Metaphysical Lessons of General Covariance

Shamik Dasgupta

Draft of May 2009, comments welcome.
Please do not cite without permission.

1 The Covariance Substantivalism Link

In this paper I will argue that if our best-confirmed physical theory satisfied General Covariance, we would have good reason to think that a certain view about the metaphysical structure of our world, i.e. substantivalism, is false. Since many of our best confirmed physical theories these days do in fact satisfy general covariance, this will amount to an argument against substantivalism.[1]

It will be useful to proceed with a particular physical theory in view. To this end I choose the General Theory of Relativity (GTR), which is uncontroversially agreed to satisfy General Covariance. The central claim I will argue for is that if GTR were our best confirmed candidate for the true and complete physical theory of our world, the fact that it satisfies General Covariance would give us good reason to think that substantivalism is false. Having focused on this case, it will be clear how the argument generalizes to other generally covariant theories.

My central claim uses two terms that require clarification: 'General Covariance' and 'substantivalism'. The first is tricky to define in full generality but easy to characterize just in the context of GTR, which is all we need.[2] In this context it is normally defined in terms of the intended models of GTR. These are triples $(M,\ g,\ t)$, where M is a four-dimensional differentiable manifold and g and t are a metric and a mass-energy tensor field, respectively, defined on M. For simplicity we can assume that all intended models

[1] In fact, it amounts to something stronger: most theories *can* be formulated in such a way as to satisfy General Covariance, and it has been argued that our best theory *should* be so formulated. If that's right, this paper will amount to an argument that we should expect substantivalism to fail.

[2] There is a large literature on the real meaning of General Covariance; see Rynasiewicz [18] for a good review. Fortunately, the issues raised in that literature will not matter to us here. If the reader worries whether my description of General Covariance gives the *real* meaning of the term, she may simply reformulate my central claim as being that if GTR were our best confirmed candidate for the true and complete physical theory of our world, the fact that it satisfies *the property that I define* (whether or not that property really is General Covariance) gives us good reason to think that substantivalism is false.

have the same manifold of points M, though of course different models will differ in the fields g and t. Now, take a model $\mathbf{M} = (M, g, t)$ and consider a diffeormophism $f{:}M{\rightarrow}M$, i.e. a bijection from M to M such that it and its inverse are both differentiable. We can construct a new "drag along" structure $f(\mathbf{M}) = (M, f(g), f(t))$, where $f(g)$ is a tensor field that results from "dragging" the value of g at each point in the original model $\mathbf{M}$ along to its image under f. Formally: for all x in M, the value of the field $f(g)$ at the point $f(x)$ is $g(x)$. The field $f(t)$ is defined similarly. Now, GTR is said to satisfy *General Covariance* because given any model $\mathbf{M}$ of GTR and any diffeomorphism f, $f(\mathbf{M})$ is also a model of GTR.

What about the term 'substantivalism'? A completely general definition of the term would take us too far afield so, as before, let us simply ask what the term means in the context of GTR. To this end, it is helpful to introduce the view by reference to the models of GTR. Models are abstract objects that represent physical scenarios. Whenever an abstract object represents a physical scenario, we need to ask which aspects of the abstract object represent something real and which are artifacts of the representation itself. In this case, we may ask: is there really a manifold of points that is represented in some way by M, or is M just an artifact of the representation? As a first pass, we can think of substantivalism as the view that M is no artifact: there really is a manifold of points represented by M.

Helpful though this characterization is at first, it is worth taking some time to refine it in three ways. First, my initial characterization might encourage the view that substantivalism is a view about what represents what, i.e. about whether the manifold M represents a manifold of points in reality. But whether x represents y is in large part a function of how we use x, and substantivalism is not a claim about how we use M. Rather, substantivalism should be understood as a purely metaphysical claim about the world, the claim that there is a manifold of points in reality. Strictly speaking, this claim is neutral on whether the manifold in reality is