will make salient different respects of similarity; different respects of similarity will correspond to different counterpart relations; and different counterpart relations will underwrite different possibilities.

¹⁵Indeed, it seems one could accept modal realism and claim that objects have their qualitative properties only relative to worlds in something like the way that an enduring object has its intrinsic properties only relative to times. See McDaniel (2004).

¹⁶Note that Lewis defends counterpart theory in several different guises. One sort of defense can be found in his (1968) and (1971), and another sort in his (1986). Here I follow the views he defends in his (1986) account.

So in a sense, what is possible for any particular thing is going to depend upon certain features of context, like what sorts of background assumptions one holds fixed. For example, if we are discussing physiology, and considering whether or not I might have been 10 feet tall, or had blue eyes, we are probably holding fixed that I am human, so all of the relevant counterparts of me are human. But let the context change. I am playing with toy dinosaurs with my kid and we are talking about what sort of dinosaur she would be. Here we'll hold different sorts of features fixed – for example, certain of her personality traits (like her penchant for stomping around, yelling, and just causing general mayhem: she's a T-rex) – and let her being human vary.

One alternative to counterpart theory is to claim that what makes a distinct individual in a different possible world count as me is simply a primitive relation – a non-qualitative counterpart relation – between me and that thing. Let's call this view *non-qualitative counterpart theory*. According to non-qualitative counterpart theory, my qualitative character plays no role in determining whether I bear such a relation to any individual in another world – it's just a brute fact about reality that certain individuals in distinct worlds are united by a primitive transworld identity relation. Here is a useful way to flesh out the distinction between this view and the counterpart theoretic account.

Ask: Could you and I have switched places? I don't mean switch places like people do in movies, by transplanting my mental life into your body and your mental life into mine. I mean, could we have swapped our qualitative roles *completely*, so that everything is qualitatively exactly as it actually is, but I am you and you are me? Call this a *qualitative swap*. Now, let's suppose that at least some qualitative swaps are genuine possibilities. Now ask: how many worlds would it take to represent these two possibilities – that is, how many worlds would it take to represent the actual state of affairs as well as the one where we switch qualitative roles?

If you are a non-qualitative counterpart theorist, then you will say that it takes two worlds to represent these two possibilities: this world represents the actual states of affairs (by being identical to it), and then it takes another possible world w that is a qualitative duplicate of this one, but where I bear a non-qualitative counterpart relation to the individual in w that occupies the qualitative role you occupy in this world, and you bear a non-qualitative counterpart relation to the individual in w that occupies the qualitative role I occupy in this world. What is needed is one world for each possibility.

Things are different for a counterpart theorist. According to the counterpart theoretic account, in order to underwrite the possibility that you and I swap qualitative roles what is needed is a counterpart relation that links me to an individual that occupies your qualitative role and you to an individual that occupies my qualitative role. But we need not look far for candidates: I am an individual that occupies my qualitative role, and you are an individual that occupies yours, so, given an appropriately permissive counterpart relation – say, one that makes every person a counterpart of every other person – you are my counterpart and I am yours. One world can do double (and triple, and quadruple, and . . .) duty in representing distinct (but qualitatively identical) possibilities.

We just saw how the non-qualitative counterpart theorist requires two qualitatively identical worlds to underwrite qualitative swap possibilities, while the counterpart theorist requires only one. It turns out that there is an important corollary to this; it is the distinction between *haecceitism* and *antihacceitism*. Haecceitism is the view that there are worlds that fail to differ in qualitative respects and yet differ with respect to representation *de re*. Antihacceitism is the view that representation *de re* supervenes on the qualitative character of worlds – it implies that, if there are worlds

w and v that fail to differ qualitatively,¹⁷ then w and v also fail to differ with respect to what they represent *de re*.¹⁸

1.3.3 The Static Shift

To begin, note that while the move from Absolute space and time to substantival spacetime (with Neo-Newtonian structure) defuses the problem associated with the kinematic shift, it does not defuse the problem associated with the static shift. Recall that the static shift argument uses Absolute space and the PII formulated in terms of qualitative properties as premises and then derives a contradiction. The same problem arises if we substitute substantival spacetime for Absolute space and time: we still end up with possibilities – different ways of locating the material universe within substantival spacetime – that differ only non-qualitatively. Now, what a substantivalist ought to say about this will depend on their prior modal commitments.

1.3.3.1 Haecceitism: Biting The Bullet

Note that the Leibnizian uses substantival spacetime and the PII formulated in terms of qualitative properties to derive a contradiction. She then takes this to show that substantival spacetime is suspect. But one could just as well claim that this shows that the PII formulated in terms of qualitative properties is suspect. Indeed, this is what a haecceitist ought to do. A haecceitist will agree that there are infinitely many shift possibilities that differ one from another only non-qualitatively. But she

¹⁷Lewis's (1986) view is supposed to be neutral with respect to whether there are such worlds.

¹⁸I want to flag a potential worry. Skow (2008) has recently argued that haecceitism and antihaecceitism are typically characterized in the literature in ways that presuppose a particular conception of possible worlds. But, so he suggests, haecceitism and antihaecceitism for an ersatzist are not the same thing as haecceitism and antihaecceitism for a realist. So, when presuppositions concerning background modal metaphysics remain unstated, we endanger ourselves to talking past one another. I think he is right. So to avoid this pitfall, I'll just assume realism about possible worlds, and offer a characterization of haecceitism suitable to this context. Most of what I say in subsequent chapters will be easy enough to transpose into the ersatzist regime as long as the transposition also includes appropriate modifications to the formulation of haecceitism.

will not think of this as particularly damning – belief in possibilities that differ only non-qualitatively is just a consequence of her general modal metaphysics.

But not every substantialist will be comfortable accepting this sort of response to the static shift argument. Perhaps one is wed to the PII and sees the tension between it and her substantialism as problematic. What then? It turns out that counterpart theory is quite well-suited to dealing with the problem.

1.3.3.2 Counterparts of Parts of Spacetime

Consider a world w with some material occupants. Now suppose we call the collection of distance relations born between the material occupants of w , R . Next, observe that there is a certain pattern of occupancy relations born between material objects and regions of spacetime at w ; call this pattern O . Now consider the group of transformations of O that preserve R – the rigid rotations and rigid translations. There will be infinitely many of them since space is continuous. Call these ‘the shift possibilities for w ’. A haecceitist will identify each shift possibility with a possible world. An *anti-haecceitist*, on the other hand, will claim that a single world can represent each of the shift possibilities. The question is how the anti-haecceitist pulls this off.

To begin, suppose you want to move a plate from one region of a tablecloth to another region one foot to the right. One way to do this is to pick up the plate, and place it one foot to the right, but another way is to pull the tablecloth one foot to the left (think of the “pull the tablecloth out from under the dinnerware” trick). The effect (at least as far as the relations between the plate and tablecloth are concerned) is the same: either way, the plate ends up at a region of the tablecloth that is one foot to the right of the region where it was. So, at first blush, it seems like there will be two ways for an antihaecceitist to account for shift possibilities: the first way uses counterpart theory to “pick up” objects and “move” them, by appealing

to counterpart relations between objects; the second way uses counterpart theory to “move the spacetime” under the objects by appealing to counterpart relations between regions of spacetime. It turns out that only the second of these strategies will work. Let’s consider why.

Begin with a case that involves a shift possibility wherein the material contents of the universe are rigidly translated one foot to the right of where they actually are. Now suppose we attempt to underwrite this possibility by appeal to counterpart relations that “pick up” the material contents of the universe and then “place” them one foot to the right. We can do this by appealing to the extremely permissive counterpart relation that makes every object a counterpart of every object. We may implement this counterpart relation so that, for example, it makes the lamp that is one foot to my right count as my counterpart. This will underwrite certain possibilities, among them will be the possibility that I be where the lamp actually is – that is, that I be one foot to the right of where I actually am. If we coordinate counterpart relations so that everything’s counterpart is the thing that is one foot to the right of where it actually is, we may be able to effect a global, one-foot-to-the-right shift of the material contents of the universe.

Essentially, the way this works is by thinking of this counterpart relation as affecting the occupancy relations born between material objects and regions of spacetime. For example, in the possibility that is actual, I bear the occupancy relation to a particular region of spacetime, call it R , and that lamp over there bears the occupancy relation to a particular region of spacetime, call it R^* . So, actually, I occupy R , and, actually, the lamp occupies R^* . However, in a one-foot-to-the-right shift possibility, I am where that lamp is. So within that possibility, I rather than the lamp bear the occupancy relation to R^* . This is so (as I said above) because there is a counterpart relation that makes that lamp my counterpart, and under that counterpart relation I

bear the occupancy relation to the region that the lamp actually bears the occupancy relation to.

But this approach is fraught. For one thing, it will only work for worlds where matter forms a plenum. Otherwise, there would be some object that is the farthest to the right, and that object would have no object one foot to its right that could be its counterpart, and thereby underwrite the possibility that it be one foot to the right. But there is another sort of problem, too.

What, exactly, is the region that the lamp occupies like? Intuitively, the notion of occupation that is relevant here is the one according to which the region the lamp occupies is a region that is exactly the right size for the lamp – none of the lamp spills out of the region, and the region doesn't have any more room for anything else. In short, the region the lamp occupies is (not surprisingly) exactly the shape of the lamp. But now consider the possibility wherein I bear the occupancy relation to this region. This would seem to be a possibility that involves me being lamp-shaped! Otherwise, how would I fit in there?¹⁹ But now, remember that the shift possibilities are supposed to be possibilities where the material contents of the universe remain qualitatively indiscernible.²⁰ But surely my actual shape and the possibility where I am lamp-shaped are qualitatively discernible!

This makes it look like the “pick the objects up” strategy won't work. So perhaps the way to cash out the shift possibilities is to “move” the spacetime, rather than the objects in it. Intuitively the way to do this is to “move” regions around: we move a region one foot to my right over to where I am, and thereby make it be the case

¹⁹I am, admittedly, intuition mongering a little here. One could maintain that objects and regions have their shapes independently, and that it follows from a commitment to a principle of recombination that objects don't always occupy regions of the same shape. I don't like this, though, and I'll say more about why in chapter 3.

²⁰Indeed, if we “move” me to where the lamp is by making the lamp my counterpart, then the possibility that I am where the lamp is is also a possibility where I am made of brass, have an electrical cord that plugs into a wall, etc. That's not what we want.

that I am one foot to the right of where I started.²¹ But it seems that there is a problem lurking here. If we want to say that I could be a foot to the right of where I actually am, and we are going to do this by “moving” (via a counterpart relation) the region one foot to my left over to where I am, we need to either ensure that the region we “move” over is the same shape as my actual region, or else we need to convince ourselves that it makes sense to have a me-shaped region be the counterpart of a non-me-shaped region.

But there is a good way to avoid this worry. It involves appealing to counterpart relations that hold between spacetime points, not between regions. We can appeal to a very general sort of counterpart relation that makes every spacetime point a counterpart of every spacetime point. Let’s walk through this.

First, take me and my region R . R is a collection of points that I bear the occupancy relation to. Now suppose we want to effect a shift that moves me one foot to the right of where I actually am. One way to do this is by starting with the points that compose R . Then, for each point p_i in R , we pick out a point q_i that is one foot to the right of p_i . Next, we appeal to our very general counterpart relation that makes every spacetime point a counterpart of every spacetime point, and we pick an arrangement licensed by this relation that makes each of the q_i the counterpart of the corresponding p_i . Then we observe that the fusion of the q_i is a new region R^* , where R^* is one foot to the right of R and the same shape and size as R .²² Next, we make me my own counterpart. Now when we consider the possibility underwritten by the counterpart relation that makes R^* the counterpart of R , and me my own counterpart, we have a possibility where everything is just as it actually is, except I

²¹Note that we’re not literally moving space around. Since motion is defined relative to space, moving space would seem to be incoherent. What we’re doing is figuratively moving by implementing counterpart relations that make one region represent another. I’ll continue to use this language since it’s convenient.

²²This follows from the fact that we are working in a flat, classical spacetime; things aren’t so tidy in the context of, for example, GR.

occupy R^* instead of R . Because R^* is a foot to the right of R , this is a possibility where I am one foot to the right of where I actually am. Repeating this procedure on a global scale will allow us to underwrite the possibility that the entire material contents of the universe be shifted one foot to the right. Then, making appropriate modifications – like changing the distance and direction of the points to be utilized as counterparts – will allow us to get all of the shift possibilities we want.

Let’s quickly summarize. Under the present account, shift possibilities are underwritten by the following counterpart mappings. First, we take points to be mapped to points that are some specified distance and direction away. Second, we take fusions of points to be mapped to the fusions of points that those points get mapped to – so the fusion of points p and q gets mapped to the fusion of p ’s counterpart and q ’s counterpart. This just makes the regions tag along with the points – that is, we get regions to “move” by “moving” the points that compose them. Last, we take every material object to be mapped to itself.

So, it looks like an antihaeccetist has in counterpart theory a nice way of responding to Leibniz’ static shift argument. Ontological profligacy is avoided by denying that there are worlds that differ merely in the ways that shifts differ, but the intuition that static shifts are genuinely possible is satisfied by the holding of certain counterpart relations among the inhabitants of a single world.²³

²³There is a general worry once we have all of this in place. Given counterpart theory a single world can represent all of the shift possibilities, and that means we are not, like the haeccetist, committed to qualitatively indiscernible worlds – we get to keep the PII as it applies to worlds. But now we need to ask: are the shift possibilities themselves qualitatively identical? Suppose they are. Then it seems that the PII should apply to them just as it applies to worlds. But if the PII applies to possibilities, then the distinction between shift possibilities collapses. But if the distinction between shift possibilities collapses, then the antihaeccetist has to claim that they are impossible, and this seems to force a strange and counterintuitive form of essentialism about occupancy relations on the antihaeccetist.

So suppose the shift possibilities are not qualitatively identical. (This may seem like a strange proposal, but let’s follow it to see where it leads). What sort of explanation could one offer here? Here is one. According to Lewis (1986), possibilities (not possible worlds!) are a certain sort of set-theoretic construction. (See Lewis (1986), pgs. 232-234, and Hazen (1979).) Take, for example, me and the lamp one foot to my left. We are a pair – the pair $\langle \text{me, lamp} \rangle$ – of compossible individuals.

1.4 Conclusion

Stepping back a bit, we can see that the problem raised for a substantialist by Leibniz’ static shift argument is both deep and simple: more ontological categories means a wider range of possibilities, and more possibilities can mean more opportunity for trouble. More specifically, buying-in to an “extra” ontological category of spacetime makes substantialism compatible with a range of possibilities that differ only non-qualitatively. For Leibniz, such possibilities were troublesome because they lead to violations of the PII and the PSR. From the modern perspective, the problem is that such possibilities constitute an explosion of mere haecceitistic differences, and this makes the question of haecceitism a pressing one for a substantialist. As we’ve seen, though, counterpart theory provides a substantialist who denies haecceitism a nice way of responding.

Lewis would say that switching the order of the pair is one way that this pair of individuals could differ. So there are at least two possibilities (in fact, there will be many more than two) for me and the lamp; the one that is actual – $\langle \text{me}, \text{lamp} \rangle$ – and the possibility wherein we switch places – $\langle \text{lamp}, \text{me} \rangle$. So perhaps one avenue of response for the antihaecceitist to the present dilemma is to say that the PII can’t apply to possibilities – and thereby collapse them – because the PII only applies to things that have qualitative properties and not to set-theoretic representations of things that have qualitative properties. But then this seems suspiciously ad hoc. Thanks to Chris Meacham for discussion about this problem.

CHAPTER 2

CLIMBING OUT OF THE HOLE

2.1 Introduction

Over the last quarter of a century, nearly every philosophical engagement with the dispute between substantivalists and relationalists has involved a discussion of the so-called *hole argument*.¹ The hole argument has its roots in Einstein's work on the field equations for general relativity (GR), but was introduced into the philosophical mainstream by Stachel (1980), and then used to argue against substantivalism by Earman and Norton (1987). The argument is supposed to impugn substantivalism in the context of GR by showing that it leads to an objectionable form of indeterminism.

There is now a consensus regarding the moral of the hole argument. The moral is that the argument does not succeed in its ambition to undermine substantivalism in the context of GR, but it does exert pressure on substantivalists who also accept *haecceitism* – roughly, the view that there are qualitatively indiscernible worlds that differ non-qualitatively.² There are two general strategies for addressing this pressure.³ The first strategy is to embrace *antihacceitism* – roughly, the view that there are no qualitatively indiscernible worlds.⁴ The second strategy notes that the notion of determinism exploited by the hole argument is sensitive to haecceitistic differences

¹ Cf. Pooley (2013, §7) and Norton (2011).

² More rigorous characterizations of haecceitism and antihacceitism appear in §3.4.

³ Cf. Dasgupta (2011, p. 115).

⁴ E.g., Butterfield (1988, 1989), Brighouse (1994), Mundy (1992), Hofer (1996), Pooley (2006, 2012), and Rynasiewicz (1994). For in-depth discussion of the variety of antihacceitist responses and the differences among them, see Rickles (2008), Pooley (2002), and Dasgupta (2011).

between worlds, and replaces this notion of determinism with one that is *insensitive* to haecceitistic differences between worlds.⁵

In this chapter I will argue that proponents of both strategies have conceded too much. It is not the case, as advocates of the first strategy would have it, that the hole argument warrants a rejection of haecceitism. And it is not the case, as advocates of the second strategy would have it, that the hole argument warrants a retooling of the thesis of determinism. Rather, the hole argument has *nothing at all* to do with haecceitism. Haecceitism has mistakenly been given a starring role in this drama because proponents of both strategies have failed to make explicit the range of background assumptions needed to level the hole argument. In particular, they have failed to pay sufficient attention to the relationship between the formulation of a physical theory, the theory's models, and the worlds at which the theory is true.

My discussion will be organized as follows. I begin in §2.2 by outlining the conceptual framework I will use throughout the chapter. Here the emphasis will be on carefully distinguishing features of mathematical models that are used to represent worlds, and the physical correlates of those representations. Conflating the representation and the thing represented in this context is both common and dangerous.⁶ In §2.3 I consider two conceptions of the relationship between GR's models and GR's worlds. In §2.4 I show how the default conception of this relationship plays a key role in generating the hole argument, and in §2.5 I show how the alternative conception allows us to dodge the hole argument.

⁵E.g., Arntzenius (2012), Melia (1999) and Skow (2005).

⁶Maudlin (1988) makes this observation, too.

2.2 Preliminaries

It will be important for our purposes to be explicit about the relationship between the mathematical apparatus of GR and the physical reality that that apparatus is used to represent. In this section I draw some distinctions that will help to make this relationship clear.

2.2.1 Worlds, Models, and Points

Begin by distinguishing worlds from models. I understand ‘world’ in roughly the way that Lewis (1986) does: there are many worlds, all alike in kind; one of them is ours; the rest are, from our perspective, merely possible. Many of these worlds are composed of a physical spacetime. The basic unit of physical spacetime is the physical point. I assume that classical mereology applies to these basic units, and I take the physical spacetime of a world w to be the fusion of spatiotemporally interrelated physical points at w . I will say that a *region* of physical spacetime is any fusion of physical points, and that a region A is a *subregion* of a region B iff A is a part of B .

Models are abstract mathematical objects. The sorts of models that will concern us here are general relativistic models. In general relativity, the basic ingredients of reality are described by appeal to a metric tensor field g and a stress-energy tensor field T , both of which are defined over a special sort of set of mathematical points called a *differentiable manifold*, \mathcal{M} .⁷

When a particular smattering of values of the metric and stress-energy tensor fields across a differentiable manifold satisfies Einstein’s Field Equations, that smattering of field values across that differentiable manifold is said to be a *model* of GR. Typically, models of GR are represented as the ordered triples $\langle \mathcal{M}, g, T \rangle$. For my purposes it will be useful to be able to unpack the manifold \mathcal{M} of GR models into the ordered pair (X, \mathcal{S}) , where X is a base set of points, and \mathcal{S} is the structure – i.e., topological,

⁷See Appendix A for details on manifolds.

affine, and differentiable – defined over X . So I will refer to manifolds as ‘ (X, \mathcal{S}) ,’ and I will refer to GR models as ‘ $\langle (X, \mathcal{S}), g, T \rangle$.’

I have appealed to a distinction between mathematical and physical points: physical points are what worlds are made of, mathematical points are what models are made of; physical points are (in some intuitive sense) concrete, mathematical points are (in some intuitive sense) abstract. An important distinction that follows on the heels of this one regards special collections of mathematical and physical points. I have in mind here a distinction between mathematical *cauchy surfaces* of GR models, and physical *maximal achronal subregions* of worlds. Both of these notions correspond in some intuitive sense to the notion of an instant of time, or an instantaneous space. To be exact, a cauchy surface of an n dimensional differentiable manifold (X, \mathcal{S}) is an $n - 1$ dimensional surface that intersects every extensible, time-like curve in (X, \mathcal{S}) at exactly one point. A cauchy surface of a model $\langle (X, \mathcal{S}), g, T \rangle$ is just a cauchy surface of (X, \mathcal{S}) along with the values of g and T defined at the points that compose it. I will define ‘maximal achronal subregion of a world’ as follows. First, I’ll say that a subregion S of a world w is *achronal* iff all points in S are spacelike separated from one another. Then I’ll say that an achronal subregion S of a worlds w is *maximal* iff every point in w that is not in S is either light-like or time-like related to points in S .⁸

2.2.2 Properties and Relations

I will accept an account of properties and relations that is very nearly that of Lewis (1986). So let us begin by recalling Lewis’ views concerning the taxonomy of properties and relations. He begins by recognizing a distinction between the *natural* (or *sparse*) and the *abundant* properties. The abundant properties, Lewis says, are “as

⁸So-called “island universes” will cause a wrinkle for this characterization of ‘maximal achronal subregion’, but we can ignore this complication. See Bricker (2011) for more on island universes.

gruesomely gerrymandered, as miscellaneously disjunctive, as you please. They pay no heed to the qualitative joints, but carve things up every which way. Sharing of them has nothing to do with similarity. (1986, 59)” The natural properties, Lewis continues “are another story. Sharing of them 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 characterize things completely and without redundancy. (1986, 60).” The natural properties, it is clear to see, are the ones that will be of interest to physics.

Next, Lewis claims that there are an elite few among the natural properties that are *perfectly natural*. Perfect naturalness is a primitive on Lewis’ view, which he uses to define several important concepts. First, he defines duplication in terms of perfect naturalness as follows: “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. (1986, p. 61). Next, Lewis defines *intrinsic* in terms of duplication: “an *intrinsic* property is one that can never differ between two duplicates. (1986, p. 62)” Lewis (1986, p. 61) thinks that there are intrinsic properties that are not perfectly natural, but that all perfectly natural properties are intrinsic.

Next, distinguish *individualistic* properties from *non-individualistic* properties. I will say that an individualistic property is a property that concerns a particular individual, whereas a non-individualistic property is one that does not. The property *being Obama*, for example, is an individualistic property since it concerns a particular individual, whereas *being president* is a non-individualistic property since it doesn’t concern any particular individual. Likewise, the property *being point p* is an individualistic property since it concerns a particular point, whereas *being a point* is a non-individualistic property since it doesn’t concern any particular point.

There are many sorts of non-individualistic properties. Qualitative properties are one sort. Given the distinction between the natural properties and relations and all the others, we can characterize qualitative properties as those that are built from the natural properties and relations. Examples of qualitative properties include properties like *being green*, *being round*, *being negatively charged*, and so on. Another sort of non-individualistic property is the sort associated with purely geometric entities like tensors. A tensor is a multilinear map from elements of vector spaces to the real numbers. As such, a tensor cannot be a qualitative property or relation since it is not built out of natural properties or relations, but it is nevertheless non-individualistic since it does not concern any particular individual.

Although tensors, because they are mathematical objects defined over mathematical points, are not themselves qualitative properties or relations, they are often used to *represent* certain qualitative properties and relations. In GR, for example, the metric tensor g and the stress-energy tensor T are used to represent distances and the distribution of mass and energy, respectively. For my purposes it will be important not to confuse tensors with the things they represent. But let's consider for a moment what, exactly, they do represent.

2.2.3 Tensors and What They Represent

To begin, note that for every mathematical point p of the manifold (X, \mathcal{S}) , there is a value of g and a value of T defined at p . So there is a sense in which values of g and T defined over mathematical points of (X, \mathcal{S}) look like they ought to be taken to represent monadic, intrinsic properties of points of physical spacetime (or of point sized bits of matter that occupy points of spacetime). But recall that tensors, like g and T , are maps from elements of a vector space associated with points of (X, \mathcal{S}) to real numbers. So there is another sense in which values of g and T defined over mathematical points of (X, \mathcal{S}) look like they ought to be taken to represent

relations: they each relate a pair of vectors to a real number. But things get a bit more complicated than that. Associated with each point p of a manifold (X, \mathcal{S}) is a vector space called the ‘tangent space of p ’. The tangent spaces associated with the points of a (X, \mathcal{S}) can be conjoined in such a way that they form their own manifold – called the ‘tangent bundle of (X, \mathcal{S}) ’. The dimension of the tangent bundle associated with an n -dimensional manifold is $2n$. Thus when we think of g and T as relations that relate vectors in the tangent bundle of a manifold, we are thinking of them as relating things that are, in a sense, not even part of the manifold we hoped to be describing, but are, rather, parts of a higher dimensional manifold that is used to represent features of the manifold we set out to describe.

In any event, because the tensors g and T each have values that are defined at a single mathematical point of (X, \mathcal{S}) , they seem in a way like they ought to represent monadic properties of physical points. But at the same time they also definitely involve things other than the points they are defined at. A particular value v of the metric field at a point p , for example, is not only about p , it’s also about other things, like some vectors in the tangent space associated with p . The most straightforward thing to do in light of this is to treat g and T as *extrinsic* properties of points.

But this now creates a cascade of problems. First, according to Lewis, all perfectly natural properties are intrinsic. So, if the sorts of properties that g and T represent are extrinsic, then they cannot be perfectly natural. But surely we *ought* to treat g and T as representing properties that are perfectly natural: if we don’t treat them as perfectly natural, then, because there aren’t any more plausible candidates, there will be *no* perfectly natural properties in a GR world. That is certainly unacceptable.

There is another, related problem with treating g and T as representing extrinsic properties of points. This one has to do with Lewis’ conception of modal plenitude. To begin, note that in order to underwrite his modal realism, Lewis requires a principle of plenitude; a principle that allows us to reason from some possibility to others. The

underlying intuition Lewis appeals to on behalf of plenitude is Hume's: no necessary connections between distinct existences. This gives Lewis what he calls the *principle of recombination*, according to which any distinct things can coexist or fail to coexist with one another.⁹ Now, if the metric and stress-energy tensors that appear in GR models are taken to represent properties that are extrinsic, then we will be unable to apply the principle of recombination to GR worlds. Here is why. Suppose some GR model represents a particular point p as instantiating a physical correlate of a particular value of the stress-energy tensor field. Call the physical correlate of this value v . So this model represents p as having v . Now, any duplicate of p will have v too, since v is a perfectly natural property. But v is *extrinsic*, which means that it implies the existence of something other than just p . In this case it will be some pair of vectors in a tangent space. For simplicity, just call the pair q . So now any world that has a duplicate of p will also have q . This means that while p and q are distinct things, p can never get away from q : they are necessarily conjoined. This is a sin against the Humean intuition.

So, we have several options open to us now. First, we could concede that the principle of recombination does not apply to GR worlds. But this would be a serious blow to the Lewisian framework: why does the principle apply only to *some* worlds, and not to all of them? Finding a principled explanation seems unlikely. Alternatively, we could reject the tensor analytic formulation of GR and hope that one of the other formulations is more amenable to the principle of recombination. This is also a bad option, though. The tensor analytic formulation of GR is the mainstream formulation of our best theory of the large scale structure of the universe. If the cost of adopting the Lewisian framework were giving it up, one might reasonably regard the framework with skepticism.

⁹Note that coexistence will be spelled out in terms of duplicates. I, for example, possibly coexist with a dragon in virtue of having some duplicate that inhabits a world with dragons.

Luckily, there is a much better alternative. Bricker (1993) points out that while tensors represent properties that are extrinsic to individual points, the same properties are also intrinsic to *infinitesimal neighborhoods* of points. This insight does several related and important things. First, it allows us to treat tensors as representing properties that are perfectly natural. Why? Because we seem to have two plausible interpretations of the sorts of properties that tensors represent: extrinsic properties of individual points, or intrinsic properties of infinitesimal neighborhoods of points. Since the Lewisian may adopt the interpretation that is amenable to her framework, she is free to say that tensors represent properties that are intrinsic to infinitesimal neighborhoods of points. Second, the insight allows us to agree with Lewis that all perfectly natural properties are intrinsic. Third, it makes GR compatible with recombination. How? By treating infinitesimal neighborhoods of points, rather than individual points, as the appropriate units of recombination.¹⁰

In summary, so long as we accept Bricker’s revision, we may continue to buy into the Lewisian framework. So we’ll be taking tensors to represent intrinsic properties of infinitesimal neighborhoods of points. This will allow us to treat tensors as representing perfectly natural properties.

2.2.4 Individualistic and Non-Individualistic Differences

Next, when models \mathfrak{M} and \mathfrak{M}' are exactly alike with respect to non-individualistic properties, but differ in which individualistic properties they possess, or in how the individualistic and non-individualistic properties are co-instantiated, I will say that \mathfrak{M} and \mathfrak{M}' differ *merely individualistically*. When worlds w and w' are exactly alike with respect to non-individualistic properties, but differ in which individualistic properties they possess, or in how the individualistic and non-individualistic properties are co-

¹⁰Note that this account does require the acceptance of non-standard analysis. See Bricker (1993) for more.

instantiated, then I will say that w and w' differ *merely haecceitistically*. Finally, I will take *haecceitism* to be the thesis that some worlds differ merely haecceitistically, and *antihaecceitism* to be the thesis that no worlds differ merely haecceitistically.

Note that Lewis (1986) and I give different and incompatible characterizations of the distinction between ‘haecceitism’ and ‘antihaecceitism’. In fact, on my usage of the term ‘haecceitism’, Lewis counts as a haecceitist. Let me explain why. First, Lewis is “agnostic about whether there are indiscernible worlds (1986, p. 87)”, and it is compatible with this agnosticism that there be a pair of worlds w and v that are qualitative duplicates. Next, Lewis denies transworld identities, so inhabitants of w and v will be distinct. It follows from the distinctness of individuals in w and v and from Lewis’ views about property individuation¹¹ that w and v differ individualistically. Here is why. Corresponding to each individual in w is the unit set of that individual, and, according to Lewis, corresponding to each set is the property of being that particular set. So for each individual in w , there is the property of being that particular individual. Likewise for v . But this now means that w and v differ with respect to individualistic properties. Further, because w and v are qualitative duplicates, they differ *only* with respect to individualistic properties – that is, they differ merely haecceitistically. Since I have characterized haecceitism as the view that there are worlds that differ merely haecceitistically, I must count Lewis as a haecceitist.¹²

2.2.5 Notation

I have been careful to draw some distinctions that we’ll need throughout the remainder of the discussion. Let us now introduce some terminology to help make

¹¹See especially (1986, p. 225).

¹²I am not the only one that thinks an intuitively compelling formulation of haecceitism counts Lewis as a haecceitist. Skow (2008) makes this observation, too.

these distinctions easy to keep track of. First, I use ‘ g ’ and ‘ T ’ to refer to the metric and stress-energy tensor fields defined over mathematical points of manifolds, and I use ‘ \mathbf{g} ’ and ‘ \mathbf{T} ’ to refer to the properties of worlds that g and T are used to represent. I use ‘ \mathcal{S} ’ to refer to various structures defined over the base set of mathematical points of GR models, and I use ‘ \mathbf{S} ’ to refer to the physical correlate of \mathcal{S} .¹³

2.2.6 Diffeomorphism

A diffeomorphism is a kind of mapping between differentiable manifolds that can be used to drag various geometric entities defined over points of one over to points of the other. GR is *diffeomorphism invariant*. GR’s being diffeomorphism invariant means that if a particular smattering of field values across a manifold satisfies Einstein’s Field Equations (EFE) for GR, then any smattering that is related to that one by a diffeomorphism will satisfy the EFE, too. In other words, if you begin with a model of GR, $\langle (X, \mathcal{S}), g, T \rangle$, and you apply a diffeomorphism d to it, you are guaranteed to get another model of GR, $\langle d(X, \mathcal{S}), d(g), d(T) \rangle$ as a result. This allows us to think of diffeomorphism as an equivalence relation that partitions the class of GR models. We can call these equivalence classes under diffeomorphism *D-classes* for short.

Diffeomorphisms can sometimes also be *automorphisms*: maps that take a manifold back to itself. When a diffeomorphism is also an automorphism, and we use it to drag geometric objects around, the effect is a redistribution of geometric objects among the very same base set of points. It is generally assumed – and I will go along with this assumption – that the sorts of diffeomorphisms that are relevant to

¹³It is an interesting and important question just which structures of \mathcal{S} have physical correlates. These are questions I will set aside since whatever answers one prefers can be plugged into the framework I offer here. See, for example, Arntzenius and Dorr (2012, ch. 8) for more.

D-class formation are also automorphisms.¹⁴ From now on when I use the term ‘diffeomorphism’ I’ll mean ‘*automorphic* diffeomorphism’. Now observe that a non-trivial diffeomorphism on a model \mathfrak{M} works by “picking up” the geometric entities defined over points of a manifold (X, \mathcal{S}) and moving them to *different* points of (X, \mathcal{S}) . We are then to think of the first arrangement of geometric entities over points of (X, \mathcal{S}) as constituting \mathfrak{M} , and the second, moved around arrangement of geometric entities over points of (X, \mathcal{S}) as constituting $d(\mathfrak{M})$. The individualistic properties of models are properties like *being point p*, *being point q*, etc., and the non-individualistic properties are values of the metric and stress-energy tensor fields. Both of these sorts of properties are operated on by diffeomorphisms. So, intuitively, a diffeomorphism takes a model with a particular pattern of co-instantiation of individualistic and non-individualistic properties and makes a new model by shuffling that pattern of co-instantiation around.

Now note that this conception of the effect of a diffeomorphism on a model requires certain assumptions. First, we must assume that the individualistic properties of points are not determined by their non-individualistic properties. For if individualistic properties were determined by non-individualistic properties,¹⁵ then diffeomorphic models would fail to differ since they involve the very same arrangement of non-individualistic properties. Second, we must recognize *trans-model identities* between points. That is, we must assume that there are facts about whether the same individualistic properties are instantiated in different models. Without this assumption, there would be no basis upon which to claim that we can generate a new model simply by moving the geometric entities around over a set of points. The upshot here is that models belonging to a common D-class, in virtue of being related by

¹⁴*Cf.* Butterfield (1988, 1989), Carroll (2004, especially Appendix B), Earman and Norton (1987, 520), Hawking and Ellis (1973, 56), and Norton (1993, 824).

¹⁵This is to treat the geometric entity fields as “individuating fields.” *Cf.* Dasgupta (2011) and Brighouse (1994).

automorphic diffeomorphisms, differ merely individualistically. In particular, models belonging to a common D-class will be exactly alike with respect to which individualistic and non-individualistic properties there are; they will differ only with respect to the pattern in which these properties are co-instantiated. This is an especially important point because (as we will see in more detail in §4) one of the premises of the hole argument involves the claim that substantivalists must regard mere individualistic differences among models of a common D-class as being correlated with mere haecceitistic differences among classes of worlds.

2.3 What do GR Models do?

Modern physical theories centrally involve collections of dynamical equations that somehow or other manage to convey something about worlds (and, most relevantly, about our own world). GR, for example, is centered around Einstein’s Field Equations, and these equations are supposed to say something about the relationship between mass and energy on the one hand, and spacetime curvature on the other. But how exactly does GR “say” anything about worlds?

One approach to this question is to ask how we are supposed to interpret the sentences that express the EFE. There seem to be two different ways of doing this. One way is to interpret the sentences in terms of GR’s $\langle(X, \mathcal{S}), g, T\rangle$ models,¹⁶ and the other way is to interpret them in terms of worlds. Think of the EFE as unfolding into infinitely long conjunctions of statements about the field values of individual points. If we pursue the interpretation in terms of models, then these statements will involve attributions of values of g and T to mathematical points of a manifold, whereas if

¹⁶Actually, while the standard formulation of GR is in terms of tensors on manifolds, it is also possible to formulate it in terms of Einstein algebras, or twistors. These alternative formulations will yield different sorts of models. So which mathematical objects are the models of GR needs to be relativized to a formulation. See Bain (2003, 2006) about this general point and about the twistor formalism, and Geroch (1972) and Earman (1989) about Einstein algebras.

we pursue the interpretation in terms of worlds, they will involve attributions of the properties \mathbf{g} and \mathbf{T} to physical points of worlds.

Now, it might initially seem that a physical theory can *say* something about worlds only if we *interpret* it in terms of worlds. It might seem that if we interpret the theory in terms of models, then it's just pure mathematics – it doesn't say anything about worlds, and doesn't really deserve to be called a *physical* theory. But that would be too hasty. Perhaps the dynamical equations of a theory like GR pick out a collection of models, and these models then bear certain relations to worlds, and it is in virtue of the theory's models being related to worlds that the theory *says* something about worlds.

Let's regiment this discussion. First, let's call the set of worlds at which a physical theory is true the *content* of that theory. Now let us distinguish the *heuristic conception* of models, from the *robust conception* of models. According to the heuristic conception of models, the content of GR is determined by interpreting the dynamical equations of GR in terms of worlds. On this view, one may still help themselves to an interpretation of these equations in terms of models for heuristic purposes – perhaps to help us visualize what sorts of transformations are permitted by the theory. But an interpretation in terms of models is not how the theory determines its content.

According to the robust conception, on the other hand, there is only one right way to interpret GR. This interpretation is the one in terms of models. Models then bear certain relations to worlds, and it is in virtue of being so related to GR's models that these worlds come to form the content of GR. On this view, models are like bridges that connect the dynamical equations of theory to the worlds at which it is true.

I will not offer a full defense of the robust conception here. What I will do instead is show what sorts of consequences would follow from accepting it in the context of the hole argument. To do this, I'll begin by examining the hole argument and some of

the more popular responses to it. Then I'll develop the robust conception in enough detail to unfold its consequences for the hole argument.

2.4 The Hole Argument

Recall that GR is diffeomorphism invariant. This means that any diffeomorphism on a model of GR will yield a new model of GR. So consider a model \mathfrak{M} and a “hole” diffeomorphism h that is an identity mapping everywhere except for an arbitrarily small region occupying a cauchy surface of \mathfrak{M} . Call this region H . Within H , h effects a shuffling of individualistic and non-individualistic properties. Now, given \mathfrak{M} and the diffeomorphism h , we are guaranteed a new model $h(\mathfrak{M})$ that differs from \mathfrak{M} only within H , and here it differs from \mathfrak{M} merely individualistically – that is, both models have the same individualistic and non-individualistic properties, but differ in how these properties are coinstantiated. Given this, \mathfrak{M} and $h(\mathfrak{M})$ will violate the following notion of determinism.

Model Determinism (MD): GR satisfies MD *iff*: for all GR models, \mathfrak{M} and \mathfrak{M}' , and for all cauchy surfaces Σ in \mathfrak{M} and Σ' in \mathfrak{M}' : if Σ and Σ' involve the same pattern of co-instantiation of individualistic and non-individualistic properties then, \mathfrak{M} and \mathfrak{M}' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

Now, it might initially strike one that violating something like MD is not obviously damning. After all, MD is about GR's models, not worlds, and presumably when it comes to violations of determinism, it's worlds we care about, not models. This is where the heuristic conception leads us into trouble.

On the heuristic conception, while models do not perform any vital role in securing a theory's content, they are nevertheless important heuristic tools. In the context of the hole argument, for example, models are appealed to as aids in visualizing the

spread of worlds compatible with GR. But the problem here is that we make this appeal to models without being sufficiently explicit about which features of worlds they represent. So it's as though models are asked to perform an important role without being told exactly what that role is. Let's consider this in more detail.¹⁷

Begin with a D-class \mathcal{D} of models. Once we attempt to use \mathcal{D} to help us visualize a collection of worlds, we immediately face a difficult question: what features of \mathcal{D} really represent physical quantities and entities? On the heuristic conception, one confronts this question by employing the following strategy: take all (or, anyways, most) features of \mathcal{D} to have physical correlates. Then look and see what sorts of consequences follow from taking various of these features to have physical correlates. If undesirable consequences follow from accepting a correlation between some mathematical feature of members of \mathcal{D} and some physical feature of some world, then consider rejecting that correlation. Then let the philosophical haggling commence over questions about which correlations are mandatory and which negotiable given a commitment to substantivalism.

So within the context of the heuristic conception, the default strategy is to use \mathcal{D} as a means of visualizing worlds in a way that maximizes respects of similarity between \mathcal{D} and some collection of worlds. This, for example, involves taking each element \mathfrak{M} of \mathcal{D} to have some world w as a physical correlate, where w has a particular physical structure \mathbf{S} *iff* \mathfrak{M} has a corresponding mathematical structure \mathcal{S} , and where physical points of w have properties \mathbf{g} and \mathbf{T} *iff* corresponding points of \mathfrak{M} have corresponding values of g and T , and where w has a particular pattern of coinstantiation of individualistic and non-individualistic properties *iff* \mathfrak{M} has a corresponding

¹⁷The most explicit general discussion of the role of models in the heuristic conception can be found in Skow (2005, §2.1) and Liu (1996), but see also Pooley (2002, 2011) and Dasgupta (2011, 126). Also note that Butterfield (1989) has an especially clear discussion of the way advocates of the heuristic conception reason from the formulation of determinism in terms of models to the physically important formulation in terms of worlds.

pattern of coinstantiation of individualistic and non-individualistic properties. This last type of correspondence – the correspondence between patterns of coinstantiation of individualistic and non-individualistic properties – is the key point of entry for the issue of haecceitism. It is this type of correspondence that causes mere individualistic differences among members of a common \mathcal{D} -class to be reproduced as mere haecceitistic differences among worlds. And, importantly, it is the heuristic conception that seems to legitimize this reproduction. It does so by tacitly appealing to a strategy of maximizing respects of similarity between models and worlds.

Given this background assumption about the way to use models to visualize worlds, a particular general principle of correspondence between models and worlds looks quite plausible.

Haec: for all elements \mathfrak{M} of $\text{Mod}(\text{GR})$ and for all worlds w , if \mathfrak{M} represents w , then

- (1) **Every:** every element \mathfrak{M}' belonging to the same \mathcal{D} -class as \mathfrak{M} represents some world, and
- (2) **Distinct:** there is no element \mathfrak{M}' of $\text{Mod}(\text{GR})$ such that $\mathfrak{M}' \neq \mathfrak{M}$ and \mathfrak{M}' represents w .¹⁸

The essence of **Haec** is quite simple: there is a one-to-one correspondence between models and worlds; where there is a distinction recognized by models, there is a corresponding distinction recognized by worlds. So if models can differ merely individualistically, then, correspondingly, worlds can differ merely haecceitistically. Now the problem is that, given **Haec**, it follows from there being GR models \mathfrak{M} and $h(\mathfrak{M})$ that there are GR worlds w and $h(w)$ that differ in a way that is exactly analogous

¹⁸Pooley (2002, p. 93) and Rickles' (2008, p. 90) use the term **Haec** for an equivalent principle. Note that **Haec** is also roughly equivalent to Earman and Norton's (1987) "denial of Leibniz equivalence", and Rynasciewicz' (1994, p. 409) "model literalism."

to the way that \mathfrak{M} and $h(\mathfrak{M})$ differ. That is, where \mathfrak{M} and $h(\mathfrak{M})$ differ merely individualistically, w and $h(w)$ differ merely haecceitistically. But just as \mathfrak{M} and $h(\mathfrak{M})$ jointly violate MD, w and $h(w)$ will violate an analogous notion of determinism in terms of worlds: WD.

World Determinism (WD): GR satisfies WD *iff*: for all GR worlds, w and w' , and for all maximal achronal subregions S in w and S' in w' : if S and S' involve the same pattern of co-instantiation of individualistic and non-individualistic properties then, w and w' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

Violations of WD cannot be dismissed as a mathematical quirk in the way that violations of MD can, since WD is quite clearly a *physical* variety of indeterminism. This now gives us everything we need to present an organized rendition of the hole argument against substantivalism.

2.4.1 The Hole Argument Against Substantivalism

The Hole Argument Against Substantivalism

- (P1) If substantivalism is true, then **Haec** is true.
- (P2) If **Haec** is true, then GR violates WD.
- (C1) So, if substantivalism is true, then GR violates WD.
- (P3) Any metaphysical thesis from which it follows that GR violates WD is false.
- (C2) So, substantivalism is false.

The most popular response to this argument is the *antihaecceitist* response. The antihaecceitist denies (P1). Denying (P1) involves showing that substantivalism is compatible with the denial of **Haec**. There are two ways to deny **Haec**: one way involves rejecting the first conjunct of **Haec** – **Distinct** – and the other way involves

rejecting the second conjunct – **Every**. Butterfield (1988, 1989) is an example of someone who denies **Every**. He does this by rejecting the view that trans-world identities are genuine identities, and accepting instead Lewis’ counterpart theory as an account of *de re* modality. According to this view, individuals, like physical points, do not exist in more than one world. Counterfactual claims about individuals, then, are made true by otherworldly counterparts of those individuals, rather than by those individuals themselves being located in distinct worlds.

Accepting this sort of metaphysical package allows Butterfield to claim that if we take a model \mathfrak{M} to represent some world w , then we have exhausted the representational resources of \mathfrak{M} ’s entire D-class. Consider a model \mathfrak{M} that involves a mathematical point m . Note that any model $d(\mathfrak{M})$ in the same D-class as \mathfrak{M} will involve m , too. Now suppose we take \mathfrak{M} to represent some world w . This will involve using m to represent some physical point p in w . But if we try to use $d(\mathfrak{M})$ to represent a distinct world $d(w)$, then, because $d(\mathfrak{M})$ involves m and m has been assumed to represent p , we seem to be committed to recognizing p as a part of $d(w)$. But this move trades in exactly the sort of “genuine” trans-world identities that counterpart theorists deny are legitimate: if p exists in w , then it cannot literally also exist in $d(w)$. In effect, Butterfield’s strategy is to deny that there can be worlds that differ in the sorts of ways (that is, haecceitistically) that lead to violations of WD.

Hofer’s (1996) response to the hole argument is similar in effect to Butterfield’s, but involves denying **Distinct**, rather than **Every**. Hofer argues that we ought to take the individualistic properties of physical points to be fixed by the non-individualistic properties of physical points. If we do this, then there is no way to “re-shuffle” the pattern in which individualistic and non-individualistic properties of worlds are coinstantiated – the individualistic properties follow the non-individualistic ones. Thus there cannot be worlds w and $d(w)$ that differ in the sort of way that leads to a violation of WD.

2.4.2 The Argument Against Haecceitistic Substantivalism

Though there are other varieties of antihaecceitism, what all varieties have in common is the supplementation of substantivalism with a metaphysical package that prevents logical space from being populated by worlds that differ merely haecceitistically. That is, while the antihaecceitist agrees that haecceitism is relevant, he urges that a substantivalist may nevertheless inoculate himself against the hole argument with a modal vaccine: antihaecceitism. This move redirects the force of the hole argument against the substantivalist who is also a haecceitist. This furnishes us with the following version of the argument.

The Argument Against Haecceitistic Substantivalism

- (P1*) If haecceitistic substantivalism is true, then **Haec** is true.
- (P2*) If **Haec** is true, then GR violates WD.
- (C1*) So, if haecceitistic substantivalism is true, then GR violates WD.
- (P3*) Any metaphysical thesis from which it follows that GR violates WD is false.
- (C2*) So, haecceitistic substantivalism is false.

There are several ways for a haecceitistic substantivalist to respond to this argument. Maudlin (1988), for example, denies (P1*). He claims that a haecceitistic substantivalist ought to supplement haecceitistic substantivalism with the doctrine of essentialism as applied to metrical relations. Consider a world w according to which physical points p and q bear the metric relation R . Because the metric essentialist takes all metrical relations to be essential to the points that have them, if p bears R to q in w , then every possible world in which p and q exist is a possible world where p bears R to q . Given this commitment to essentialism, the metric essentialist can take at most one member of a D-class to represent a possible world: if one model

represents a possibility, then the others, because they differ in the metrical relations they represent physical points as having, cannot.

The most popular way for a haecceitistic substantialist to respond to the argument is to deny (P3*) and argue that the notion of determinism exploited by the hole argument is not the sort of determinism that we ought to care about. Arntzenius (2012), Melia (1999), and Skow (2005) are examples of this sort of strategy. They argue that the hole argument shows us that the sort of determinism relevant to GR is a sort that is blind to haecceitistic differences between worlds.

2.4.3 The Argument Against Haecceitistic Determinism

This furnishes us with yet another way of redirecting the hole argument.

The Argument Against Haecceitistic Determinism

- (P1**) Haecceitistic substantialism is true.
- (P2**) If Haecceitistic substantialism is true, then **Haec** is true.
- (P3**) If **Haec** is true, then GR violates WD.
- (P4**) If GR violates WD, then WD is false.
- (C**) So, WD is false.

It is incumbent on the haecceitistic substantialist who understands the force of the hole argument in this way to offer an alternative formulation of determinism.¹⁹ For our purposes, we need not go any further. For I will argue that all of the lines of argument we have considered go wrong by emphasizing the relevance of **Haec**. The standard dialectic revolves around the question of who is committed to **Haec** and how, if one is, they are to defuse that commitment. But I'll argue **Haec** really isn't relevant. It merely *looks* relevant under the assumption of the heuristic conception. But what if we reject the heuristic conception?

¹⁹See Arntzenius (2012), Melia (1999), and Skow (2005).

2.5 The Robust Conception

Recall that advocates of the heuristic conception confront the question of which features of models really represent physical quantities and entities by appeal to a strategy of maximization: first, you maximize the respects of similarity between a collection of models and a collection of worlds; next, you see if any of these respects of similarity cause trouble. If there is trouble, then consider rejecting the troublemaking correlation. This is a rather oblique way of addressing the question, and it has the effect of making **Haec** look relevant, when it really isn't.

Advocates of the robust conception confront the question of which features of models represent physical quantities and entities head-on. In the following sections I sketch one way that friends of the robust conception may proceed. This involves using a relation that I call *determination* as a means by which to construct GR's content. After characterizing this relation, I consider its effect on the hole argument, and show why **Haec** is irrelevant.

2.5.1 Determination

I'll now introduce a relation that I will call 'determination' which we will use to build GR's content from its models. The basis for this relation will be a mapping between the set of mathematical points of models and the fusion of physical points of worlds. There are a few features that this mapping must have. First, it will have to be one-to-one and onto, otherwise a GR model may be mapped to a world at which which spacetime is not a continuum. Second, it will have to involve a means of correlating the structure \mathcal{S} of a GR model with the structure \mathbf{S} of a world. Otherwise, for example, models with certain global topologies may mapped to worlds with radically different global topologies. Finally, the determination relation must form the basis of a correspondence between the metric and stress-energy tensor fields g and T of models, and their physical correlates, the perfectly natural properties \mathbf{g} and \mathbf{T} .

Intuitively, what we want to require is that a certain mathematical point in a model has a particular value of g or T *iff* the physical correlate of that point has *that same value*. But we can't quite do that since physical points don't ever have values of g and T . Rather, they have instances of \mathbf{g} and \mathbf{T} . So what we need here is to appeal to some previously established correlation between values of g and T and instances of \mathbf{g} and \mathbf{T} . It won't matter all that much exactly what this correlation is, so long as we stick to it for the purposes of evaluating the relation under consideration. That said, there are surely some features that would be nice. For example, it would be nice if, given values a and b of g , where a is least, and b greatest, corresponding values, \mathbf{a} and \mathbf{b} , of \mathbf{g} were such that \mathbf{a} was greatest and \mathbf{b} least. There are many other features that would also be nice, although what these features are taken to be will depend upon ones position with respect to issues in the theory of measurement.²⁰ In any even, let's call a correlation that has all of these nice features – what ever they are taken to be – a *natural* correlation

With this in place we are now ready to introduce our relation. Let us suppose that w is a world, that Y is the set of physical points that compose w and that \mathbf{g} and \mathbf{T} are perfectly natural properties of w that are naturally correlated with g and T . Let us also define a predicate $\text{FUS}(a_1, \dots, a_n)$ which is the fusion of a_1, \dots, a_n . Now we can say that the model $\mathfrak{M} = \langle (X, \mathcal{S}), g, T \rangle$ **determines** w *iff*:

There exists a bijective function $f : X \rightarrow Y$ such that:

1. For every subset x of X and $y = \text{FUS}(f(a) | a \in X)$ of Y , x has \mathcal{S} *iff* y has \mathbf{S} .
2. Suppose v_g and v_T are values of g and T , respectively, and $d_{\mathbf{g}}$ and $d_{\mathbf{T}}$ are the corresponding perfectly natural properties. Then, for every $x \in X$ and $y \in Y$:

²⁰See, for example, Krantz et al (1971), and Eddon (2013).

- 2.1. $v_g[x]$ only if $d_{\mathbf{g}}[f(x)]$, and $d_{\mathbf{g}}[y]$ only if $v_g[f^{-1}(y)]$, and
- 2.2. $v_T[x]$ only if $d_{\mathbf{T}}[f(x)]$, and $d_{\mathbf{T}}[y]$ only if $v_T[f^{-1}(y)]$, and
- 2.3. if v_g and v_T are coinstantiated by x , then the corresponding $d_{\mathbf{g}}$ and $d_{\mathbf{T}}$ are coinstantiated by the image under f of x , and if $d_{\mathbf{g}}$ and $d_{\mathbf{T}}$ are coinstantiated by y , then the corresponding v_g and v_T are coinstantiated by the inverse image under f of y .

To sum up: the determination relation holds between \mathfrak{M} and w *iff* there is some way of putting mathematical points of \mathfrak{M} into correspondence with physical points of w , so that under that correspondence, corresponding collections of points have corresponding structures, and corresponding points have corresponding values of g/T and instances of \mathbf{g}/\mathbf{T} .

Now recall that on the robust conception of models, GR gets its content through the medium of its models. I propose that we can spell out GR's content by invoking the determination relation. We may say the following.

Content: for all worlds w , w is part of the content of GR *iff* there is some model of GR that determines w .

The overall picture that emerges is one according to which GR picks out a collection of models. These models then enter into the determination relation with various possible worlds. And it is in virtue of being so related to GR's models that these worlds come to form GR's content. So, now that we have this apparatus at our disposal, what does it do for us?

2.6 Consequences of The Robust Conception

Under **Content**, (P1) of the hole argument is false. But so too are (P1*) and (P2**). This is because while members of a common D-class agree with one another about which individualistic and non-individualistic properties are instantiated, they

differ in how these properties are *co*-instantiated. But the determination relation tracks individualistic and non-individualistic properties without tracking how these properties are co-instantiated. So the determination relation respects the features that elements of a common D-class have in common, but is blind to the features with respect to which they differ. Let's consider the consequences of this.

2.6.1 Determinism

To begin, observe that the direct link between models and worlds advocated by the robust conception licenses two different ways of viewing determinism. On the one hand, we can begin by thinking of our models as picking out the worlds that form the content of our theory. Once these worlds have been picked out, we can forget how we got them, and go about whatever sort of philosophizing we like. We may, for example, formulate a thesis of determinism like WD that is purely in terms of worlds, and that ignores the way that those worlds were picked out. But we can also formulate a notion of determinism that is sensitive to the way that the worlds are picked out by the models. Arguably, on the robust conception it is this later sort of determinism that is most closely associated with the physical theory. Recall that, according to the robust conception, physical theories are interpreted in terms of models. The models then function as devices by which the theory gets its content. It therefore seems plausible to suppose that the notion of determinism that respects this chain of dependence is the one that is most relevant to GR. Consider the following proposal.

Robust Determinism (RD): GR satisfies RD *iff*: for all GR models, \mathfrak{M} and \mathfrak{M}' , for all cauchy surfaces Σ in \mathfrak{M} and Σ' in \mathfrak{M}' , and for all classes of worlds W determined by \mathfrak{M} and W' determined by \mathfrak{M}' : (1) if Σ and Σ' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in Σ and Σ' are

related by a smooth transformation, then \mathfrak{M} and \mathfrak{M}' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in \mathfrak{M} and \mathfrak{M}' are related by a smooth transformation, and (2) $W = W'$.

GR satisfies RD. When cauchy surfaces of a pair of GR models are such that their individualistic and non-individualistic properties are related by a smooth transformation, then it follows that the models from which these cauchy surfaces are selected are such that their individualistic and non-individualistic properties are globally related by a smooth transformation. But it also follows that when models satisfy this criterion, they are guaranteed to determine the same classes of worlds. That is, the features that allow models to satisfy RD also force those models to determine the same classes of worlds.

2.6.2 Haecceitism and Antihacceitism

Given **Content**, haecceitism is irrelevant to the hole argument. To see why, consider a model \mathfrak{M} and the class W of worlds that it determines, and the model \mathfrak{M}' and the class of worlds W' that it determines. Now, under the assumption of haecceitism, W and W' will have many members – at least one member for each distinct way of permuting individualistic and non-individualistic properties. Under the assumption of antihacceitism, on the other hand, each of these classes will have only one member. This is because the antihacceitist does not think that worlds ever differ merely in virtue of differences in how the individualistic and non-individualistic properties are co-instantiated. But the cardinality of the classes of worlds is besides the point. What matters, so far as satisfying RD is concerned, is just whether \mathfrak{M} and \mathfrak{M}' can ever differ with respect to the worlds they determine. If \mathfrak{M} and \mathfrak{M}' are

members of a common D-class, then the answer is “no”. This means that it makes no difference whatsoever whether one is a haecceitist or an antihaecceitist.²¹

2.7 Conclusion

The heuristic conception has shaped the dialectic around the hole argument. The effect of this has been to make haecceitism appear relevant when it is not. This happens as the result of the following pattern. First, the heuristic conception does not give models a clear role to play in the relationship between the formulation of a physical theory and that theory’s content, but it nevertheless employs them in non-trivial ways. In particular, advocates of the heuristic conception appeal to models for purposes of visualizing GR’s content. The way this visualization proceeds is by initially maximizing respect of similarity between models and worlds. This involves taking all of the distinctions recognized among GR models to be recognized among GR worlds. Of particular importance is the reproduction of mere individualistic difference among D-classes of models as mere haecceitistic difference among classes of worlds. Importing this distinction among models to the domain of worlds makes **Haec** seem a plausible principle of correspondence between models and worlds. **Haec** may then be used as the main premise in the hole argument against substantivalism.

The philosophical community has now reached a stable consensus: the way to respond to the hole argument against substantivalism is to stage one among a variety of metaphysical interventions, each of which has the effect of redirecting the force of the

²¹Note that the haecceitist substantivalist will find herself in situation where she can acknowledge several notions of determinism. On the one hand, she may observe that there is a notion of determinism – WD – that is sensitive to haecceitistic differences that her preferred metaphysical package violates. But on the other hand, she may acknowledge that there is another sort of determinism – RD – that her preferred metaphysical package does to violate. Moreover, the haecceitist substantivalist may point out that RD is the notion of determinism that is most closely associated with GR. She is licensed to make this point in virtue of her acceptance of the robust conception which gives models a clear and central role in connecting a theory to its content. So the haecceitist substantivalist who accepts the robust conception finds herself in a situation where she can have her cake and eat it, too.

hole argument. Either – like Butterfield, Hofer, Maudlin, and Pooley – we buttress substantivalism with some modal thesis that blocks the generation of problem-causing worlds by fiat. Or else – like Arntzenius, Melia, and Skow – we fiddle with the formulation of determinism to make it weak enough not to bother haecceitism. What all of these responses have in common is that they accept the heuristic conception, and agree that haecceitism is relevant to the hole argument.

But there is another way. We can undermine the hole argument by rejecting a more fundamental assumption: the heuristic conception. We may instead accept the robust conception that takes models to play a central role in connecting theories to worlds. If we accept the account I give, and take the content of GR to be formed according to **Content**, then whether GR is deterministic or not will not depend upon whether one is a haecceitist or an antihaecceitist. The question of haecceitism as opposed to antihaecceitism – the central point of contention in the aftermath of the hole argument – becomes irrelevant.

CHAPTER 3

DUALISM AND THE PRICE OF HARMONY

3.1 Introduction

Spacetime substantivalism is the view that the basic parts of spacetime are fundamental entities. There is a live debate among substantivalists that concerns the relationship between material objects and parts of spacetime. *Supersubstantivalists* maintain that material objects are identical to parts of spacetime; *dualists* maintain that material objects are distinct from the parts of spacetime, but that they bear a fundamental relation of *occupation* to some of them.

Supersubstantivalists and dualists alike find it plausible to suppose that there is a principle of harmony that governs the relationship between material objects and parts of spacetime. This principle says, in effect, that if a material object is appropriately related to a part of spacetime, then that object and that part of spacetime have exactly the same geometric, topological, and mereological structure. For the supersubstantivalist, this appropriate relation is *identity*: if a material object is identical to a part of spacetime, then that object and that part of spacetime have exactly the same shape. For the dualist, the appropriate relation is *occupation*: if a material object occupies a part of spacetime, then that object and that part of spacetime have exactly the same shape. But whereas for the monist this sort of shape harmony is a consequence of Leibniz' Law, for the dualist it seems to be a brute fact. I will argue that this consideration in conjunction with adequately developed Humean intuitions about recombination provides a strong argument against dualism.

In this chapter I will proceed as follows. In §3.2, I introduce some familiar metaphysical notions that we will need throughout. Then, in §3.3, I describe the dualist’s fundamental occupation relation and some related notions. In §3.4 I develop the principle of harmony, and show that there are three ways that a dualist may explain it. I call these explanations ‘shape principles.’ In §3.5 I introduce a framework that fleshes-out some broadly Humean intuitions about modal plenitude. Then, in §3.6, I show how this framework undermines each of the dualist’s shape principles. I take this to show that the dualist’s occupation relation is a troublemaker. Finally, I conclude that if the occupation relation is a troublemaker, then we have a reason to favor a theory that doesn’t require it: supersubstantivalism

3.2 Preliminaries

I will suppose that the basic parts of spacetime are spacetime points, and that classical mereology applies to them. This allows us to characterize spacetimes as fusions of spacetime points, and it allows us to characterize *regions* of spacetime in terms of the parthood relation.

Next, I will assume a modal framework that is roughly that of Lewis (1986). I take worlds to be maximally interrelated sums¹; I take *de dicto* modal claims to be claims about what worlds there are; and I take *de re* modal claims to be analyzable in terms of counterparts. This allows me to understand modally qualified claims as involving quantification either over worlds or over counterparts.

I take there to be a primitive notion of *perfect naturalness* that applies to properties and relations. The perfectly natural properties and relations are special in the sense that they are “joint carving”: when things share a perfectly natural property, they are objectively similar, and when pairs of relata share a perfectly natural relation,

¹Lewis thinks of the relevant sort of interrelatedness as being spatiotemporal. I don’t agree with this assumption. To see why, read on.

the pairs are objectively similar. Given the notion of perfectly natural properties and relations, we may characterize duplication. Two individuals are *duplicates* iff there is a similarity map from parts of one to parts of the other such that corresponding parts have the same perfectly natural properties and stand in the same perfectly natural relations.

The concept of duplication can be used to characterize the notion of an intrinsic property. A property P is *intrinsic* iff no two duplicates anywhere in logical space disagree with respect to P , otherwise P is *extrinsic*. Then we may say that the *intrinsic character* of a thing is an exhaustive catalogue of that thing's intrinsic properties. Next we can say that a relation R is *internal* iff R supervenes on the intrinsic character of its relata. Similarity is the paradigm: whether A and B are similar depends only on the intrinsic characters of A and B , it does not depend on their locations, or their arrangement among other things. We may say that a relation R is *external* iff it does not supervene on the intrinsic character of its relata taken individually, but it does supervene on the intrinsic character of its relata taken together. Spatiotemporal distance relations are the paradigm: whether A and B are five feet apart does not supervene upon the intrinsic characters of A and B , but it does supervene on the intrinsic character of their fusion.² Finally, any relation that is internal or external is *intrinsic*, otherwise it is *extrinsic*.³

Finally, recall that I characterized worlds as “maximally interrelated sums.” Lewis thinks of worlds as maximally *spatiotemporally* interrelated sums. This has the effect of making spatiotemporality an essential ingredient of reality. But there are compelling grounds for skepticism about this. One reason for skepticism comes from physics. Quantum mechanics and general relativity are famously successful physical

²Although, note that there may be some reasons for thinking distance relations are not external. See, for example, Bricker (1993) and Maudlin (2007).

³The characterization I give follows Bricker (1996).

theories; equally famous is their apparent incompatibility. Physicists' attempts to reconcile the two theories have spawned radical new proposals for the fundamental structure of the world. Among these are proposals according to which there are no natural distance relations.⁴ This strongly suggests that we ought to think that there are worlds at which there are no natural distance relations – indeed, ours might be one of them.

But note that if there are non-spatiotemporal worlds, then the Lewisian analysis of 'world' needs to be augmented. The analysis that I accept appeals to *external* rather than spatiotemporal interrelatedness. We should say that worlds are maximal externally – rather than *spatiotemporally* – interrelated sums.⁵ My assumption that there are worlds unified by natural, external, non-spatiotemporal relations will be required for several of the arguments I consider below.

3.3 Occupation

There are a variety of occupation relations that different philosophers have considered as candidates for the sort of fundamental relation a dualist needs to unite objects and regions of spacetime.⁶ These include notions like *weak occupation*, which, when it obtains between an object O and a region R , requires only that some part of O be “in” some subregion of R . Consider, for example, the region within my office. I weakly occupy this region as soon as I put my arm through the door.⁷ There are other notions as well. Like, for example, *whole occupation*, according to which an object O occupies a region R just in case there is no part of O “outside” of R . According

⁴See Butterfield and Isham (2001) for an overview of the various research projects in quantum gravity, include ones according to which there are no natural distance relations.

⁵See Bricker (1996) for detailed discussion of this proposal.

⁶Parsons (2007), Hudson (2006), Saucedo (2011), Gilmore (2008).

⁷This is Parsons' (2007) “weak location”.

to this notion, I wholly occupy the region within my office when I stand inside of it, despite their being subregions of it that are “free of” me. Finally, the notion I prefer is *exact occupation*.⁸ Intuitively, an object O exactly occupies a region R just in case every part of O is “in” some part of R , and every subregion of R has some part of O “in” it. As Parsons (2007) puts it, the region an object exactly occupies is like its “shadow in spacetime”.

Notice that I used quotes around words like ‘in’ and ‘outside’ in the preceding paragraph. I did this because these sorts of notions are ultimately to be characterized in terms of the particular occupation relation we choose to treat as a primitive. So, although we can use such words to get a feel for things, we cannot innocently use these words in characterizing the occupation relation. So what we’ll do now is pick one of the notions of occupation discussed above, and treat it as a primitive. Then, appealing to this primitive in conjunction with the resources of mereology, we can define other sorts of occupation relations in terms of it. I already expressed a preference for exact occupation, so we’ll take that to be our primitive. This then enables us to define weak and whole occupation. An object O *weakly occupies* a region R just in case there is some region R^* such that O exactly occupies R^* and R^* overlaps R . And an object O *wholly occupies* a region R just in case there is some region R^* such that O exactly occupies R^* and R^* is a subregion of R .

3.4 Harmony

I’ll call the sum of the features of an object or a region that are determined by its geometry its *shape*.⁹ So to say that such-and-such and so-and-so are the same

⁸This is also the preferred primitive occupancy relation of Casati and Varzi (1999), Saucedo (2011), and Leonard (forthcoming).

⁹McDaniel (2007) characterizes shape this way, too.

shape entails that such-and-such and so-and-so have the same topological, affine, and metrical structure.¹⁰ Now we may express the principle of harmony as follows.

Harmony: for every world w , if there is some material object O and some region of spacetime R that are both parts of w , and O exactly occupies R , then O and R have the same shape.

The business of this paper is to elaborate on premise (1) of the following argument.

The Argument Against Dualism

1. If harmony is true, then dualism is false.
2. Harmony is true.
3. So, dualism is false.

But I anticipate that some committed dualists may be tempted to reverse my argument and offer the following one instead.

The Argument Against Harmony

- 1*. If harmony is true, then dualism is false.
- 2*. Dualism is true.
- 3*. So, harmony is false.

¹⁰Here is the place to flag a related notion of harmony. There is a much-discussed issue regarding the relationship between the mereological structure of material objects and the mereological structure of the regions those objects exactly occupy. This issue regards the question: is it metaphysically possible for these respective structures to be in misalignment? Is *mereological disharmony* possible? Simmons (2004), McDaniel (2007) Saucedo (2011), and Uzquiano (2011) argue that disharmony is possible; Skow (2007) and Schaffer (2009) argue that it is not possible; and Leonard (forthcoming) offers a taxonomy of restrictions one might place on the sorts of disharmony that are possible. While the motivations behind *mereological* harmony are similar to those behind *shape* harmony, the two notions are different enough to be considered separately.

But it is important to note that premise (2) enjoys far more intuitive support than premise (2*). The intuition that round things won't fit in square places, for example, is more robust than the intuition that objects are related to regions of spacetime by an occupancy relation. The relative weight of these intuitions is important. It strongly militates against using dualism as a premise in an argument against **Harmony**. So, while I neither offer a positive argument for **Harmony** nor assume that an argument against **Harmony** is necessarily doomed, I will insist that a simple reversal of my argument – like The Argument Against **Harmony** – is illegitimate. This is as strong a claim as I will need.

3.4.1 The Shapes of Objects and their Regions

Given the dualist's bifurcation of the material and the spatiotemporal, what might explain **Harmony**? Is it simply a miraculous coincidence or brute necessity that objects happen to find the right regions to exactly occupy? Or is there some relationship between the shapes of objects and regions that makes **Harmony** inevitable?

There are three mutually exclusive and jointly exhaustive principles concerning the relationship between the shape properties of objects and the regions they exactly occupy. They are as follows:

Imposition: for every world w , every material object O that is part of w , and every region of spacetime R that is part of w : if O exactly occupies R , then O *imposes* its shape on R .

Inheritance: for every world w , every material object O that is part of w , and every region of spacetime R that is part of w : if O exactly occupies R , then O *inherits* its shape from R .

Independence: for every world w , every material object O that is part of w , and every region of spacetime R that is part of w : O and R have their shapes *independently*.

Note that under **Inheritance** the shape properties of regions are intrinsic, while the shape properties of material objects are extrinsic, and under **Imposition** the shape properties of regions are extrinsic, while the shape properties of material objects are intrinsic. Because of this, it is clear that both **Inheritance** and **Imposition** are capable of providing genuine explanations of **Harmony**. **Harmony** is satisfied in the case of **Inheritance** because shapes are had by material objects only in virtue of their relations to regions, and **Harmony** is satisfied in the case of **Imposition** because shapes are had by regions only in virtue of their relations to material objects. But under **Independence** shape properties of both regions and material objects are intrinsic. So there is no clear sense in which **Independence** *explains* **Harmony**. What we'll see presently is that problems with **Inheritance** and **Imposition** push dualists towards **Independence**. But **Independence** will force dualists to concede that **Harmony** is just a brute necessity.

To see how these arguments get off the ground, we need to consider how metaphysical theories couple to possibilities. This will first involve us in a consideration of *principles of plenitude*.

3.5 Modal Plenitude

A principle of plenitude is a principle according to which logical space is constituted. It is a principle that ensures that all the possibilities are accounted for; it ensures that there are no gaps – no holes in logical space where a possibility should have been, but isn't. Lewis (1986) discusses and defends one such principle, inspired by the Humean denial of necessary connections between distinct existents. This is his principle of recombination. A first approximation of the principle says that “anything can coexist with anything else [and] anything can fail to coexist with anything else. (1986, p. 88).” But Lewis' principle of recombination is problematic in a number of ways. It is problematic, first, because Lewis never articulates it in sufficient detail,

and second, because all attempts to develop the principle as a comprehensive account of plenitude seem to fail.¹¹

Bricker (1991, forthcoming), however, has proposed to extend Lewis' picture by taking plenitude to require not one principle, but four. To see how these principles work, begin by thinking of worlds as generally being divisible into two parts: their structure, and their contents. In familiar cases, structure is spatiotemporal structure: a system of distance relations instantiated by parts of substantial spacetime. (Although note that I assume that some worlds have non-spatiotemporal structure.) A world's contents, in familiar cases, are the material objects or material fields that are parts of that world. Bricker proposes a principle of plenitude for structures, a principle of plenitude for contents, and a principle of recombination that redistributes contents across structures. These three principles put meat on the bones of the intuition that "anything can coexist with anything else." The fourth principle – the principle of solitude – puts meat on the bones of the intuition that "anything can fail to coexist with anything else."¹²

I will need three of these principles: plenitude of structures, plenitude of recombinations, and the principle of solitude.

3.5.1 A Plenitude of Structures

What motivates the principle of plenitude for structures? Here is one sort of motivation: we don't know exactly what shape our universe is. For all we know, it might be finite or infinite; it might be bounded or unbounded; it might be flat, or either positively or negatively curved; it might be simply connected, or have holes. These are just some of the possible structural features that are compatible with our

¹¹See, for example, Forrest and Armstrong (1984), Divers and Melia (2002), Effird and Stoneham (2008), and Bricker (forthcoming).

¹²Bricker (forthcoming) argues that the principle of solitude is not a fourth, fundamental principle of plenitude, but a consequence of the principles of plenitude of structures and recombinations.

best large-scale physics: general relativity (GR). But perhaps our best physics isn't right – indeed, it probably isn't. In that case, the sorts of possibilities I just mentioned may not even begin to plumb the depths of what's compatible with what we know about our world. This suggests that we ought to think logical space includes a wide range of structures.

To begin, we need the notion of a *natural generalization* of a class of structures.¹³ Consider some class of structures \mathcal{S} . A natural generalization \mathcal{G} of \mathcal{S} is a superclass of \mathcal{S} that contains elements that naturally generalize certain features possessed by members of \mathcal{S} . (Here what makes a generalization *natural* is just that it be accepted as a natural or intuitive generalization by mathematicians.) Some examples will help to clarify. Consider 3 dimensional Euclidean space. There are several features of 3 dimensional Euclidean space that may be naturally generalized. First, we may generalize the dimension, which will give us the class of all finite dimensional Euclidean spaces. Second, we may generalize the metric, which will give us all of the 3 dimensional Riemmanian spaces. Or we may generalize both at once to get all of the finite dimensional Riemmanian spaces.¹⁴ Given the notion of a natural generalization of a class of structures, we can state the principle of plenitude of structures as follows:

Plenitude of Structures: if \mathcal{S} is a class of possible structures, and \mathcal{G} is a natural generalization of \mathcal{S} , then any member of \mathcal{G} is a possible structure.

Several brief comments are in order. First, this principle does not supply initial justification for any particular possibilities.¹⁵ Rather, it takes a commitment to a cer-

¹³Cf. Bricker (1991), and Belot (2011).

¹⁴See Belot (2011, chapter 2) for relevant discussion.

¹⁵Initial justification for believing certain structures to be possible is a function of those structures' roles in theorizing about the actual world. If a "structure plays, or has played, an explanatory role in our theorizing about the actual world (Bricker, 1991, p. 609)", then we are justified in believing that it is a possible structure. Once a structure earns this status, it then becomes an appropriate basis for justifying the possibility of other, related structures.

tain collection of possibilities as input, and generates another collection of suitably related possibilities as output. So, for example, **Plenitude of Structures** does not commit one to the possibility 3 dimensional Euclidean space. However, given a prior, independent commitment to the possibility of 3 dimensional Euclidean space, **Plenitude of Structures** commits one to the possibility of all suitably related spaces. These, as I mentioned, will include all of the finite dimensional Euclidean and Riemannian spaces. Finally, that a structure is possible does not mean that it is a world structure. **Plenitude of Structures** only entails that natural generalizations of \mathcal{S} are instantiated by *parts* of worlds.

3.5.2 A Plenitude of Recombinations

The principle of recombination is meant to provide for a plenitude of arrangements of contents across structures. Here in this world we have trees and cars, and there in some other world is some exotic structure. Shouldn't we think that trees and cars might have found their way into that exotic structure? If so, that possibility won't come from either the principle of plenitude for structures or for contents. It will come from a principle that recombines contents across structures.

To state the principle, we need to introduce a few terms. First, we need the notion of an arrangement of a class of world-contents within a structure. We'll say that an *arrangement* of a class C of world-contents within a structure S is a category-preserving, one-to-one or one-to-many mapping from elements of C to places in S . We'll say that an arrangement A is *non-overlapping* iff A never maps distinct contents to overlapping places in a structure. Next, we need the notion of a world recombining a class of world-contents according to an arrangement. We'll say that a world w with structure S *recombines* a class C of world-contents *according to* an arrangement A iff (1) A is an arrangement of C in S , and (2) for all elements $c \in C$ and places $s \in S$,

if A maps c to s , then s is occupied in w by a duplicate of c . Now we are ready to state the principle.

Principle of Recombination (PR): For any class C of world-contents, and any non-overlapping arrangement A of C , there exists a world that recombines C according to A .¹⁶

3.5.3 The Principle of Solitude

The principles we have so far considered won't allow us to cut worlds apart and thereby make new worlds. They won't, for example, allow us to reason from a world with a pair of iron spheres to a world with a single iron sphere. To do that we need a principle that underwrites the denial of necessary connections between distinct existents. The principle of solitude is meant to do just that.

Principle of Solitude (PS): for any possible individual a , there is a world w containing a duplicate of a , and containing no individual that is not a part of a .

3.6 Problems for Exact Occupation

We now have a collection of adequately developed principles of modal plenitude in hand. I have already shown that, given dualism, there are three ways to explain **Harmony**. I will now show that, given the principles of plenitude, each of these explanations leads to trouble.

3.6.1 Against Imposition

The problem with **Imposition** is that it makes material objects do too much of the work, and in so doing it incurs all the costs associated with relationalism without reaping any of its benefits. Let's be precise about how and why this is the case.

¹⁶This section reproduces parts of Bricker (forthcoming).

According to lore, the substantivalist commitment to the existence of spacetime has provided an elegant and straightforward foundation for classical mechanical physics. Physicists have traditionally offered explanations of mechanical phenomena in terms of material bodies situated within a medium of spacetime points, where these spacetime points instantiate a variety of structures (e.g., topological, metrical, affine, and differential). By accepting these spacetime points into their ontology, substantivalists are able to take the physicists' explanations at face-value. The relationalist, on the other hand, regards spacetime points with suspicion; she would prefer to provide an interpretation of our physics that does not involve reference to them. But it turns out that providing such interpretations, or rewriting the physics so as to avoid mention of spacetime, is extremely difficult. So what we have here is a sort of a trade-off. In exchange for an ontological cost, the substantivalist gets an elegant foundation for physics, and in exchange for explanatory acrobatics, the relationalist gets a lean ontology. But for a substantivalist to accept **Imposition** is for her to embrace the worst of both worlds: she is stuck with the ontological burden, but she also saddles herself with the explanatory acrobatics.

To see exactly why, first distinguish the *material shape* of a world from the *spatiotemporal shape* of a world. The spatiotemporal shape of a world is the sum of the topological, affine, and metrical properties and relations instantiated at that world. The material shape of a world is the sum of the topological, affine, and metrical properties and relations instantiated by the parts of matter at that world. It is intuitive and uncontroversial that many worlds have both spatiotemporal and material shape. The question is what sorts of entities instantiate spatiotemporal shape, and what sort of relationship there is between spatiotemporal shape and material shape.

The typical substantivalist claims that the spatiotemporal shape of a world w is instantiated by the fusion of spacetime points at w , and the material shape of w is the shape of the fusion of spacetime points at w that are occupied by matter.

The relationship between spatiotemporal and material shape is thus straightforward: the latter is embedded within the former. This conception of spacetime and its relationship with material objects grounds certain highly intuitive possibilities. For example, it seems very plausible to suppose that empty space is at least possible. Even if worlds like ours are entirely pervaded by fields, and this makes them count as entirely non-empty, surely there are classical particle worlds that have empty regions. The possibility of empty space is easy for a typical substantivalist to ground: it's just a region that is unoccupied by matter. Consider another sort of possibility. Begin with a world w the spacetime of which is 3 dimensional and Euclidean, and whose lone material occupant is an iron sphere. If there is a world like w , then it follows from Plenitude of Structures and Plenitude of Recombinations that there is a world v the spacetime of which is 4 dimensional and Euclidean, and whose lone occupant is a duplicate of the iron sphere in w . Again, the substantivalist explanation is an easy one: spatiotemporal shape is entirely independent of material shape, so a single material configuration is compatible with many different spatiotemporal shapes.

These sorts of possibilities are problematic for the both the relationalist and the **Imposition** substantivalist because both accounts seem to require that spatiotemporal shape be identified with material shape. In the case of the relationalist, this identification seems inevitable because material objects are the only things around to instantiate shapes – there aren't any other entities around to do the job. In the case of the **Imposition** substantivalist, this identification seems inevitable since unoccupied regions are shapeless regions. Thus, in both cases, the only shapes around are those that are possessed by material objects, so the spatiotemporal shape of a world must be the material shape of that world – there aren't any other obvious candidates.

Of course, a lack of *obvious* candidates doesn't mean a lack of candidates *simpliciter*. In fact, the strategy most often pursued by relationalists is to claim that the spatiotemporal shape of a world is fixed, not by matter, but by facts about where

matter *could have* been. This involves the invocation of a *sui generis* form of modality. Whether this sort of modality can be reduced to non-modal notions depends on the details of one's account. However, Belot (2011) argues that a relationalist who wants to have a one-to-one correspondence between the space of relationalist possibilities and the space of substantivalist possibilities will have to take this modality to be primitive.

It is fair to expect that the **Imposition** substantivalist, *qua* substantivalist, be in a position to accommodate the full range of possibilities typically associated with substantivalism. However, just as the relationalist requires primitive modality in order to underwrite a one-to-one correspondence between the space of relationalist possibilities and the space of substantivalist possibilities, so too will the **Imposition** substantivalist require primitive modality in order to underwrite a one-to-one correspondence between typical substantivalist and **Imposition** substantivalist possibilities. But whereas the cost of accepting this sort of primitive modality might be tolerable for a relationalist, since it is offset by the benefit of a lean ontology, it is absolutely intolerable for a substantivalist, who has no analogous means of righting the balance of theoretical costs and benefits.

3.6.2 Against Inheritance

Consider the fusion M of actual material objects. It follows from the Principle of Solitude that there is some world w that contains a duplicate of M and contains no individual that is not a part of M . That means there is some world w with just material objects, no spacetime. If **Inheritance** is true, then material objects in such worlds have no shapes at all. Pause for a moment to reflect on how truly odd that would be.

In a world where material objects have no shape, since shape comprehends such things as topological structure, there is nothing to ground notions of continuity and

connectedness. It's not necessary to go too far into technical details to see the consequences of this. It would mean, for example, that given such a world and a duplicate of me within that world, there is no fact of the matter concerning that duplicate whether his heart is in his chest or on his elbow, whether the duplicates of neurons that are adjacent in my brain are adjacent in his, and so forth. Indeed, it's hard to see how in such worlds there could be things like people.

In a world where material objects have no shape, not only are they neither square, nor round, nor . . . , but there is also no fact of the matter about whether they have one side longer than the other, or whether one object is larger than another. So, for example, if we consider a spacetime-free world populated by a duplicate of my car and a duplicate of my garage, there is no fact of the matter about whether, in that world, my car fits in my garage. The relation of comparative size is dependent upon the shape properties of the objects being compared, so with no shape, there are no facts about comparative size. Odd!

Now, I suspect that one might be tempted to reply on behalf of **Inheritance** that it is only natural for a substantivalist to think of spacetime as carrying as much of the explanatory burden as possible. At the very least, it is natural for a substantivalist to think that spacetime points and regions instantiate all the spatiotemporal properties and relations, like topological, affine, and metric properties and relations. So given this, it is only natural to think that in a world with no spacetime, material objects will lack all those sorts of paradigmatically spatiotemporal properties. Of course material objects behave strangely in non-spatiotemporal worlds, the reply continues; non-spatiotemporal worlds are strange parts of logical space!

But I think this reply misses the point. If you have already accepted the **Inheritance** view, then you will likely be inclined to adjust to whatever possibilities it entails. The point is, consequences that are pre-theoretically very odd have to count against the view when we are surveying options. And it cannot be denied that the

sorts of possibilities discussed above are pre-theoretically very odd. It's not that we could never be justified in living with possibilities of this sort. It's just that, other things being equal, it would be better if our theory didn't entail them.

3.6.3 Against Independence

Dualism is the default version of substantivalism. This is so primarily because it seems to accord better with common sense. Monism, so the story goes, seems weird since it entails, for example, that regions of spacetime fight and love and sing. Regions of spacetime, so common sense tells us, don't *do* anything; they are – as the dualist says – the silent and immobile containers for things that fight and love and sing.¹⁷

But when dualism is thought of as enjoying the advantage of agreeing with common sense, it cannot be dualism under the guise of the **Inheritance** or **Imposition** view; both of these views are wildly counterintuitive. The version of dualism that agrees with common sense is dualism under the guise of the **Independence** view. This is the version of dualism according to which the universe is like a cosmic jello mold: just as little grapes and marshmallows fit neatly into little grape and marshmallow-shaped cavities in the jello, people and houses and so forth fit neatly into little people-shaped and house-shaped regions of spacetime.

But **Independence** is just as problematic as **Inheritance** and **Imposition**. The problem is that **Independence** together with Plenitude of Recombinations entails that **Harmony** is false. Confronted with this, the dualist may either restrict Plenitude of Recombinations in ways that block the entailment, or else she may employ some semantic trickery to dodge the entailment.. Each of these options leads to its own sort of trouble. In the following section we'll see why, exactly, **Independence** together with Plenitude of Recombinations entails that **Harmony** is false, then we'll

¹⁷ Cf. Schaffer (2009, p. 133), Sider (2001, p. 111), Skow (2005, p. 58).

examine the troubles the dualist faces when she restricts Plenitude of Recombinations or engages in semantic trickery.

3.6.3.1 Independence, Harmony and PR are Inconsistent

According to dualism, material objects and regions of spacetime are ontologically distinct and are bound by a fundamental external relation of exact occupation. It follows from Plenitude of Recombinations, then, that we may mix and match regions with objects, since we may mix and match the relata of external relations. To illustrate, where R is some external relation, if $R[A, B]$ and $R[C, D]$, then there is some world where a duplicate of A bears R to a duplicate of D , and some world where a duplicate of A bears R to a duplicate of C , . . . and so on. Now, when R is the relation of exact occupation, and the relata are material objects on one hand and regions of spacetime on the other, then strange possibilities ensue.

Some strange possibilities are these: if a material object O exactly occupies some region S , then there is some world where a duplicate of O exactly occupies another duplicate of O , and some world where a duplicate of S exactly occupies a duplicate of S , and some world where a duplicate of S exactly occupies a duplicate of O . Perhaps these can be got around by appeal to the logical properties of exact occupation: it is intransitive, irreflexive, and asymmetric. In essence, the relation respects ontological categories – it is a relation that necessarily holds between a thing in the category of material object and a thing in the category of spacetime region.

That sounds reasonable, and that does block some of the outlandish possibilities, without really seeming to count as much of a restriction on Plenitude of Recombinations. However, this still leaves plenty of room for trouble so long as we are assuming **Independence**. According to **Independence** the shapes material objects and regions of spacetime are logically independent – neither depends in any way on the other. So now consider a world wherein a rectangular object exactly occupies a rect-

angular region of spacetime, and a circular object exactly occupies a circular region of spacetime. It follows from this possibility that there is some world where the objects trade regions. That is, it follows that there is a world where a rectangular object exactly occupies a circular region and a circular object exactly occupies a rectangular region. This is an obvious violation of **Harmony**.

3.6.3.2 Restricting PR

Once one has formulated general principles of plenitude like those outlined in §3.5, it is worrisome to introduce amendments that constrain the range of possibilities they generate. The worry is that many sorts of constraints would be ad hoc, and that the theory that requires them is likely to be false. That said, it is natural to think that the principles of plenitude have certain constraints already built in. As I mentioned in the previous section, for example, once one has parceled out reality into a cluster of distinct ontological categories, it is reasonable to expect principles of plenitude to respect those categories. So, to take another example, given certain views about the relationship between particulars and properties, it is natural to think that it is impossible for a property to instantiate a particular – elements of the category of particulars instantiate elements of the category of properties, never the other way around.

So could an **Independence** substantialist plausibly argue that shape properties are deep enough features of reality that principles of plenitude always respect them in ways that maintain **Harmony**? This is highly implausible. But even if we grant this, the **Independence** substantialist looks bad in comparison to the supersubstantialist. For the supersubstantialist, **Harmony** is grounded in occurrent facts within each world: every material object in every world is identical to some region of spacetime at that world. But for the **Independence** substantialist it is simply a brute fact about logical space that it won't permit violations of **Harmony**. So

to the extent that explanations are preferable to brute facts, supersubstantivalism is preferable to **Independence** substantivalism.

3.6.3.3 Preserving Harmony?

Is there a way to make **Independence**, **Harmony**, and PR consistent? Perhaps an **Independence** substantivalist can find inspiration for a reconciliation in a debate about the relationship between properties and their causal powers. On one side of this debate are necessitarians who think that the causal powers properties confer are essential,¹⁸ and on the other side are contingentists¹⁹ who think that they are not. A typical necessitarian approach may be characterized as follows.²⁰

[Suppose] that properties are individuated by their nomological roles. The essence of a property, on this view, is its place in the Ramsified lawbook. To derive the Ramsified lawbook, conjoin the law statements, uniformly replace each property name by a variable, and prefix the result with a unique-existential quantifier $\exists!$ for each variable. To find the place of a given property, delete its associated quantifier. The resulting open sentence describes the essence of this property. To be that property is to satisfy that sentence. On this view, like charges could not attract because to be charge is to satisfy the place of

According to the sort of view of properties described here, it is of the essence of the property *charge* that it satisfy Coulombs' law. What if a dualist were to think of the relation of exact occupation along these lines? Could one plausibly argue that it is of the essence of the relation *exactly occupies* that it satisfy **Harmony**?

¹⁸ Cf. Shoemaker (1980), Ellis (1999).

¹⁹ Cf. Lewis (2009), and Schaffer (2005).

²⁰ There are other necessitarian approaches. See Schaffer (2005) for discussion of a few of them.

This sort of response is really just an attempt to apply a superficial semantic remedy to a deep metaphysical problem. It is not successful. Consider two different angles on the issue. First, recall the distinction between abundant properties and relations and natural properties and relations. The natural properties and relations carve reality at the joints, they ground cases of similarity, and are always intrinsic. The abundant properties and relations are not like this; they are totally indiscriminating. For every possible predicate, no matter how bizarre, there is an abundant property or relation. Given this, we are guaranteed to refer to some property or relation any time we formulate a predicate, but we are *not* guaranteed to refer to a *natural* property or relation any time we formulate a predicate. So, while abundance guarantees that there is some relation whose essence involves the satisfaction of **Harmony** – which we may agree to refer to with the predicate ‘exactly occupies’ – there is no guarantee that this relation is a *natural* relation. Indeed, ‘exactly occupies’, when restricted in this way, says something less than we ordinarily take it to say. For, arguably, what we ordinarily take it to say involves a perfectly natural relation, a relation that is not the one we have defined, because that defined relation, unlike perfectly natural relations, does not satisfy PR.

Here is the second angle on the same issue. Suppose that we do accept the strategy under consideration and take it to be of the essence of *exactly occupies* that it satisfy **Harmony**. Now, I’ve argued that **Independence** substantialists should not restrict PR because, if they do, their view looks ad hoc. But not restricting PR means that the relata of external relations can be mixed and matched any which way. But now, according to the **Independence** substantialist, there is a natural external relation – never mind for a moment what we’re calling it – that binds material objects to regions of spacetime. It follows from PR that this relation sometimes obeys **Harmony** and sometimes it doesn’t. The solution under consideration, then, boils down to this: we agree to call this relation ‘exact occupation’ in worlds where it satisfies **Harmony**,

and we agree to call it something else in worlds where it violates **Harmony**. But this is the most superficial of solutions. The problematic worlds are still there no matter how we agree to use our words to pick out relations. Simply agreeing not to use the predicate ‘exactly occupies’ to refer to the relation that binds material objects to regions of spacetime at the problematic worlds does not make these worlds any more palatable.

Trying to reconcile **Independence**, **Harmony**, and PR by appeal to necessitarianism is a semantic trick. It is simply a distraction from the fact that the important metaphysical problems remain unsolved and as troubling as ever.

3.7 Conclusion

I have argued that dualism is problematic because it requires a fundamental external relation – exact occupation – that is a troublemaker. Exact occupation is a troublemaker because it conflicts with **Harmony**, and **Harmony** is independently plausible. The conflict between exact occupation and **Harmony** is brought into focus by appeal to some broadly Humean principles of plenitude. There are, of course, many opportunities for a committed dualist to block my arguments. She may reject **Harmony**, she may reject or restrict various of the principles of plenitude, or she may happily accept some of the counterintuitive results I discuss in §6.1 and §6.2. But all of these options have costs. These costs ought to motivate substantialists who are undecided between dualism and supersubstantialism to decide in favor of supersubstantialism.

CHAPTER 4

PROSPECTS FOR SUPERSUBSTANTIVALISM

4.1 Introduction

Our world is made either mostly or completely of material things – material objects, or material fields – separated by various spatiotemporal distance relations. But this leaves much open. In particular, it leaves open the question of just what sort of entity material objects are, and what it is in virtue of which they bear various distance relations. Perhaps the only fundamental, concrete objects there are are the material ones. And perhaps the way to explain spatiotemporality – the fact that material objects are separated by various spatiotemporal distances – is in terms of relations that these material things bear directly to one another. This is *relationalism*. Relationalism is typically traced back to Leibniz, and has enjoyed the support of people like Mach, and, at times, Einstein.

But perhaps relationalists have it wrong. Perhaps the relational ontology omits an entire category of entity. Perhaps material objects are not the only sort of fundamental, concrete objects; perhaps spacetime (or its basic parts) is also a fundamental, concrete object. On this sort of view, spatiotemporality is in the first place a feature of spacetime, and it is in virtue of their relationship to spacetime that material objects are spatiotemporal. This view also has an impressive pedigree, which includes people like Newton, and, at times, Einstein, as well as most contemporary philosophers.

Given substantivalism, there is a further question regarding the nature of the relationship between material objects and parts of spacetime. Is this relationship like the relationship between a muffin-tin and its muffins – the former being distinct from

the later, but with the two united by a relation of occupation. Or is it more like the relationship between an eddy in a body of water – the former being something like a property of the latter. Substantivalists who favor the muffins-in-muffin-tins analogy are *dualists*. According to dualists, material objects and spacetime are fundamentally different sorts of thing, the two are united by an occupation relation that binds them together and underwrites their collaboration in composing physical reality. Substantivalists who favor the eddies-in-water analogy are *supersubstantivalists*. Supersubstantivalists think that spacetime (or its basic parts) is the only fundamental, concrete object, and that material objects are reducible to regions of spacetime and their properties.

Supersubstantivalism has enjoyed recent interest, and can arguably also claim an impressive pedigree. For example, there are grounds for understanding Plato¹, Alexander², Descartes³, Newton⁴, Clifford⁵, and Spinoza⁶ all to be supersubstantivalists in some way or other. The contemporary interest in the view, however, seems to be the result of two things. First, supersubstantivalism seems to be more parsimonious than dualism, since it does not require a fundamental occupancy relation, and does not require two distinct categories of fundamental, concrete objects. Second, modern field theories seem to treat fields as properties of spacetime, and a literal reading of these theories yields supersubstantivalism. Against these considerations dualists typically press two sorts of argument. The first is based on the idea that co-location may be possible, and it is not obvious how a supersubstantivalist can

¹*Cf.* Graves (1972).

²*Cf.* Thomas (2013).

³*Cf.* Skow (2005, p. 55), and Schaffer (2009, p. 133)

⁴*Cf.* Thomas (2013), and Skow (2005).

⁵*Cf.* Graves (1972).

⁶*Cf.* Bennett (1984).

accommodate co-location. The second sort of argument is based on the idea that supersubstantivalism is counterintuitive.

This chapter will proceed as follows. In §4.2 I introduce some concepts that we will need as we proceed. In §4.3 I consider a variety of ways to formulate supersubstantivalism and consider some of the advantages and disadvantages of each formulation. Then, in §4.4, I consider the alleged advantages of supersubstantivalism over dualism. I show that field theory turns out not to supply a big advantage to the supersubstantivalist since a dualist can offer a nice reformulation of her view with which it is compatible. I then show that it is not clear that supersubstantivalism is more parsimonious than dualism. In §4.5, I consider some of the alleged problems for supersubstantivalism. I concluded that supersubstantivalists have a variety of interesting and compelling responses to arguments from common sense. I also argue that, while there is indeed real pressure to accommodate the possibility of co-location, this is an accommodation that a supersubstantivalist can make.

The main conclusion I reach is somewhat deflationary. The standard list of pros and cons appealed to in the debate between dualists and supersubstantivalists does not strongly decide between the two views. Rather, it brings into focus a particular version of dualism – *bundle-theoretic dualism* – and a particular version of supersubstantivalism – *compresence supersubstantivalism*. Moreover, on inspection, the distinction between bundle-theoretic dualism and compresence supersubstantivalism seems to be merely semantic, rather than robustly metaphysical.

4.2 Preliminaries

4.2.1 Spacetime

In general, I will suppose that the basic parts of spacetime are spacetime points. The main question will be whether there are, distinct from, and in addition to spacetime points, any material objects. Sometimes I will talk like there are. When I do, I

will often assume that the basic parts of material objects are material points. This allows me to apply classical mereology to both. This, in turn, allows us to characterize spacetimes as fusions of spacetime points, and material objects as fusions of material points. It also allows us to characterize *regions of spacetime* in terms of the parthood relation. If we suppose spacetime is a fusion of spacetime points Ω , then: r is a *region of spacetime* =_{def.} r is a part of Ω .

I will have occasion to challenge many of these assumptions. For example, in §7.1, I will consider the prospects of treating the fundamental constituents of reality as *neutral* entities – point-like things that are neither material, nor spatiotemporal. I will also sometimes talk about fields, where these are taken to be the material constituents of reality. The question here will be whether we ought to think of fields as occupying spacetime in the way that material points – if there are any – do, or whether we ought to think of them as properties of spacetime. In every case, the context will make it clear which background assumptions I am appealing to.

4.2.2 Natural Properties and Relations

I take there to be a primitive notion of *perfect naturalness* that applies to properties and relations. The perfectly natural properties and relations are special in the sense that they are “joint carving”: when things share a perfectly natural property, they are objectively similar, and when pairs of relata share a perfectly natural relation, the pairs are objectively similar. Given the notion of perfectly natural properties and relations, we may characterize duplication. Two individuals are *duplicates iff* there is a similarity map from parts of one to parts of the other such that corresponding parts have the same perfectly natural properties and stand in the same perfectly natural relations.

The concept of duplication can be used to characterize the notion of an intrinsic property. A property P is *intrinsic iff* no two duplicates anywhere in logical space

disagree with respect to P , otherwise P is *extrinsic*. Then we may say that the *intrinsic character* of a thing is an exhaustive catalogue of that thing's intrinsic properties. Next we can say that a relation R is *internal* iff R supervenes on the intrinsic character of its relata. Similarity is the paradigm: whether A and B are similar depends only on the intrinsic characters of A and B , it does not depend on their locations, or their arrangement among other things. We may say that a relation R is *external* iff it does not supervene on the intrinsic character of its relata taken individually, but it does supervene on the intrinsic character of its relata taken together. Spatiotemporal distance relations are the paradigm: whether A and B are five feet apart does not supervene upon the intrinsic characters of A and B , but it does supervene on the intrinsic character of their fusion.⁷ Finally, any relation that is internal or external is *intrinsic*, otherwise it is *extrinsic*.⁸

4.2.3 Modality

Next, I will assume a modal framework that is roughly that of Lewis (1986). Like Lewis, I take worlds to be maximally interrelated sums. However, unlike Lewis, I take the relevant sort of interrelatedness to be underwritten by a network of external relations.⁹ I take *de dicto* modal claims to be claims about what worlds there are; and

⁷Although, note that there may be some reasons for thinking distance relations are not external. See, for example, Bricker (1993) and Maudlin (2007).

⁸The characterization I give follows Bricker (1996).

⁹Note that Lewis thinks of worlds as maximally *spatiotemporally* interrelated sums. This has the effect of making spatiotemporality an essential ingredient of reality. But there are compelling grounds for skepticism about this. One reason for skepticism comes from physics. Quantum mechanics and general relativity are famously successful physical theories; equally famous is their apparent incompatibility. Physicists' attempts to reconcile the two theories have spawned radical new proposals for the fundamental structure of the world. Among these are proposals according to which there are no natural distance relations. (See Butterfield and Isham (2001) for an overview of the various research projects in quantum gravity, including ones according to which there are no natural distance relations.) This strongly suggests that we ought to think that there are worlds at which there are no natural distance relations – indeed, ours might be one of them. See Bricker (1996) for discussion.

I take *de re* modal claims to be analyzable in terms of counterparts. This allows me to understand modally qualified claims as involving quantification either over worlds or over counterparts.

4.2.4 The Occupation Relation

I have already mentioned that dualists require an occupation relation that binds material objects to regions of spacetime. There are a variety of these relations on the market, and there does not seem to be a consensus concerning which is the best one. The idea, however, is that all of these relations are related and inter-definable. So, for example, if we settle on one of these occupation relations and treat it as a primitive of our theory of occupancy, we can then define the other occupancy relations in terms of it. This is analogous to the situation in classical mereology: although it is standard to take parthood to be a primitive and then define the other mereological notions in terms of it, one could just as well start with overlap and use it to define parthood and the other notions.

My own preference, if I were a dualist, would be to accept *exact occupation* as a primitive in my theory of occupancy. Because this is a primitive, I can't offer you an analysis, but I can say a bit about the relation. When a material object O exactly occupies a spacetime region R , then O and R are the same size, shape, and dimension. An elephant will not generally be able to exactly occupy a region that is the shape and size of a mouse; a cube will not generally be able to exactly occupy a toroidal region.¹⁰ Now, as I write, I will have in mind exact occupation as the dualists occupancy primitive, but I will just write 'occupation' to avoid unnecessary complication, and because not much in this chapter turns on the specifics of the dualist's occupancy relation.

¹⁰Why just "generally"? Because some dualists will accept all sorts of mayhem from their occupancy relation. See chapter 3 for more about this.

One last point. A supersubstantialist can understand talk of occupancy relations. He can understand, for example, the question “Is Tolstoy identical to the region he occupies?” Here it is just that for the supersubstantialist, occupancy isn’t a *fundamental* relation, as it is for a dualist. The supersubstantialist will analyze occupation in terms of the reduction thesis he applies to material objects.

4.3 Supersubstantialisms

Substantialism is the view that spacetime is a fundamental, concrete object¹¹ that does not depend for its existence on anything else.¹² *Dualistic* substantialism is substantialism in conjunction with the thesis that material entities – material objects, or material fields – belong to a *distinct* category of fundamental, concrete object, and that material objects are related to regions of spacetime by a fundamental relation of occupancy. *Supersubstantialism* is substantialism in conjunction with the thesis that material entities are reducible to spacetime and its properties. A variety of ways of cashing out ‘reducible’ gives rise to a variety of supersubstantialisms.

There is a top-level distinction among varieties of supersubstantialism that we may approach by considering the question: What sorts of fundamental properties and relations does spacetime instantiate? According to the *modest supersubstantialist*, spacetime instantiates not only the fundamental topological and geometric properties and relations typically associated with spacetime, but also fundamental properties that are paradigmatically material, like mass and charge. According to the *radical supersubstantialist*, the only fundamental properties and relations instantiated by

¹¹What if you are a mereological nihilist? The mereological nihilist maintains that there are no composite objects, so, because spacetime is the composite of certain basic parts (points, regions, etc.), there is no spacetime. Does it follow that mereological nihilism and substantialism are incompatible? Of course not. The mereological nihilist merely needs to substitute ‘basic units of spacetime’ for ‘spacetime’ in the above characterization.

¹²*C.f.*, Lehmkuhl (2015), Pooley (2012), Schaffer (2009).

spacetime are topological and geometric properties. On this view, properties like mass and charge are not fundamental. Rather, they are reducible to the geometric and topological properties of spacetime.¹³

So substantivalism comes in two variants: dualism and supersubstantivalism. Supersubstantivalism itself comes in two variants: modest and radical. A typical dualist will claim that there are material properties and there are spatiotemporal properties, and spacetime instantiates the spatiotemporal properties, while matter instantiates the material properties. The modest supersubstantivalist agrees with the dualist that there is a distinction between these sorts of properties, but claims that spacetime is capable of being the bearer of not just the spatiotemporal properties, but also the material ones. The radical supersubstantivalist thinks that, fundamentally, there is no distinction between material and spatiotemporal properties. According to him, there are really just a collection of geometric properties instantiated by spacetime, and everything is reducible to these. So the modest supersubstantivalist thinks that the dualist is mistaken about how many types of fundamental entity are required to instantiate the fundamental properties, while the radical supersubstantivalist thinks the dualist is wrong about how many types of fundamental entity are required to instantiate the fundamental properties, and that both the dualist and the modest supersubstantivalist are mistaken about what properties there are.

Most philosophical engagement with supersubstantivalism has involved modest supersubstantivalism. Lewis (1986), Sider (2001), Skow (2005), and Schaffer (2009) all defend versions of modest supersubstantivalism¹⁴ But some philosophers of physics have suggested that all the focus on modest supersubstantivalism is unfortunate. Sklar, for example, is dismissive of the view, he suggests that “[t]he identification

¹³The terms ‘radical supersubstantivalism’ and ‘modest supersubstantivalism’ are due to Skow (2005).

¹⁴Their defenses seem to vary in vigor with Lewis occupying the less vigorous side of the spectrum and Schaffer the more vigorous side. See the following sections for more on this.

of all of the material world with the structured world of spacetime [should not be] interpreted as the linguistic trick of simply replacing objects by the region of spacetime they occupy and some novel “objectifying feature” – say replacing ‘There is a desk in the (X,T) region’ by ‘The (X,T) region desks.’ (1974)” This sort of linguistic trick is supposedly something the modest supersubstantialist is guilty of. Lehmkuhl (2015) has a similarly dim view of moderate supersubstantialism. He suggests that the philosophical community would be better served by shifting the focus to radical supersubstantialism, since only radical supersubstantialism has the potential to interact fruitfully with live issues in physics. Presumably, he thinks that philosophical issues are important only to the extent that they interact fruitfully with empirical science.

4.3.1 Modest Supersubstantialisms

Modest supersubstantialism is the view that spacetime is a fundamental, concrete object that does not depend for its existence on anything else, that material entities are reducible to spacetime, and that paradigmatically material properties like charge and mass can be instantiated by spacetime. This still leaves open how to cash out ‘reducible’. There are several ways that modest supersubstantialists have proceeded here. Perhaps the most popular way has been by taking the reduction of the material to involve the identification of the material objects with regions of spacetime. Another strategy is to take material objects to be composed by parts of spacetime. Each of these varieties of supersubstantialism have further variants. I’ll consider some of these below. Later in the chapter (§7.3.2), I’ll offer a novel version of modest supersubstantialism that I think is preferable to those discussed here.

4.3.1.1 Identity Supersubstantialism

Identity supersubstantialism is the view that we ought to take the reduction of material objects to regions of spacetime to involve the identification of material ob-

jects with regions of spacetime. Now, it may seem odd to treat the *identification* of A with B as a *reduction* of A to B , since reduction is typically taken to be an asymmetric relation, while identity is symmetric. But in the present context it's not so odd, since most people who accept the view that material objects are identical to regions of spacetime, tend to treat spacetime as taking over the job of material objects, rather than the other way around. Schaffer, for example, gives the following picturesque description: “[On the dualist view] [w]hen God makes the world, she must create the receptacle, and fill it with material. Then she can pin the fundamental properties onto the material substrata that fill the receptacle... [According to the supersubstantialist alternative] [w]hen God makes the world, she need only create spacetime. Then she can pin the fundamental properties directly to spacetime. (2009, p. 133)” It's clear from Schaffer's remarks that at least his version of identity supersubstantialism has a built-in asymmetry: it's spacetime that gets the material properties, not matter that gets the spatiotemporal properties.¹⁵

Now, the identity view may be qualified in any number of ways; it does not, *per se*, say anything about *which* regions of spacetime are to be identified with material objects. One might, for example, claim that only regions that are chosen by God qualify as material objects, or only regions that involve biological processes, or only regions that instantiate consciousness, etc. However, as Schaffer notes, there are several qualifications that are especially natural ones. They are as follows:

Massy Identity Supersubstantialism: every region with non-zero mass-energy is identical to a material object.

Connected Identity Supersubstantialism: every connected region is identical to a material object.

¹⁵See also Schaffer's (2009) remarks on p. 146. He claims that his version of Unrestricted Identity Supersubstantialism entails “a one-one mapping between material objects and spacetime regions, which is *a perfect opportunity for reduction*. (My emphasis)”

Unrestricted Identity Supersubstantivalism: every region is identical to some material object.

Schaffer himself defends unrestricted identity supersubstantivalism.¹⁶ He is motivated to accept this view because it maximizes a sort of mirroring between material objects and spacetime. Consider the following principles:

Monopolization: necessarily, if a material object O exactly occupies a region R , then there is no material object distinct from O that exactly occupies R .

Materialization: necessarily, every material object exactly occupies some region of spacetime.

Exhaustion: necessarily, if a material object O exactly occupies a region R , then there is no region distinct from R that O exactly occupies.

As Schaffer puts it: “Materialization and exhaustion together show that every material object occupies one and only one spacetime region, and monopolization and [the thesis of unrestricted identity supersubstantivalism] together show that every spacetime region is occupied by one and only one material object. That yields a one-one mapping between material objects and spacetime regions, which is a perfect opportunity for reduction. (2009, p. 146)”

But there are some problems here that we need to flag. First, the problem with monopolization and materialization is that they are false (just a trifling problem!). Begin with monopolization. Monopolization rules out the possibility of co-locating material objects. But, even leaving aside familiar issues having to do with statues and

¹⁶Note that Schaffer uses the term ‘monistic substantivalism’ and ‘dualistic substantivalism’ rather than ‘supersubstantivalism’. I confess that I much prefer Schaffer’s terminology, but he seems to be just about the only one to use it; ‘supersubstantivalism’ seems have established itself as the standard nomenclature.

lumps, there are strong motivations for regarding cases of co-location as genuinely possible. I defer further discussion of these motivations to §4.5.2.

Now consider materialization. Our two best physical theories – quantum mechanics and general relativity – are famously incompatible, and it is not clear what a reconciliation will look like. Among the live options for a reconciliation are accounts according to which there are no fundamental distance relations. On this sort of view, material objects exist, and they occupy parts of the fundamental structure of the world, but because the structure is not spatiotemporal, they do not occupy parts of spacetime. Because this is a live option among physicists – something that may, for all we know, be true of our world – we ought surely countenance its possibility. But this means materialization is false.¹⁷

Next, consider the phenomenon of empty regions. *Prima facie*, it seems that a supersubstantialist who accepts an identity view would analyze emptiness as follows: a region of spacetime R is *empty* iff there is no material object that is identical to R . But there can be no such region according to unrestricted identity supersubstantialism (although there can be according to the other identity views). The problem here is that it is supposed to be a decisive advantage of substantialism over relationalism that it be in a position to explain the possibility of empty regions – a possibility that is surely one we ought to accommodate. But it seems that Schaffer’s unrestricted supersubstantialism is forced to concede this advantage.

Schaffer is aware of this problem, and he does attempt to mitigate against the conclusion I draw by claiming that “empty regions are impossible in one sense but possible in another. Empty regions are impossible – at least within field theory – in the sense in which an empty region is one that lacks field values. . . . Empty regions are possible in the sense in which an empty region is one with null values for certain

¹⁷*C.f.*, Lehmkuhl (2015).

fields. (2009, p. 145.)” But, presumably, field theories, if true, are contingently true. So, for example, there are field-free, Newtonian particle worlds. But at worlds like this we cannot appeal to an analysis of ‘empty region’ in terms of null field values, so it seems we’d be stuck back with the analysis of ‘empty region’ I gave in the preceding paragraph. But again, according to unrestricted identity supersubstantivalism, this sort of emptiness is impossible. That means that matter at all field-free Newtonian worlds forms a plenum. But that’s incredible!

This actually is part of a more general issue that sometimes surfaces in debates between dualists and supersubstantivalists. The question of dualism as opposed to supersubstantivalism is often cast as a question that faces someone who has already decided in favor of substantivalism. The problem is that in the haggling between these two views, one may concede something to the other that ends up shifting the balance back towards relationalism. Schaffer’s unrestricted identity supersubstantivalism is an example of this phenomenon: the move to unrestricted identity substantivalism requires that we take empty regions to be impossible (at least, impossible within the sphere of Newtonian worlds). And this seems to undermine one of the advantages of substantivalism over relationalism.

4.3.1.2 Constitution Supersubstantivalism

According to another version of modest supersubstantivalism, material objects are *constituted* by spacetime regions, though they are not identical to them. Thomas (2013) argues that this was actually a position that Newton held early in his career.

One advantage the constitution view has over the identity view is that it makes room for a material object to differ from its associated regions of spacetime with respect to non-categorical properties. Compare with the case of the statue and the lump of clay from which it was made. Suppose the clay came into being at some time t , and that it was subsequently fashioned into a statue of David at t_1 . At t_1 it is

tempting to identify the statue and the lump of clay: if you were to pick up the statue, it would be difficult for you to look at it and claim that you held more than one thing in your hands. But other considerations press in the opposite direction. The statue and the lump differ with respect to temporal properties: the lump exists at both t and $t1$, while the statue exists at $t1$ but not t ; they differ with respect to modal properties: the lump, but not the statue, has the property *being possibly spherical*; they differ in kind: the statue is a statue, while the lump is not a statue; they differ with respect to persistence conditions: the lump, but not the statue, could persist through a squashing.¹⁸ These sorts of considerations have lead many philosophers to deny that constitution is identity.

As I said, applying a constitution view to supersubstantivalism allows a supersubstantialist to claim that material objects and their associated regions differ with respect to non-categorical properties. The big advantage of this is that it gives the supersubstantialist a way to avoid one of the main arguments against supersubstantivalism: the argument from *de re* modal difference. I will discuss this argument and the constitution supersubstantialist's reply to it in §8.2. Perhaps an obvious disadvantage to constitution supersubstantivalism is that whereas the idea that a statue is composed of a lump of clay is intuitively accessible, the idea that a tree, person, or electron is composed by spacetime is intuitively opaque.

4.3.2 Radical Supersubstantivalism

Radical supersubstantivalism is the view that spacetime is a fundamental, concrete object, and that material properties are reducible to spatiotemporal properties.¹⁹

¹⁸See Gibbard (1975), and Wasserman (2013).

¹⁹See Lehmkuhl's (2009) dissertation for excellent and detailed discussion of radical supersubstantivalism.

There are many ways that this reduction might proceed, but the underlying intuition on its behalf is nicely summarized by Clifford (1876):

I hold in fact

1. That small portions of space are in fact of a nature analogous to little hills on a surface which is on the average flat; namely, that the ordinary laws of geometry are not valid in them.
2. That this property of being curved or distorted is continually being passed on from one portion of space to another after the manner of a wave.
3. That this variation of the curvature of space is what really happens in that phenomenon which we call the motion of matter, whether ponderable or etherial.
4. That in the physical world nothing else takes place than this variation, subject (possibly) to the law of continuity.

This general thought was the inspiration for Wheeler and Misner's geometrodynamics. Geometrodynamics is a reconstruction of GR as a theory of 3-spaces evolving in time. Material properties, according to this view, are reduced to various of the geometric properties of space (not spacetime). In particular, different sorts of material properties are taken to correspond to different structural features of space. Thus we have Wheeler echoing Clifford when he asks whether "curved empty geometry [is] a kind of magic building material out of which everything in the physical world is made: (1) slow curvature in one region of space describes a gravitational field; (2) a rippled geometry with a different type of curvature somewhere else describes an electromagnetic field; (3) a knotted-up region of high curvature describes a concentration of charge and mass-energy that moves like a particle? (Wheeler (1962), p. 361)"

Geometrostatics was agreed to have been a failure on the grounds that it lacks a well-posed initial value problem, and fails to predict fermions.²⁰ Nevertheless, the radical supersubstantialist program has carried on in physics under various banners. For example, classical Kaluza-Klein theory unifies the gravitational and electromagnetic fields in the context of a five-dimensional manifold. Modern versions of the Kaluza-Klein theory have successfully reduced further material properties to the geometric properties of a five-dimensional manifold, and so appear to be promising habitats for radical supersubstantialism.²¹ There are also ongoing projects in loop quantum gravity that seem to offer respite for radical supersubstantialism. According to loop quantum gravity approaches, space (which is distinct from time, and evolves through it) is discrete, and the properties of the fundamental particles of the Standard Model are reduced to states of these fundamental, discrete units of space.

But it is important to observe that while the radical supersubstantialist program may perhaps be substantially developed in the context of some future physics, all current attempts are incomplete, and all past attempts have stalled. It thus remains an unfulfilled ambition. So the right attitude to take, I think, is to remain hopeful of the prospects for radical supersubstantialism, but to continue to examine in greater depth the case for or against modest supersubstantialism.

4.4 Advantages of Supersubstantialism

There are two considerations that are often touted as important advantages of supersubstantialism over dualism. The first comes from physics: most relatively modern physical theories are field theories, and field theories seem to treat physical fields as *properties* of spacetime, rather than as things that occupy spacetime. The

²⁰See Stachel (1974) about the the fermion problem, and Schafner and Cohen (1974) for the problem with the initial value problem. *C.f.* Lehmkuhl (2009).

²¹See, for example Wesson (2007).

second consideration is that supersubstantivalism is more parsimonious than dualism: where the dualist requires two types of fundamental entity – spacetime and material objects – the monist requires only one – spacetime; and where the dualist requires the instantiation relation to pin properties to material objects, and then an occupation relation to situate material objects in regions of spacetime, the monist requires only instantiation – the properties stick directly to the spacetime. In this part of this chapter, I will consider these two sorts of considerations and see how well they stand up to scrutiny.

4.4.1 Field Theories

There are lots of different sorts of fields, and there are lots of different sorts of field theories, but the complexities need not detain us. The important point is that *material fields* are one sort of field. The stress-energy field in GR is a good example of one of these. So let's focus on the difference between dualist and supersubstantivalist accounts of this sort of field. For the rest of this section I'll just use 'field' to refer to the sorts of material fields at issue here.

So, do field theories really decide in favor of supersubstantivalism? I think that they do not. But they do introduce an interesting twist: they turn what appeared to be a substantive metaphysical debate between dualists and supersubstantivalists into a semantic issue. Here is why.

First, field theories really do not offer much that the supersubstantivalist can lord over the dualist. While it is true that it is tempting to treat fields as properties of spacetime, there is no reason we can't think of fields as *occupying* spacetime, rather than being instantiated by it. Nothing about the physics could possibly bear on whether to prefer one over the other relation. So, let's suppose that fields occupy points of spacetime. How, exactly, would that work? Well, it turns out that it will depend on what sort of account of properties you accept. The basic idea will involve

treating field values as material point objects that occupy points of spacetime. But the problem will be to account for the “materiality” of these point objects in a way that doesn’t smack of redundancy.

To see the problem, begin by supposing that the dualist accepts a substance-attribute theory of material point objects. According to a substance-attribute theory, material objects are analyzed in terms of an element of particularity – a *thin particular* – that instantiates some property or collection of properties. The thin particular with its full complement of properties is a *thick particular*. So perhaps material point objects are to be analyzed as thick particulars that are composed by point-sized material thin particulars together with a property corresponding to a value of a field.

For example, consider a particular value of the stress-energy field, call it v . Suppose we think of v as a property. The dualist might analyze the attribution of v to a point of spacetime in terms of v ’s instantiation by a material thin particular. The combination of this thin particular together with the property v that it instantiates is a thick particular. Material point objects are these sorts of thick particulars, and these sorts of thick particulars are what bear the occupancy relation to points of spacetime. It’s like a layer cake: first you take the field-value *qua* property, and you have this instantiated by a material thin particular, then you take the resulting thick particular and you have this occupy a point of spacetime – the instantiation and occupation relations are like frosting, and property (field value), the material thin particular, and the spacetime point are like the layers of cake.

The problem with this is that it seems awfully bloated. This makes it tempting to find ways to get rid of layers of cake and frosting. The most obvious way would be to say that the only element of materiality that is needed is supplied by spacetime: just take the field value v to be a property, but have it bear the instantiation relation directly to the point of spacetime rather than introduce all these intervening

layers of thin particulars and occupancy relations. This of course is to go over to supersubstantivalism. But this isn't the only way to go.

A dualist who is a bundle theorist could offer an account that is very similar to that of the supersubstantivalist. A bundle theorist thinks of material objects as mere bundles of properties (either universals or tropes), rather than as thin particulars together with collections of properties. So a dualist who is a bundle theorist may think of the field value v itself as a material point object, and can then take this material point object to bear the occupancy relation to a point of spacetime. This gets rid of a lot of the bloat that a substance-attribute theorist seems committed to: there is no need for thin particulars, and there is no need for the instantiation relation to bind v to a thin particular.

There are two ways to read all of this: a weaker way, and a stronger way. According to the weaker way, we have found that field theories put some pressure on a dualist to be a bundle theorist rather than a substance-attribute theorist, and that if this move is made, the bundle theoretic dualist and the supersubstantivalist seem to come out pretty neck-and-neck. The supersubstantivalist thinks that field values are properties that are instantiated by spacetime points, and bundle theoretic dualists think that field values are material point objects (that happen to also just be properties since bundle theorists take material objects to just be bundle of properties) that bear an occupancy relation to spacetime points. There may still be a bit of a tilt in favor of supersubstantivalism here since supersubstantivalists don't need an occupancy relation. They get by with instantiation alone, which, presumably is a freebie on the grounds that it, as a part of first-order logic, is part of everyone's ideology.

But there is also a stronger, more interesting way to read all of this. The difference between bundle theoretic dualism and supersubstantivalism seems to come down to whether we ought to call the relation that binds fields to spacetime 'occupation' or 'instantiation', and whether to call the fields themselves 'material properties' or

‘material objects.’ But this seems to me like a mere semantic issue. If that’s right, then maybe the *two* most plausible version of substantivalism in the context of field theories are, fundamentally, *one*.

4.4.2 Parsimony

Supersubstantivalism is supposed to be more parsimonious than dualism.²² Consider Sider:

“There is considerable pressure to [identify material with regions of space-time], for otherwise we seem to gratuitously add a category of entities to our ontology. All the properties apparently had by an occupant of space-time can be understood as being instantiated by the region of spacetime itself. The identification of spatiotemporal objects with the regions is just crying out to be made. (2001, p. 109-110)”

Schaffer adopts a similar attitude:

“Where the [supersubstantivalist] of any stripe has one sort of substance . . . the dualistic substantivalist needs two sorts of substances (plus the containment relation to link them). What is the necessity for the dualistic doubling of substance types? [. . .] It is as if the dualist has not just pins (properties) and pincushions (material objects), but also a sewing table (the spacetime manifold) on which the pincushions sit. But once one has the sewing table, the pincushions seem superfluous. Why not push the pins directly into the table? (2009, p. 137-138)”

Note that the supersubstantivalist’s alleged conservatism is really twofold: he has a leaner ontology, since he gets by with one type of fundamental, concrete object;

²²*C.f.*, Lewis (1986, p. 76), who “oppose[s] the dualist conception as uneconomical”, and Quine who is critical of “a redundant ontology containing both physical objects and place-times” (1981, p. 17).

he also has a leaner ideology since he does not require a fundamental occupancy relation. I think there are two things to think about in relation to this observation. The first is whether parsimony really counts for much, the second is whether there is some serious theoretical cost associated with this sort of parsimony. Parsimony is one among the theoretical virtues. If you maintain that it is important to respect the theoretical virtues, and gloat about how parsimonious your theory is, you had better be prepared to respond if the lump pops up in some other part of the rug.

Is supersubstantivalism really more parsimonious than dualism? I think it is clear from what we saw in §4.4.1 that this really depends on what version of dualism you are considering. Clearly, supersubstantivalism is more ontologically parsimonious than substance-attribute dualism, although it seems to be about on a par with bundle theoretic dualism. But what about the alleged advantage of ideological parsimony that supersubstantivalists are supposed to enjoy over dualists? Is this a real or imagined advantage?

Well, the idea is supposed to be that whereas the dualist ideology involves both instantiation and occupation – material objects are related to properties by a relation of instantiation, and to regions of spacetime by a relation of occupation – the supersubstantivalist ideology gets by with instantiation alone. But the bundle-theoretic dualist doesn't mention instantiation in their analysis of the relationship between material objects and regions of spacetime. They say that material objects just are bundles of properties, and that these bundles occupy spacetime. But here is a worry: what about the properties and relations that spacetime is supposed to instantiate, like metric and topological properties and relations? Presumably the dualist requires an instantiation relation for those properties and relations. But does she? What if the dualist just gets rid of instantiation altogether? There are two ways to go about this sort of thing.

One way involves treating regions of spacetime as bundles of properties, and then claiming that the spatiotemporal bundles get together with the material bundles via an occupancy relation. There's no need for instantiation here. But (and I'll say a little more about this later), this might make a substantivalist feel like she is sliding uncomfortably close to relationalism. According to the sort of view we're considering, while it's true that there are such things as spacetime regions, it is also true that reality is made out of nothing but properties and relations, since material objects and spacetime regions are nothing but bundles of properties standing in various sorts of relations.

Another way to pull off this sort of trick is to maintain that spacetime is a genuine substance (not just bundles of properties and relations) but that all of its properties and relations are mediated by an occupancy relation. That is, it's not just material objects that occupy spacetime, but regular properties and relations occupy spacetime, too. So, to be clear, the idea here is this: (1) spacetime is a substance; (2) there is no such fundamental relation as *instantiation*; (3) so the properties and relations of spacetime are not *instantiated* by spacetime, but *occupy* it; (4) material objects are also related to spacetime by an occupancy relation.²³ There is yet another way to pull off this trick. We could revert to a substance-attribute theory, and say that the relationship between thin particulars and their properties is *occupation* rather than *instantiation*. On this sort of view, properties and relations occupy material objects, and material objects occupy regions of spacetime. Again, no need to mention instantiation. Although, I've already said why I don't like this third view: it is bloated.²⁴

²³Here I would not deny that there are properties in the abundant sense, and that an instantiation relation holds between them and particulars. Rather, I would deny that the relation that sparse, immanent properties have to spacetime is anything other than occupation.

²⁴Cowling (2014) develops the idea of replacing *instantiation* with *occupation*.

So, there seem to be ways for a dualist to match not just the ontological parsimony enjoyed by the supersubstantialist, but also the ideological economy. I think the two best ways are the bundle-theoretic account that does away with instantiation and treats spacetime regions as bundles, and the account that takes material objects to be bundles of properties, and treats the properties and relations of substantial spacetime as being related to spacetime by an occupancy, rather than an instantiation, relation. If pressed, I'd opt for the former, despite, as I noted earlier, being suspiciously relational.

4.5 Problems for Supersubstantialism

There are two sorts of problems for supersubstantialism that have received some attention in the literature. The first has to do with the fact that the view seems counterintuitive, and the second has to do with the possibility of co-location. In this part of the paper I will discuss both of these problems in detail. What we will find is that these “problems” are actually salutary to the supersubstantialist: they force him to clarify his position.

4.5.1 Counterintuitiveness

Supersubstantialism is supposed to be counterintuitive. Schaffer, for example, notes that “[d]ualistic substantialism is by far the more natural and popular view. . . [supersubstantialism], by contrast, is a revisionary and unpopular view.” (2009, p. 133) Lehmkuhl observes that “[supersubstantialism] may seem logically possible but slightly non-common-sensical.” (2015)²⁵ Hawthorne makes this point in significantly more detail:

²⁵*C.f.* Skow (2005), and Sider (2001, p. 111).

“One objection one might make against [supersubstantivalism] would be to point out that its identity claims – say that a certain poached egg is identical to a region of space-time – seem to directly violate common sense. More subtly, one might offer an objection based on modal considerations that runs as follows. It is natural to think of space-time regions as having their spatiotemporal profile essentially, but of their occupants as having their spatiotemporal profile only accidentally. For example, it seems obvious that the space-time region I occupy could not have been shorter in temporal extent, but that I could. If so, and if we accept the principle concerning identity known as Leibniz’ law (that if x has some characteristic that y lacks, then x is not identical to y), then I am not identical to the space-time region that I occupy.

A third argument can be made on similar grounds employing non-modal properties. Many of the [predicates] that seem true of me sound very odd when combined with noun phrases that explicitly pick out space-time regions that I occupy. For example:

The space-time region that I occupy walked to the fish and chip shop last night.

Here the problem is not so much that the claim conflicts with a previous explicit commitment about the nature of space-time regions (say, that space-time regions do not walk to fish and chip shops); rather, it’s that predicates like ‘walk’ sound very odd in combination with noun phrases that are explicitly about spacetime regions, whereas they don’t sound odd at all in combination with noun phrases concerning their occupants. (2008, p. 265)”

I don't really think there are three arguments here. I think there are two. One of these is an argument to the effect that common sense does not regard the occupants of spacetime as identical to the regions they occupy: on the one hand, common sense straightforwardly rejects identity claims involving spacetime regions and their occupants, and on the other, it recoils from sentences that involve predicating certain properties of spacetime regions. In both cases the key premise would be the same: if common sense rejects or recoils from X , then there is some reason to think that X is false.

The second argument involves modal differences between spacetime regions and their occupants. This argument is related to the first – in so far as these modal differences are generated by our common sense intuitions about regions and their occupants – but it is also quite importantly different from it, since it requires a different sort of response. So, what we'll do now is formulate the two arguments in a concise way, and then consider how a supersubstantialist ought to respond. One of the interesting things to note is that there will be different sorts of responses available to different sorts of supersubstantialists.

4.5.2 The Argument From Common Sense

We will begin with the argument from common sense. This argument encapsulates Hawthorne's first and third complaints from the passage above.

The Argument From Common Sense

- P1. Supersubstantialism entails that the author of War and Peace is a spacetime region.
- P2. Common sense tell us that the author of War and Peace is not a spacetime region.
- P3. If common sense tells us that P, then P.

C1. So, the author of War and Peace is not a spacetime region.

C2. So, supersubstantivalism is false.

I admit that I am initially tempted to dismiss this as an unimportant argument on the grounds that we are under no pressure to accept (P3). But I think this temptation must be resisted. Part of what motivates supersubstantivalism is a respect for the theoretical virtues. Typically these are taken to involve such properties as explanatoriness, parsimony, and intuitiveness. Because this framework is implicated in motivating supersubstantivalism, we can't very well dismiss it when it fails to suit us. So we ought to have something fairly substantial to say in reply to the charge that supersubstantivalism is counterintuitive. As it turns out, I think there is quite a lot to say. Indeed, some other philosophers have attempted to offer replies on behalf of supersubstantivalism to the argument from counterintuitiveness. However, it seems that no one has recognized (1) how *much* there is to say, and (2) that what sort of thing is the right thing to say depends on whether one wishes to defend radical or modest supersubstantivalism.

4.5.2.1 Modest Supersubstantivalism and Neutral Monism

Let us formulate neutral monism like this. *Neutral monism* is the view that, fundamentally, reality is composed of one sort of neutral entity, but that there may be other non-fundamental, non-neutral entities. This view is typically understood in the context of the philosophy of mind as one concerning the relationship between the mental and the physical. Neutral monists, like Spinoza, James, Mach, and Russell have seen neutral monism as a *tertium quid* in the debate between materialists and idealists. The idea, in outline, is that both mental and physical phenomena are non-fundamental, but are reducible to some basic entities that are neither mental, nor physical, but neutral. These neutral entities all belong to a single fundamental ontological category.

Not many people are neutral monists. But that is because it is difficult to properly deploy in the context of the distinction between the mental and the physical. Might it be easier to deploy in the context of the distinction between spacetime and matter? What if we regard the fundamental constituents of reality as points, but think of them as neither fundamentally material, nor as fundamentally spatiotemporal, but rather as *neutral*. These neutral points are nevertheless spatiotemporal in a non-fundamental sense. They are rightly regarded as spatiotemporal points when we think of them as things that instantiate a certain subset of their properties and relations – i.e., the spatiotemporal properties and relations, like metric, topological, affine, etc. And they are rightly regarded as material points when we think of them as things that instantiate another subset of their properties and relations – i.e., the material properties and relations, like mass, charge, and so on.

What this strategy allows us to do is claim, for example, that it is not literally true that the author of War and Peace is a region of spacetime, since, literally, the author of War and Peace is an amalgamation of properties that we do not ordinarily attribute to spacetime. This means that the supersubstantialist is not committed to the truth of the sorts of counterintuitive claims he is typically accused of being committed to. This undermines the claim that supersubstantialism is counterintuitive. Note however, that this sort of strategy is most appropriate to the modest supersubstantialist, since it seems to require that there be a legitimate distinction between material and spatiotemporal properties and relations – a distinction that the radical supersubstantialist denies.

4.5.2.2 Radical Supersubstantialism and the *A Posteriori*

But there is another avenue of response available to the radical supersubstantialist. To begin, consider Lehmkuhl's (2015) characterization of the distinction between modest and radical supersubstantialism:

“[M]odest super-substantivalism is not a scientific research programme. It is not a stance that could motivate research in physics, or serve as guiding principle for such research. Modest super-substantivalism is a purely metaphysical standpoint that can be taken quite independently from the physical theory we find to be true, and it is motivated by purely philosophical advantages. . . . That is not bad in itself. But it cannot be denied that a philosophical standpoint like radical super-substantivalism that can be fruitful for physics, motivate it and in turn be questioned by it, is a very desirable thing. . . . This is what radically super-substantivalist positions offer: they are programmes that pose a real challenge to physics, offering fruitful heuristics for scientific research, and can in turn be challenged by it.”

Modest supersubstantivalism, if true, is true *a priori*; arguments for it are either true or false no matter how the world is. The truth of radical supersubstantivalism, however, is an *a posteriori* matter. Either radical supersubstantivalism is true, or it is false, but to tell, we’d have to go look at what the world is actually like. We’d have to go look to see if properties like mass and charge are in fact reducible to the sorts of properties and relations that spacetime actually instantiates. This is an important difference, and it bears on the charge that supersubstantivalism is counterintuitive.

Consider Skow’s following observation:

“[One] might oppose supersubstantivalism because it sounds odd to say that I bought my coffee from a region of spacetime. But, in general, the demise of ordinary language philosophy has taught us to be suspicious of arguments that start with it would sound odd to say....’ And in particular, history should make us suspicious: there was a time when it sounded odd to say that I bought my coffee from a swarm of elementary particles. (2005, p. 58-59)”

The thought here seems to be that the oddness of phrases like ‘I bought my coffee from a region of spacetime’ derives from our ignorance of the actual state of affairs. Just as ‘I bought my coffee from a swarm of elementary particles’ sounded odd before we came to find that things like people are actually swarms of elementary particles, ‘I bought my coffee from a region of spacetime’ will sound odd to the ear that is ignorant of the fact that things like people are actually regions of spacetime. But just as coming to find that things like people are actually swarms of elementary particles should mitigate against the oddness of the phrase ‘I bought my coffee from a swarm of elementary particles’, so too should coming to find that people are actually regions of spacetime mitigate against the oddness of the phrase ‘I bought my coffee from a region of spacetime’.

However, it is important to observe what the actual force of this line of reasoning is. All this talk of ‘coming to find that such-and-such’ shows us that this line of reasoning is only appropriate to an *a posteriori* thesis; it is not appropriate to an *a priori* one. So while Skow’s point is inappropriately deployed on behalf of modest supersubstantivalism, it is just the ticket for the radical supersubstantivalist.

4.5.3 The Argument From *De Re* Modal Difference

Now we’ll consider an argument against supersubstantivalism based on the intuition that there are some spatiotemporal properties that are contingent to material objects but essential to the regions they occupy.

The Argument From *De Re* Modal Difference

- P1. If there is some property P , such that Tolstoy and the region that Tolstoy occupies differ with respect to P , then supersubstantivalism is false.
- P2. The region that Tolstoy occupies has, while Tolstoy lacks, the property *being necessarily 82 years in temporal extent*.
- C. So, supersubstantivalism is false.

Once again, the correct response to the argument depends on what sort of supersubstantivalism one accepts. The relevant distinction here is between two versions of modest supersubstantivalism: the constitution view or the identity view.

4.5.3.1 The Constitution View

The constitution view takes material objects – like Tolstoy – to be constituted by regions of spacetime. The constitution relation, unlike identity, is supposed to allow for a thing and the stuff it is constituted by to share all of their categorical properties – like their shape, size, mass, charge, and mereology structure – and yet differ with respect to their non-categorical properties – like their temporal, or modal properties. It seems to follow that a constitution supersubstantivalist is in a position to reject (P1): a region of spacetime and the material object that it constitutes may differ with respect to the modal property *being necessarily 82 years in temporal extent*.

But the problem I have with constitution supersubstantivalism is the same problem that I have with constitution views more generally. It is one thing to stipulate that there is a relation that allows its relata to share categorical properties while differing with respect to non-categorical ones, and it is another thing to explain what it is in virtue of which things may share categorical properties while differing with respect to non-categorical ones. The latter would be satisfying; the former, which is the strategy of constitution theorist, is simply mysterious.

4.5.3.2 The Identity View

There are at least two ways for an identity supersubstantialist to respond to the argument from *de re* modal difference. One way is to be a counterpart theorist, the other way is to reject metric essentialism.

Counterpart Theory and Identity Supersubstantialism Proponents of an identity view (any of the versions we discussed) could accept counterpart theory as an account of *de re* modality.²⁶ Modality *de re* concerns the relationship between individuals and certain of their properties. Consider me. I am 6'2". But my being 6'2" seems contingent. It seems like I could easily have been a little taller or a little shorter: maybe if I'd eaten more broccoli, I'd have been taller; maybe if I'd eaten less, I'd have been shorter. I am also human. But perhaps my being human is an essential property. Perhaps I could not exist without being human. So although I'm 6'2", I am possibly 6'4", or 5'11", or what have you, but I am essentially human. Counterpart theory is an account of this sort of modality.

According to counterpart theory, I am possibly taller or shorter, not because I literally exist in some other possibility wherein I am either taller or shorter, but because I bear a counterpart relation to an individual that is either taller or shorter. Similarly, if I am essentially human, all of my counterparts are human. Whether something is one of my counterparts, and thereby represents a possibility for me, will depend upon qualitative similarities between me and that thing. But qualitative similarity is a somewhat loose notion, one that seems to be sensitive to context. The way to factor this in is to take the possibilities for me (and other individuals) to be sensitive to context: different contexts will make salient different respects of similarity; different respects of similarity will correspond to different counterpart relations; and different counterpart relations will underwrite different possibilities.

²⁶*C.f.*, Skow (2005) p. 70-71, and Schaffer (2009) p. 145.

So in a sense, what is possible for any particular thing is going to depend upon certain features of context, like what sorts of background assumptions one holds fixed. For example, if we are discussing physiology, and considering whether or not I might have been 10 feet tall, or had blue eyes, we are probably holding fixed that I am human, so all of the relevant counterparts of me are human. But let the context change. I am playing with toy dinosaurs with my kids and we are talking about what sorts of dinosaurs they would be. Here we'll hold different sorts of features fixed – for example, certain of Ella's personality traits (like her speed and loud vocalizations: she's a velociraptor), and certain of Ben's (like his penchant for stomping around, yelling, and just causing general mayhem: he's a T-rex) – and let their being human vary.

So now we can apply this to our case. When we think of a spacetime region as Tolstoy we invoke a context where temporal extent seems to be contingent. This is simply an artifact of our standard way of thinking of people as being such that they could possibly have lived a little longer, or died a little sooner. But when we think of the region as a region of spacetime, we invoke a different context, one where temporal extent seems to be necessary. This, presumably, is also an artifact of our standard way of thinking of a region of spacetime as a thing that is individuated by its extent. So although in both cases we're talking about the very same entity, one conceptualization of it suggests one set of possibilities, while another suggests another set.

Denying Metric Essentialism Metric essentialism is the view that spacetime points have their metric relations essentially. This is a position that is defended by Maudlin (1988) in the context of the hole argument against substantivalism. It is not a popular view among philosophers, and it is rightly eschewed. But, apparently, if we have the intuition that spacetime regions have their temporal extent essentially, it is because we are somehow motivated by metric essentialism. I really don't think

that we have any strong or clear pre-theoretic intuitions about things like spacetime regions, at least not any that might bear on the question of whether they do or don't have their temporal extent essentially. So it is probably just confusion that leads to something like (P3) in the first place. But there are further reasons to reject it. Chief among these it that general relativity and metric essentialism seem at odds. Recall that according to general relativity, matter covaries with the metric. One consequence of this is that if you have an unoccupied region, and you plunk a mass down in it, you thereby change its metric structure. So if you think that metric relations are essential to spacetime, then (at least in the context of general relativity) you also have to think that the distribution of masses is essential to spacetime.

4.5.4 Co-Location

Co-location occurs when multiple distinct material objects exactly occupy the same region of spacetime. This is not like when you and I come to be in the same office at the same time, or even cram into the same small closet at the same time. This is more like when you put two different size 12 feet into the same size 12 shoe at the same time – intuitively, it seems like there just isn't room for both of them. Moreover, I am not talking about co-location in the sense that a constitution theorist might talk about the co-location of statues and lumps of clay. A constitution theorist might say something like this:

My seated body occupies at a seated-body-shaped region r . Could anything *else*, you ask, occupy r ? Well, look, there are certain identity conditions associated with me, and there are others associated with the stuff of which I am constituted right now. I, for example, do not cease to be me upon the gain and loss of parts. I am still the same person before and after I cut my hair, and before and after I metabolize some food. On the other hand, the stuff that makes me up right now – that is, a certain quantity

of molecules – does not survive the gain and loss of parts. If one molecule is subtracted from a sum of molecules, then that sum becomes a different sum. So because I and the stuff that constitutes me have different identity conditions, we are not identical. Because we are not identical, but are nevertheless such that a temporal part of me exactly occupies the same region as a temporal part of a certain sum of molecules, exact co-location is clearly possible. In fact, it's actual!

But this is importantly different from the sort of co-location I am worried about. To see this, observe that we're considering a case where a temporal part of me, call it 'tJ', is co-located with a temporal part of a certain sum of molecules, call it 'tM'. Now note that tJ and tM have exactly the same parts. The constitution theorist does not deny this, she merely denies that this makes them identical. The constitution theorist will typically attempt to explain spatiotemporal co-location with material co-location. But the sorts of cases of co-location I am worried about are cases where the objects that exactly occupy the same region do not have all their parts in common; they are cases of spatiotemporal co-location *without* material co-location. From now on I'll assume that co-location entails not just that distinct material objects exactly occupy the same region, but that the objects in question also differ with respect to some of their parts.

Now why think that this sort of co-location is a genuine possibility? There are really two sorts of reasons. The first reason is that the causal powers possessed by material objects are what prevent them from co-locating in the actual world, but there are good reasons for thinking that causal powers are not essential to the objects that have them. The second reason is that a reasonable philosophical extrapolation from contemporary quantum mechanics entails that co-location is not just a metaphysical possibility, but a physical one – it might even happen all the time right under our noses!

Now the problem. Dualists have an easy time of it when they are asked if their spacetime metaphysics is compatible with co-location. Since material objects and regions of spacetime are distinct sorts of things which are united by a fundamental relation of occupation, the dualist can simply claim that a single region can stand in this relation to multiple, distinct material objects. The supersubstantialist, on the other hand, does not have things so easy. If, as the supersubstantialist seems to require, material objects are to be analyzed in terms of regions of spacetime, it seems to follow that co-location is a metaphysical impossibility. A superficial reason for this is that regions of spacetime cannot occupy entirely distinct regions of spacetime. The deeper problem here, however, is easy to see when we consider just what co-location would involve for a supersubstantialist.

Co-location for a supersubstantialist cannot involve distinct material objects occupying the same region of spacetime. This is because the supersubstantialist does not think that there is such a relation to be born between material objects and regions of spacetime (at least not a *fundamental* relation). Co-location for a supersubstantialist will be analyzed in terms of patterns of property instantiation. Here is what I mean. Consider a possibility wherein a black and a white billiard ball are set in motion towards one another. Suppose that at the moment of contact, when the laws of the actual world would cause them to repel one another, they instead come to rest in exactly the same place. They perfectly overlap. Now a dualist will be able to explain the situation by claiming that two billiard balls, one black, and one white, have come to be located at the exact same region of spacetime. The supersubstantialist on the other hand has to say that a region of spacetime has come to be both black and white (and otherwise billiard ball like). But note that this region is not both black and white in the way that a tuxedo is – by having some black parts here, and some white parts there. This region is black *all over* and white *all over*. But, so it

seems, nothing can be both black all over and white all over. So, if co-location is possible, supersubstantialist must be false.

So, in the following several sections I'll develop two lines of argument meant to support the possibility of co-location. Then I'll consider how the supersubstantialist may reply.

4.5.4.1 The Argument from Causal Powers

Return to the billiard balls. Suppose you have two of them, and you roll them towards one another on a collision course. What will happen when they collide? They will bounce off of one another and change course. Why does that happen? Because the chemical bonds that prevail among the constituting particles of the billiard balls are pretty tough to sever, and the force of one ball rolling towards the other is insufficient to break them. If we caused the balls to collide with enough force, we could get them to break, and even to disintegrate completely, but we could not get them to pass through one another unaltered in the way that, say, a ghost might pass through a wall leaving itself and the wall unaltered. Why can't the billiard balls act like ghosts? Because the billiard balls instantiate certain properties, and (in the actual world) these properties realize certain causal powers, and *being deflected when struck* is among the effects of these causal powers. But what if those causal powers are contingent to the properties that realize them? What then?

If the causal powers associated with the relevant properties – in this case, probably the charge properties of fundamental particles – are contingent, then, arguably, there are worlds where duplicates of these billiard balls exist, but where these very same properties do not cause the billiard balls to bounce off of one another, but rather allow them to pass through one another like ghosts. But if there are worlds where

this sort of thing happens, then there are worlds where things co-locate, since the billiard balls, or parts of them, co-locate as they pass through one another.²⁷

Now, this way of arguing for the possibility of co-location hinges on the idea that properties have internal essences, that their being the property that they are goes beyond simply a propensity to bring about such-and-such causes. This is *quidditism*. Quidditism is a controversial thesis. The main alternative to it is the view that properties are individuated by, and so essentially connected to, the causal powers they confer. Call this *necessitarianism*.²⁸ Shoemaker (1980), for example, opposes quidditism, and accepts a version of necessitarianism. He claims that “[o]nly if some [necessitarian] theory of properties is true [...] can it be explained how properties are capable of engaging our knowledge, and our language, in the way they do.” We’d be shrouded in ignorance, so far as properties go, if quidditism were true.

But quidditism has its proponents. Among them, David Lewis (2009), who agrees with Shoemaker’s assessment that quidditism entails a kind of ignorance, but disagrees that this warrants its rejection. Instead, he says, it requires humility concerning our knowledge about the fundamental properties. Schaffer (2005) argues that quidditists can, after all, know something of the fundamental properties. He considers a variety of versions of quidditism and necessitarianism and ends up defending a view that he calls *quiddistic contingentism*. Quiddistic contingentism says several things. First, it says that properties are trans-world entities, rather than world-bound ones. This means that the very same property P – not, mind you, a *counterpart* of P – exists in more than one world. Second, properties can recombine with any “lawmakers”,

²⁷I anticipate some nit-picking here. Someone might say: while billiard balls *seem* pretty dense, really they are mostly empty space, their constituting particles maintaining their distances from one another in virtue of the strong and electroweak forces. To this I reply: fine; then substitute the case where we are considering a single fundamental particle. If that still won’t satisfy, then consider a non-actual case involving something like a billiard ball that is perfectly continuous – i.e., its material parts are point-sized and can be mapped one-to-one and onto to the spacetime points of the region it occupies.

²⁸*C.f.* Schaffer (2005).

where lawmakers are understood to be whatever entities one accepts in the analysis of laws of nature.²⁹

I'm not sure how many people accept quidditistic contingentism. But Schaffer does. And, considering he seems to be one of the most vocal proponents of super-substantivalism, and also one who is dismissive of the possibility of co-location *qua* challenge to supersubstantivalism, this seems quite enough. So, let us assume that quidditistic contingentism is true. It seems to follow that co-location is possible, since the only thing – at least the only things aside from a prior commitment to supersubstantivalism – that prevent co-location are the actual causal powers associated with certain of the properties of matter.

4.5.4.2 The Argument from Quantum Mechanics

Begin by distinguishing *state dependent* properties from *state independent* properties. A state dependent property is a quantity of a physical system that changes in some manner over time. Examples include things like position and momentum. State independent properties are properties of a physical system that do not change over time. Examples include things like mass and charge. So physical systems are first categorized by their state independent properties, and then their dynamics are characterized in term of their state dependent properties.

The state dependent properties of a non-relativistic quantum mechanical system may be represented either by a high-dimensional vector space over the complex numbers,³⁰ or by a wave function over configuration space.³¹ To begin, suppose that A is a quantum system composed of a single electron, that \mathbf{V} is the associated vector

²⁹For Lewis (1973, 1983) the lawmakers will be those regularities that correspond to the set of true, deductive systems with the most favorable overall balance of strength and simplicity.

³⁰A complex number is a number of the form $a + bi$, where $a, b \in \mathbb{R}$, and $i = \sqrt{-1}$

³¹There is an exception to this rule. The *spin state* of a particle is a state independent property of that particle, but it is nevertheless characterized by a vector space.

space, and that \mathbf{C} is the associated configuration space. A state of A , then, may either be represented by a state vector, $|\psi\rangle$, in \mathbf{V} or by a wave function $\Psi(r, t)$ over \mathbf{C} . Once again, the vector space \mathbf{V} , and the configuration space \mathbf{C} are merely alternative means of encoding the same information. As we'll see, it is sometimes more natural to appeal to the wave function formulation, and other times more natural to appeal to the state vector formulation.

To begin, observe that the state vector or wave function representing the state of a quantum system provides an exhaustive description of that system in terms of its state-dependent properties. Among these state-dependent properties is a property that corresponds to location. This will be our focus.

So how is it that the information regarding the location of a quantum system is generated? –Well, it turns out that the sort of information quantum mechanics generates about the locations of particles is *probabilistic*. That is, given a system A , we are told, for each region of space r , and for a time t , what the probability of finding A at r *would* be, *were* we to measure A 's location at t .

When doing this sort of calculation, the wave function formulation is most convenient. So what we need is a means of generating a probability measure from the wave function $\Psi(r, t)$. The problem is that because the configuration space over which the wave function operates is complex, the value of the wave function will be a complex number. But a complex number is not the sort of thing that can straightforwardly be taken to encode a probability value, that is, some number n , such that $n \in [0,1]$. So what we do is take the absolute square of the wave function summed over the relevant regions of space. This will give us a positive real in the interval $[0,1]$.

Now suppose we want to know where a quantum system A is, but we haven't done any measurements yet. Here is what we know: we know that *if* we measure A for exact location, we will find A 's exact location. Here is what we don't know: we don't know, prior to measuring, which region is A 's exact location, and, what's more

puzzling, we don't even know if there is *any* region that is A 's exact location. All we know is that, if we do make a measurement, it will turn up an exact location. So, while nothing prior to measuring A for its exact location will allow us to derive A 's exact location, what we can do is appeal to a formula that will give us a *probability distribution* corresponding to A 's exact location upon measurement.

Now suppose we want to consider the probability of finding A 's exact location, r , to be in some region R . In that case, to get the probability that the region r , at which A is exactly located, is a subregion of R , we integrate over the probability amplitudes of each subregion of R using the formula:

$$P_{r \in R} = \int_R |\Psi(r, t)|^2 dr \quad (4.1)$$

Note that this is essentially just summation of the probabilities associated with points in R ; the complication is induced by the fact that there are continuum many points in a region, each with probability 0, so countable additivity will not be helpful in generating a non-zero probability measure. This is why we use integration rather than summation.

4.5.4.3 Co-Locating Bosons

According to the standard model of particle physics, there are two classes of fundamental particles: fermions and bosons. Fermions are mass-carrier particles which are associated with matter. Bosons are force-carrier particles and mediate all of the fundamental forces (i.e., electromagnetism, and the strong and weak forces). What these two classes of particle have in common is that individual instances of each class are indistinguishable from one another. That is, if you take an aggregate of photons – which are a type of boson – each will be qualitatively indistinguishable from all the others. Two photons never differ with respect to any state-independent property. This is actually a general feature of quanta.

There is, however, one very important respect in which bosons and fermions differ. Fermions and bosons differ with respect to the statistics associated with composite systems of which they are members. Composite fermionic systems always occupy antisymmetric states, and composite bosonic systems always occupy symmetric states. Here is what this means. First, note that the state vector of a composite system is built up out of the state vectors of the components of that system by appealing to the following rule (given a simple two-particle case):

$$|\psi_{composite}\rangle = \int_{x,y} A(x = r_1, y = r_2) |x = r_1, y = r_2\rangle \quad (4.2)$$

Now consider a law that applies to Fermions:

$$A(x = r_1, y = r_2) = -A(x = r_2, y = r_1) \quad (4.3)$$

This is ultimately a claim about the relationship between the probabilities associated with the component vectors in a composite system. To see what this claim amounts to, consider a composite system with two component fermions, x and y . Suppose that fermion x is exactly located at region r_1 , and fermion y is exactly located at region r_2 . This state of affairs will be represented by the vector, $|x = r_1, y = r_2\rangle$, and will have an amplitude $A(x = r_1, y = r_2)$. What (4.3) says is that if the amplitude $A(x = r_1, y = r_2)$ of $|x = r_1, y = r_2\rangle$ is equal to n , then the amplitude of the vector describing the state of affairs consisting in x being exactly located at r_2 and y being exactly located at r_1 must be $-n$. But note that there is only one case where $n = -n$, and that is when $n = 0$. But now this straightforwardly carries over to probability values. If the amplitude of a vector $|\psi\rangle$ is n , then the probability that $|\psi\rangle$ describes that state of the system it corresponds to is $|n|^2$. So when the amplitude of a vector $|\psi\rangle$ is 0, that entails that the probability that $|\psi\rangle$ describes the state of the system it corresponds to is 0 as well (since $|0|^2 = 0$). The upshot here is that the antisymmetry requirement given by (4.3) entails that fermions never exactly occupy the same region.

But now consider the parallel law for bosons:

$$A(x = r_1, y = r_2) = A(x = r_2, y = r_1) \quad (4.4)$$

The importance of this law is this: if the amplitude of the vector representing the state of affairs consisting in two bosons, x and y , being arranged with x exactly located at r_1 and y exactly located at r_2 is n , then the amplitude of the vector representing the state of affairs consisting in x and y being arranged with x exactly located at r_2 and y exactly located at r_1 is n as well. To make this clear, consider just r_1 for a moment. This law says that if x is exactly located at r_1 , then y may be exactly located there as well.

The state vector of a bosonic system is still well-defined in the case where the bosons are exactly co-located. This means that among the elements in the configuration space of non-relativistic quantum mechanics are elements that involve co-located bosons.

4.5.4.4 Co-Located Bosons and Modal Plenitude

So, I've just taken pains to show that it is compatible with an orthodox, collapse version of non-relativistic quantum mechanics that bosons co-locate. So what? Quantum mechanics is simply a mathematical algorithm for making predictions about the outcomes of measurements of physical systems. It does not come packaged either with a compulsory formalism – there are non-collapse theories, there is relativistic field theory, etc – or with a compulsory ontology. So we need more than the bare observation that a particular version of quantum mechanics entails such-and-such before we are under any pressure to recognize such-and-such as a genuine possibility.

What we need is an account of modal plenitude according to which the coherency of a physical theory gives us reason to believe that each of the states that are compat-

ible with it represent genuine possibilities.³² Rather than defend such a principle in detail, I'll just consider what one might say *against* it. To begin, I am asking that we begin with a physical theory and consider whether or not it is coherent. If we have reason to think it is, then, so I say, we have reason to believe in the possibilities that are compatible with it. Rather than confront the sophisticated case of co-locating bosons directly, begin by considering the following simple case: in classical mechanics, dynamical interactions are sometimes modeled by point particles. Sometimes we want to consider what happens when point particles collide, or, from the four-dimensional perspective, what happens when then their world-lines intersect. The standard approach to modeling this sort of phenomenon involves treating the point particles as being co-located at the spacetime region where their world-lines intersect. Such cases represent straightforward case of co-location.

So, I say: if classical mechanics is coherent, then such cases give us reason to think that co-location is possible. Moreover, classical mechanics *is* coherent, so these cases do give us reason to think that co-location is possible.

But am I putting the cart before the horse here? Perhaps what we ought to do, at least in some cases, is to start by asking what possibilities are entailed by a theory, and then use these possibilities as guides to the theory's coherence. If the theory entails outlandish possibilities, then it isn't coherent; if it doesn't entail outlandish possibilities, then it is coherent. But co-location – I can imagine someone saying

³²As it turns out, Bricker (1991) offers an account of plenitude that entails the sorts of possibilities I am concerned with here. Consider Bricker's principle (B).

(B) We have warranted belief that a structure is logically possible if that structure plays, or has played, an explanatory role in our theorizing about the actual world.

Because orthodox quantum mechanics and Newtonian mechanics with point particles have played explanatory roles in our theorizing about the actual world, we are warranted in believing the structures posited by these theories to be logically possible. These theories posit structures that involve co-locating entities, so we are warranted in believing that co-located entities are logically possible. Note also that while I use the concept of coherency in my discussion above, this is compatible with Bricker's (B): a theory's having genuine explanatory power entails that it is coherent.

– isn’t coherent. So classical mechanics isn’t coherent – at least not when we take these sorts of cases literally, and not as “convenient fictions”, idealizations that make calculations easier.

This is actually a difficult issue. Here is what I think: there are standards for consistency within mathematics (which may themselves be grounded in logic), and these standards are more rigorous and less subjective than the standards set by our intuitions alone. So when we attempt to undermine, or drastically reconstrue a physical theory that meets these standards for mathematical consistency, and we do so on the grounds that this physical theory generates counterintuitive possibilities, we do so at our peril. It is legitimate to use intuition as a guide to rejecting or accepting certain views in the absence of other standards, but it is not legitimate to allow intuition to trump something like mathematical consistency. For these reasons, I think there is real pressure to accommodate the possibility of co-location.

So, we’ve seen that the orthodox, collapse version of quantum mechanics entails that co-location is possible. Moreover, I’ve given some reasons to think that, if a coherent physical theory entails that such-and-such is possible, then we have reason to think that such-and-such is possible. I assume that the orthodox, collapse version of quantum mechanics is coherent in the sense I think is relevant. So we have reason to think co-location is possible.

4.5.5 Supersubstantivalism and Co-Location

So, we seem to have two compelling reasons for thinking that co-location is a genuine possibility. How should a supersubstantivalist respond? It seems to me there are two sorts of responses available, but only one of them is good. The first response – the one I do not like – is to try to deny that co-location is a genuine possibility, the second is to embrace this possibility and see where it leads. I’ll address the first of these responses first, and the second one second.

4.5.5.1 Tollensing The Argument

Here is an argument:

The Argument from Co-location

(P1.) If co-location is possible, then supersubstantivalism is false.

(P2.) Co-location is possible.

(C.) So, supersubstantivalism is false.

This is an argument that a dualist might give against a supersubstantivalist.³³ I have just argued that (P2) is true. Schaffer (2009) disagrees. He is dismissive of the Argument From Co-Location on the grounds that (P2) is false. What is the argument against (P2)? Something like this:

The Argument Against Co-location

(P1*.) If co-location is possible, then supersubstantivalism is false.

(P2*.) Supersubstantivalism is true.

(C*.) So, co-location is not possible.

I do not think responding to the argument from co-location by appeal to the argument against co-location is a viable strategy. The primary reason for this has to do with the source of the pressure to accommodate the respective theses: the pressure to accommodate co-location comes both from physics and metaphysics, while the pressure to be supersubstantivalists (at least of the moderate version under consideration) is purely metaphysical. If the benefits to total theory of supersubstantivalism were great, and unable to be met elsewhere, I might be persuaded to accept the argument

³³Indeed, this is an argument that Hawthorne (2008) *does* give.

against co-location. But this is not the state of things. We have seen that the bundle-theoretic dualist is pretty close to equal with the supersubstantialist in the context of field theory, and so the main advantage of supersubstantialism is that it is more parsimonious. But denying a possibility that physics countenances on the grounds of parsimony alone seems to me a bad move.

4.5.5.2 Supersubstantialism and The Compresence Relation

What a supersubstantialist should do is accept something a lot like a bundle theory. Bundle theories typically require something called a *compresence* relation.³⁴ Compresence is a primitive, second-order equivalence relation on n -tuples of properties (either universals or tropes). Now, I've earlier suggested that a dualist is well-advised, in the context of field theory, to be a bundle theorist. So, let us restate the version of dualism that is suggested by field theory. First, material objects are to be analyzed in terms of bundles of properties. To get the properties bundled into discrete objects, we need the compresence relation. Next, we take the occupation relation to hold between bundles of compresent properties and regions of spacetime. Now, the supersubstantialist, to deal with the problem posed by co-location, can ape the dualist.

First, the supersubstantialist should appeal to a relation of compresence among properties in order to underwrite the distinctness of material objects. Consider a case of two 1g point masses, a and b , with perpendicular trajectories that intersect at some point p . Because a and b are distinct before and after p , we are under some real pressure to regard them as distinct at p , even though they are co-located. The dualist is in a position to make this accommodation. Indeed, the dualist who is also a bundle theorist can make the accommodation twice over! Here is what I mean. At p the bundle-theoretic dualist can claim that a and b are distinct on the grounds that p

³⁴This is Russell's terms. See Armstrong (1989, Chapter 4) for discussion.

bears the occupation relation to a , and a separate instance of the occupation relation to b . But she can also claim that, at p , a is distinct from b on the grounds that there are two instances of the compresence relation – one that holds between a and itself, and the other that holds between b and itself. The supersubstantialist who accepts compresence – let’s call him the *compresence supersubstantialist* – can appeal to this second strategy, and derive the same result.

This compresence supersubstantialist must reject the identity view: if a is identical to p and b is identical to p , then it immediately follows that a is identical to b , but that’s just what we’re trying to avoid. But what positive thesis does he accept? The basic idea is simple: material objects are to be analyzed as bundles of compresent properties instantiated by spacetime. This is sort of a hybrid of a substance-attribute theory and a bundle theory, since we have bundles of compresent properties, but we also have a substance – spacetime regions – that instantiate them. On this view material objects are not identical to regions of spacetime, rather, they are identical to regions of spacetime together with property bundles. So, the property bundles associated with a and b are distinct on the grounds that one includes the relational property *being compresent with b* while the other includes the relational property *being compresent with a* . This keep a and b distinct at p despite being co-located at p .³⁵

4.6 Conclusion

We’ve considered a variety of arguments that bear on the issue of dualism as opposed to supersubstantialism. There are several important themes that we ought to walk away with. First, radical supersubstantialism is an intriguing thesis, and it

³⁵Note that these relational properties – *being compresent with a* , etc. – are not fundamental. But perhaps only the fundamental properties are included in bundles. If that is so, then what’s to keep a and b distinct? Deciding in favor of tropes along with primitivism about their individuation will may solve this problem.

seems that some optimism for its prospects is warranted. But it remains an unfulfilled ambition, and we shouldn't hold supersubstantivalism hostage to the development of its radical variant. Second, it is far less clear than usually advertised that supersubstantivalism is less commonsensical, but more parsimonious, than dualism.

We have seen that the supersubstantivalist has a number of responses to the arguments derived from common sense. We have also seen that field theory makes substance-attribute dualism seem bloated. This should prompt a dualist to be a bundle-theoretic dualist. The contingency of causal powers, and some cases from physics, create pressure to accommodate the possibility of co-location. In order to make the accommodation, the supersubstantivalist should be a compresence supersubstantivalist. So how do these two views compare? Well, a bundle-theoretic dualist uses the term 'material object' to refer to bundles of compresent material properties, and he uses the term 'occupation' to refer to the fundamental relation that binds material objects to regions of spacetime. The compresence supersubstantivalist, on the other hand, uses the term 'instantiation' to refer to the relation born between bundles of compresent material properties and regions of spacetime, and he uses the term 'material object' to refer to regions of spacetime together with the bundles of compresent properties they instantiate. As I have said, this debate does not strike me as a substantive metaphysical debate, it is merely a debate about which are more appropriate terms for various relations and entities.

APPENDIX A

MANIFOLDS

Let's consider how to build a manifold. Begin by recalling that a topological space (X, τ) is a set of points X equipped with a characterization of which sets are open and closed, given by τ . Now we characterize some basic notions concerning means of comparing topological spaces. To do this we begin with the notion of a **map** between topological spaces. A map $\phi : M \rightarrow N$ is a correlative relationship between the space M – called the **domain** of ϕ – and the space N – called the **image** of M under ϕ .

Next, we need the notion of **composition** of maps. Begin with the map $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$. Here we may think of the map from \mathbb{R}^m to \mathbb{R}^n as one that takes the m -tuple (x^1, x^2, \dots, x^m) to the n -tuple (y^1, y^2, \dots, y^n) . This allows us to think of $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ as a collection of functions. In particular, it allows us to think of ϕ as a collection of n functions ϕ^i , one for each of the y^i , that acts on the m -tuple. So, to get y^1 , for example, you allow ϕ^1 to operate on (x^1, x^2, \dots, x^m) ; to get y^2 , you allow ϕ^2 to operate on (x^1, x^2, \dots, x^m) , and so on.

Next we consider each of these functions ϕ^i . If the P^{th} derivative of ϕ^i exists, then we say that ϕ^i is a C^P function – that is, continuous to degree P . To determine the degree of continuity of the map $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ you check to see to what degree each of the functions – the ϕ^i – are continuous. If it turns out that each of the ϕ^i are *at least* P times differentiable, then the map $\phi : \mathbb{R}^m \rightarrow \mathbb{R}^n$ is said to be P times differentiable, and if the map is infinitely differentiable, the map is said to be **smooth** or C^∞ .

The next notion we require is that of a chart. A **chart** is a pair (U, ϕ) , where U is a subset of a topological space (X, τ) and ϕ is a one-to-one map $\phi : U \rightarrow \mathbb{R}^n$,

where $\phi(U)$ is open in \mathbb{R}^n . In English: a chart is a pair, the first member of which is a specified subset of a topological space, and the second member of which is a means of correlating that subset with an open subset of a coordinate representation of Euclidean space. Now consider two charts (U_α, ϕ_α) and (U_β, ϕ_β) on an arbitrary (not necessarily Euclidean) topological space (X, τ) . Suppose that $(U_\alpha \cap U_\beta) \neq \emptyset$. Now, ϕ_α maps $(U_\alpha \cap U_\beta)$ to an open subset of \mathbb{R}^n . Let us call this open subset Φ . ϕ_β maps $(U_\alpha \cap U_\beta)$ to an open subset of \mathbb{R}^n . Let us call this open subset Ψ . Now, because U_α and U_β share a common set of points as subsets, we expect that there be a relationship between Φ and Ψ . At a raw, intuitive level, because Φ and Ψ are coordinatizations of U_α and U_β , respectively, and because U_α and U_β share a common set of points, Φ and Ψ involve two alternative descriptions of the very same points. Because they are alternative ways of describing the same things, we expect that they themselves are related in some way. It turns out that the relationship between Φ and Ψ is given by a special mapping that is the composition of ϕ_α and the inverse of ϕ_β . This gives us the notion of a **transition map**, $\tau_{\alpha,\beta} : \phi_\alpha(U_\alpha \cap U_\beta) \rightarrow \phi_\beta(U_\alpha \cap U_\beta)$, which is given by $\tau_{\alpha,\beta} = (\phi_\alpha \circ \phi_\beta^{-1})$. Such maps are defined for all and only such charts as involve subsets whose intersections are non-empty.

Now that we have the notions of continuity, of a chart, and of a transition map between charts, we are ready to characterize the notion of an atlas.

Definition A.0.1. A collection of charts $\{(U_\alpha, \phi_\alpha)\}$ is a C^∞ **atlas** if and only if:

1. $\bigcup U_\alpha = M$.
2. For any two charts (U_α, ϕ_α) and (U_β, ϕ_β) , such that $(U_\alpha \cap U_\beta) \neq \emptyset$, there exists a C^∞ transition map $(\phi_\alpha \circ \phi_\beta^{-1}) : \phi_\alpha(U_\alpha \cap U_\beta) \rightarrow \phi_\beta(U_\alpha \cap U_\beta)$.

We will say that an atlas is **maximal** if it is not possible to add another chart that is compatible with those already included. Now, finally, we are ready to give the definition of a manifold.

Definition A.0.2. A C^∞ **manifold** \mathcal{M} is a topological space equipped with a maximal atlas.

...

APPENDIX B

MODEL-DETERMINISM AND WORLD-DETERMINISM

The first step in the hole argument involves showing that GR models violate a type of model-theoretic determinism. The second step involves saddling substantivalism with an interpretation of models in terms of worlds that causes this model-theoretic determinism to be reproduced among collections of worlds. Butterfield, for example, begins his discussion of the hole argument by claiming that “[t]he basic idea of determinism is this: a spacetime theory is deterministic if any two of its models that agree on the physical state at one time agree on the physical state at any other time. (1989, p. 2)” He then goes on to claim that “the basic idea of determinism is that a single physically possible world is specified by the physical state on a certain region of spacetime. . . . Since [my] definitions are cast in terms of models . . . rather than possible worlds . . . I need to consider the relation between models and possible worlds. (1989, p. 11)”¹ As we can see, the idea is that we begin with a conception of determinism in terms of a theory’s models. Then we implement a principle of interpretation that allows us to transpose that notion to the realm of the theory’s worlds.

Let us begin with the following notion of determinism in terms of models.

Model Determinism (MD): For all GR models, \mathfrak{M} and \mathfrak{M}' , and for all cauchy surfaces Σ found in \mathfrak{M} and Σ' found in \mathfrak{M}' , if Σ **agrees** with Σ' , then \mathfrak{M} **agrees** with \mathfrak{M}' .

Then we may formulate a corresponding notion of determinism in terms of worlds.

¹See also Melia (1999, p. 656) Brighouse (1994, p. 118), and Butterfield (1988).

World Determinism (WD): For all GR worlds w and w' , and for all maximal achronal subregions S of w and S' of w' , if S **agrees** with S' , then w **agrees** with w' .

In both cases the task now is to spell out what agreement ought to consist in. We'll begin with MD.

I talked a lot about how the distinction between members of a common D-class of models is grounded in the relationship between individualistic and non-individualistic properties: while all members of a common D-class are alike in which individualistic and non-individualistic properties they possess, they differ in how those individualistic and non-individualistic properties are co-instantiated. Now because both of these varieties of property are relevant to the identity conditions for models, we might anticipate that what it means for models to agree is for them to fail to differ with respect to either of them. In other words, we might expect to cash-out agreement between models in the following way:

M-Agreement 1: models \mathfrak{M} and \mathfrak{M}' **M-agree 1** *iff:* \mathfrak{M} and \mathfrak{M}' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

Likewise, we might expect to cash-out agreement between cauchy surfaces in the following way:

C-Agreement 1: cauchy surfaces Σ and Σ' **C-agree 1** *iff:* Σ and Σ' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

Next, we parlay these notions into the following notion of determinism:

Model Determinism (MD1): GR satisfies MD1 *iff:* for all GR models, \mathfrak{M} and \mathfrak{M}' , and for all cauchy surfaces Σ found in \mathfrak{M} and Σ' found in \mathfrak{M}' : if

Σ and Σ' involve the same pattern of co-instantiation of individualistic and non-individualistic properties then, \mathfrak{M} and \mathfrak{M}' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

But GR's models will violate MD1. Here is why. Recall that the effect of a (non-trivial automorphic) diffeomorphism is a rearrangement of the pattern of co-instantiation of individualistic and non-individualistic properties across points. So if you start with a model \mathfrak{M} and you apply a diffeomorphism d to it, the model you end up with $d(\mathfrak{M})$ will involve the same individualistic and non-individualistic properties as \mathfrak{M} , but will differ with \mathfrak{M} over the pattern in which the individualistic and non-individualistic properties are co-instantiated. Now, among the sorts of diffeomorphism one could consider in this capacity are so-called "hole diffeomorphisms." A hole diffeomorphism h is a trivial diffeomorphism over some initial segment of a model, but then becomes non-trivial within some segment. The segment of the model where h is non-trivial is called 'the hole.' The overall effect of h on a model \mathfrak{M} is the generation of a new model $h(\mathfrak{M})$, where \mathfrak{M} and $h(\mathfrak{M})$ are exactly alike with respect to individualistic and non-individualistic properties *and* with respect to how those properties are co-instantiated, *except* for within an arbitrarily small region, the hole. Within the hole \mathfrak{M} and $h(\mathfrak{M})$ differ over the pattern of co-instantiation of individualistic and non-individualistic properties.

So, now here is why this causes a violation of MD1. Consider two models \mathfrak{M} and $h(\mathfrak{M})$ related by a hole diffeomorphism. Consider a cauchy surface Σ from \mathfrak{M} and a cauchy surface Σ' from $h(\mathfrak{M})$ such that Σ' is prior to the hole in $h(\mathfrak{M})$. Because h is trivial over all segments outside of the hole, Σ and Σ' will involve the same pattern of co-instantiation of individualistic and non-individualistic properties. It does not follow from this, however, that \mathfrak{M} and $h(\mathfrak{M})$ will globally involve the same pattern of co-instantiation of individualistic and non-individualistic properties, since there is a

region – the hole – where \mathfrak{M} and $h(\mathfrak{M})$ differ in this respect. So GR’s models violate MD1.

But so what? Perhaps this by itself isn’t all that worrying, since a violation of MD1 doesn’t automatically have any physical significance. The models of GR, after all, are just “heuristic devices”; they’re useful for helping us picture what really matters: GR’s worlds.² But the problem is, with LI in place, given diffeomorphic models \mathfrak{M} and $h(\mathfrak{M})$, we are committed to diffeomorphically related worlds w and $h(w)$. Diffeomorphically related worlds will violate a notion of determinism that *is* important. To see this, let’s first extract the relevant notion of determinism. We’ll begin with a notion of agreement between worlds.

W-Agreement 1: worlds w and w' **W-agree 1** *iff:* w and w' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

Next, we need a notion of agreement that corresponds to the notion of agreement between cauchy surfaces. This will be a notion of agreement between maximal achronal subregions.

S-Agreement 1: maximal achronal subregions S and S' **S-agree 1** *iff:* S and S' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

Now we can use these notions to formulate a notion of determinism for worlds that corresponds to the notion we developed earlier in terms of models.

World Determinism 1 (WD1): GR satisfies WD1 *iff:* for all GR worlds, w and w' , and for all maximal achronal subregions S in w and S' in w' : if

²Norton, for example, points out that a notion of determinism that remains confined solely to GRs models “is usually dismissed as a purely mathematical gauge freedom associated with active general covariance. (1993, 825)” See also Butterfield (1989, 10-11).

S and S' involve the same pattern of co-instantiation of individualistic and non-individualistic properties then, w and w' involve the same pattern of co-instantiation of individualistic and non-individualistic properties.

As I have noted in 2.4, haecceitist substantivalism violates WD1 (although, I called it ‘WD’). Let’s walk through this more carefully. Begin with two models \mathfrak{M} and $h(\mathfrak{M})$ related by a hole diffeomorphism. Now, under the assumption of **Haec**, given \mathfrak{M} and $h(\mathfrak{M})$, we commit ourselves to worlds w and $h(w)$ that differ in an analogous way. Recall that \mathfrak{M} and $h(\mathfrak{M})$ are exactly alike save for within some arbitrarily small region. With this region \mathfrak{M} and $h(\mathfrak{M})$ differ with respect to the pattern in which the individualistic and non-individualistic properties are co-instantiated – that is, they differ with respect to which points have which values of g and T . Similarly, w and $h(w)$ will be exactly alike save for within some arbitrarily small region, H . Within H , w and $h(w)$ will differ with respect to the pattern in which the individualistic and non-individualistic properties are co-instantiated – that is, they differ with respect to which points have which determinates of \mathbf{g} and \mathbf{T} . Given this, it will be possible to select some maximal achronal subregion S in w , and another maximal achronal subregion in S' in $h(w)$ prior to H , such that S and S' involve the same pattern of co-instantiation of individualistic and non-individualistic properties. And yet, because w and $h(w)$ differ within H , it does not follow from the fact that S and S' agree in this way, that w and $h(w)$ do, too.

As I noted in 2.4.3, the haecceitist substantivalist can respond to this issue by substituting a new notion of determinism. Here is one way to do this. First, it turns out that there is a notion of determinism in terms of models that GR satisfies. Instead of appealing to a notion of agreement that requires the patterns of co-instantiation of individualistic and non-individualistic properties to be identical, we appeal to a notion that requires them to merely be related *smoothly*. This will give us the following notion of agreement.

M-Agreement 2: models \mathfrak{M} and \mathfrak{M}' **M-agree 2** *iff:* \mathfrak{M} and \mathfrak{M}' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in \mathfrak{M} and \mathfrak{M}' are related by a smooth transformation.

By way of explaining what I mean here by a *smooth transformation*, let's consider the following game. You have a pegboard covered with hooks. There are two rings hanging on each hook – a blue one and a red one. The rings are tied to their immediate neighbors with pieces of string. The object of the game is to rearrange rings across hooks while satisfying three constraints. First, if a particular blue and a particular red ring are together on a hook, you cannot move them independently – they are a pair, they have to stay together. Second, every hook has to have one blue and one red ring after rearranging – no hook can be naked, and no hook can have more than two rings. Third, you cannot ever break any of the strings holding rings to their neighbors. Now, I'll say that a smooth transformation, in the context of the game, is a transformation of rings across hooks that satisfies the three constraints.

Its likely obvious, but I'll be explicit anyways: a pegboard with hooks is like a model, the hooks are like the individualistic properties of manifold points, and the red and blue rings are like the non-individualistic properties of models – the values of the metric and stress-energy tensors fields. The first constraint forces the metric and stress-energy to be coordinated. You cannot pair stress-energy and metric values willy-nilly and still satisfy the field equations. The second constraint ensures that every point gets a field value, and no point gets more than one value of a particular field. The third constraint forces the topology of the non-individualistic properties to remain the same while allowing the relationship between the individualistic and non-individualistic properties to change – the strings force neighboring points to remain neighboring points. So a smooth transformation in the context of a model is a trans-

formation of non-individualistic properties across individualistic ones that satisfies the above three constraints.

Next, we can make the corresponding amendment to C-agreement 1, which gives us the following:

C-Agreement 2: cauchy surfaces Σ and Σ' **C-agree 2** *iff:* Σ and Σ' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in Σ and Σ' are related by a smooth transformation.

This then gives us MD2:

Model Determinism 2 (MD2): GR satisfies MD2 *iff:* for all GR models, \mathfrak{M} and \mathfrak{M}' , and for all cauchy surfaces Σ found in \mathfrak{M} and Σ' found in \mathfrak{M}' : if Σ and Σ' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in Σ and Σ' are related by a smooth transformation then, \mathfrak{M} and \mathfrak{M}' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in \mathfrak{M} and \mathfrak{M}' are related by a smooth transformation.

Now, just as we mimicked MD1 for the purposes of generating WD1, we can mimic MD2 for the purposes of generating a physically significant notion of determinism that GR's worlds will satisfy. We begin with a notion of agreement between worlds.

W-Agreement 2: worlds w and w' **W-agree 2** *iff:* w and w' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in w and w' are related by a smooth transformation.

Next, we characterize agreement between maximal achronal subregions.

S-Agreement 2: maximal achronal subregions S and S' **S-agree 2** *iff*: S and S' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in S and S' are related by a smooth transformation.

This will give us the following notion of determinism.

World Determinism 2 (WD2): GR satisfies WD2 *iff*: for all GR worlds, w and w' , and for all maximal achronal subregions S of w and S' of w' : if S and S' involve the same individualistic and non-individualistic properties, and the pattern of co-instantiation of individualistic and non-individualistic properties in S and S' is related by a smooth transformation then, w and w' involve the same individualistic and non-individualistic properties, and the patterns of co-instantiation of individualistic and non-individualistic properties in w and w' are related by a smooth transformation.

GR's worlds will satisfy WD2. But this victory requires the concession on the part of the haecceitist substantialist that haecceitistic differences between worlds are irrelevant. The solution that I offer gives a haecceitist substantialist a way to have two notions of determinism: WD1 (again, in 2.4 I called it 'WD'), which can see haecceitistic differences, and RD which is blind to haecceitistic differences.

...

BIBLIOGRAPHY

- [1] Armstrong, D. M. (1989). *Universals: An Opinionated Introduction*. Boulder, CO: Westview Press.
- [2] Arntzenius, F. (2012). *Space, Time, and Stuff*. Oxford, UK: Oxford University Press.
- [3] Bain, J. (2003). Einstein algebras and the hole argument. *Philosophy of Science*, 70, 1073-1085.
- [4] Bain, J. (2006). Spacetime structuralism. Manuscript.
- [5] Belot, G. (2011). *Geometric Possibility*. Oxford, UK: Oxford University Press.
- [6] Bennett, J. (1984). *A Study of Spinoza's Ethics*. Indianapolis: Hackett Publishing Company.
- [7] Bricker, P. (1991). Plenitude of possible structures. *Journal of Philosophy*, 88, 607-619.
- [8] Bricker, P. (1993). The fabric of space: Intrinsic vs. extrinsic distance relations. *Midwest Studies in Philosophy*, 18, 271-294.
- [9] Bricker, P. (1996). Isolation and unification: the realist analysis of possible worlds. *Philosophical Studies*, 84, 225-238.
- [10] Bricker, P. (2001). Island universes and the analysis of modality. In Preyer, G., & Siebelt, F., (eds.), *Reality and Humean Supervenience: Essays on the Philosophy of David Lewis*. Rowman & Littlefield, 27-55.
- [11] Bricker, P. (forthcoming). *Modal Matters*. Oxford, UK: Oxford University Press.
- [12] Brighouse, C. (1994). Spacetime and holes. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 1, 117-125.
- [13] Butterfield, J. (1988). Albert Einstein meets David Lewis. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 2, 65-81.
- [14] Butterfield, J. (1989). The hole truth. *British Journal for the Philosophy of Science*, 40, 1-28.

- [15] Butterfield, J., & Isham, C. (2001). Spacetime and the philosophical challenge of quantum gravity. In Callender, C., & Hugget, N., (eds.), *Physics Meets Philosophy at the Plank Scale*, Cambridge, UK: Cambridge University Press.
- [16] Carroll, S. (2004). *Spacetime and Geometry: An Introduction to General Relativity*. Boston, MA: Addison-Wesley.
- [17] Casati, R., & Varzi, A. (1999). *Parts and Places: The Structures of Spatial Representation*. Cambridge, MA: MIT Press.
- [18] Clifford, W. K. (1876). On the space-theory of matter. In *William Kingdon Clifford: Mathematical Papers*. London, UK: Macmillan.
- [19] Cowling, S. (2014). Instantiation as Location. *Philosophical Studies* 167, 667-682.
- [20] Dasgupta, S. (2011). The bare necessities. *Philosophical Perspectives*, 25, 115-160.
- [21] DeRosset, L. (2009). Production and necessity. *Philosophical Review*, 118.
- [22] Eddon, M. (2013). Quantitative properties. *Philosophy Compass*, 8, 633-645.
- [23] Earman, J. & Norton, J. (1987). What price substantivalism? The hole story. *British Journal for the Philosophy of Science*, 38, 515–525.
- [24] Earman, J. (1989). *World Enough And Space-Time*. Cambridge, MA: MIT Press.
- [25] Ellis, B. (1999). Causal powers and laws of nature. In Sankey, H., (ed.), *Causation and Laws of Nature*, Dordrecht, DE: Kluwer.
- [26] Friedman, M. (1983). *Foundations of Space-Time Theories: Relativistic Physics and the Philosophy of Science*. Princeton, NJ: Princeton University Press.
- [27] Frigg, R. (2006). Scientific representation and the semantic view of theories. *Theoria*, 55, 49-65.
- [28] Geroch, R. (1972). Einstein algebras. *Communications of Mathematical Physics*, 26, 271- 275.
- [29] Graves, J. C. (1972). *The Conceptual Foundations of Contemporary Relativity Theory*. Cambridge, MA: MIT Press.
- [30] Hawking, S. & Ellis, G. F. R. (1973). *The Large Scale Structure of Space-time*. Cambridge, UK: Cambridge University Press.
- [31] Halvorson, H. (2012). What scientific theories could not be. *Philosophy of Science*, 79, 183-206.
- [32] Hawthorne, J. (2008). Three-dimensionalism vs. four-dimensionalism. In Sider, T., Hawthorne, J., & Zimmerman D., (eds.), *Contemporary Debates in Metaphysics*. Malden, MA: Blackwell.

- [33] Hazen, Allen (1979). Counterpart-theoretic semantics for modal logic. *The Journal of Philosophy*, 76.
- [34] Hofer, C. (1996). The metaphysics of space-time substantivalism. *The Journal of Philosophy*, 93, 5-27.
- [35] Krantz, D., Luce, D., Suppes, P., & Tversky, A. (eds.) (1971). *Foundations of Measurement, Vol. I: Additive and Polynomial Representations*. New York, NY: New York Academic Press.
- [36] Lehmkuhl, D. (2009). Spacetime matters: on supersubstantivalism, general relativity, and unified field theories. Dissertation. Oxford University.
- [37] Lehmkuhl, D. (2015). The metaphysics of supersubstantivalism. (Preprint).
- [38] Leonard, M. (forthcoming). What is mereological harmony? *Synthese*.
- [39] Lewis, D. (1968). Counterpart theory and quantified modal logic. *The Journal of Philosophy*, 65, 113-126.
- [40] Lewis, D. (1971). Counterparts of persons and their bodies. *The Journal of Philosophy*, 68, 203-211.
- [41] Lewis, D.(1973): *Counterfactuals*. Oxford, UK: Basil Blackwell.
- [42] Lewis, D.(1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61, 343-377.
- [43] Lewis, D. (1986). *On the Plurality of Worlds*. Cambridge, UK: Blackwell.
- [44] Lewis, David (2009). Ramseyan humility. In Braddon-Mitchell, D., & Nola, R. (eds.), *Conceptual Analysis and Philosophical Naturalism*. Cambridge, MA: MIT Press.
- [45] Liu, C. (1996). Realism and spacetime: of arguments against metaphysical realism and manifold realism. *Philosophia Naturalis*, 33, 243-63.
- [46] Maudlin, T. (1988). The essence of space-time. *The Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 82-91.
- [47] Maudlin, T. (1993). Buckets of water and waves of space: why spacetime is probably a substance. *Philosophy of Science*, 60, 183 – 203.
- [48] Maudlin, T. 2007). *The Metaphysics Within Physics*. New York, NY: Oxford University Press.
- [49] Maudlin, T. (2012). *Philosophy of Physics: Space and Time*. Princeton, NJ: Princeton University Press.
- [50] McDaniel, K. (2007). Extended simples. *Philosophical Studies*, 133, 131 – 141.

- [51] Melia, J. (1999). Holes, haecceitism and two conceptions of determinism. *British Journal for the Philosophy of Science*, 5, 639-664.
- [52] Misner, C. (1972). Some topics for philosophical inquiry concerning the theories of mathematical geometrodynamics and of physical geometrodynamics. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 7-29.
- [53] Mundy, B. (1992). Space-time and isomorphism. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 515-527.
- [54] Norton, J. (1993). General covariance and the foundations of general relativity: eight decades of dispute. *Reports of Progress in Physics*, 56, 791-861.
- [55] Norton, J. (2011). "The Hole Argument". The Stanford Encyclopedia of Philosophy (Fall 2011 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2011/entries/spacetime-holearg/>>.
- [56] Parsons, Josh (2007). Theories of location. In Zimmerman, D. (ed), *Oxford Studies in Metaphysics* 3. Oxford, UK: Oxford University Press
- [57] Pooley, O. (2002). The Reality of Spacetime. Dissertation. Oxford University.
- [58] Pooley, O. (2006). Points, particles and structural realism. In Rickles, D., French, S., & Saatsi, J. (eds.), *The Structural Foundations of Quantum Gravity*. Oxford, UK: Oxford University Press. 83-120.
- [59] Pooley, O. (2013). Substantivalist and relationalist approaches to spacetime. In Batterman, R. (ed.), *The Oxford Handbook of Philosophy of Physics*. Oxford, UK: Oxford University Press.
- [60] Pooley, O. (forthcoming). *The Reality of Spacetime*. Oxford, UK: Oxford University Press.
- [61] Penrose, R. (2004). *The Road to Reality: A Complete Guide to The Laws of The Universe*. Random House.
- [62] (1951). Quine, W. V. O. (1981). Two dogmas of empiricism. *Philosophical Review*, 60, 20-43.
- [63] Quine, W. V. O. (1981). Things and their place in theories. In Quine, W. V. (ed.), *Theories and things*, 1-23. Cambridge, MA: Harvard University Press.
- [64] Rickles, D. (2008). *Symmetry, Structure, and Spacetime*. Amsterdam, NL: Elsevier.
- [65] Rynasiewicz, R. (1992). Rings, holes, and substantivalism: on the program of Leibniz algebras. *Philosophy of Science*, 59, 572-589.
- [66] Rynasiewicz, R. (1994). The lessons of the hole argument. *British Journal for the Philosophy of Science*, 45, 407-436.

- [67] Rynasiewicz, R. (1996). Absolute versus relational space-time: an outmoded debate? *The Journal of Philosophy*, 93, 279-306.
- [68] Saucedo, R. (2011). Parthood and location. In Zimmerman, D. (ed), *Oxford Studies in Metaphysics* 6.
- [69] Schaffer, J. (2005). Quiddistic knowledge. *Philosophical Studies*, 123, 1-32.
- [70] Schaffer, J. (2009). Spacetime: the one substance. *Philosophical studies*, 145, 131-148.
- [71] Schafner, K. and R. Cohen, (eds). Proceedings of the 1972 Biennial Meeting of the Philosophy of Science Association, 64, 1974. D. Reidel Publishing Company, *Synthese* Library.
- [72] Shoemaker, S. (1980). Causality and properties. In van Inwagen, P. (ed.), *Time and Cause*, Dordrecht, DE: Reidel.
- [73] Sider, T. (2001). *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford, UK: Oxford University Press.
- [74] Simons, P. (2004). Extended simples: A third way between atoms and gunk. *The Monist* 87 (3): 371-384.
- [75] Sklar, L. (1974). *Space, Time and Space-Time*. Los Angeles, CA: University of California Press.
- [76] Skow, B. (2005). Once upon a spacetime. Dissertation. New York University.
- [77] Skow, B. (2008). Haecceitism, anti-haecceitism, and possible worlds. *The Philosophical Quarterly*, 58, 98-107.
- [78] Stachel, J. (1974). The rise and fall of geometrodynamics. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 31-54.
- [79] Stachel, J. (1980). Einstein's search for general covariance. In Howard, D., & Stachel, J. (eds.), *Einstein and the History of General Relativity (Einstein Studies, Volume 1)*. Boston, MA: Birkhuser, 63-100.
- [80] Suárez, M. (2003) Scientific representation: against similarity and isomorphism. *International Studies in The Philosophy of Science*, 17, 225-244.
- [81] Suppe, F. (1989). *The Semantic Conception of Theories and Scientific Realism*. Urbana, IL: University of Illinois Press.
- [82] Teller, P. (1991). Substances, relations and arguments about the nature of spacetime. *The Philosophical Review*, 100, 363-397.
- [83] Thomas, E. (2013). Space, time, and Samuel Alexander. *British Journal for the History of Philosophy*, 21, 549-569.

- [84] Uzquiano, G. (2011). Mereological harmony. In Zimmerman, D. (ed), *Oxford Studies in Metaphysics* 6, 199-224. Oxford, UK: Oxford University Press.
- [85] Wasserman, R. “Material Constitution”, The Stanford Encyclopedia of Philosophy (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2015/entries/material-constitution/>.
- [86] Wesson, P. S. (2007). *Space-Time-Matter*. World Scientific.
- [87] Wheeler, J. A. (1962). *Geometrodynamics. Italian Physical Society: Topics of Modern Physics*, 1.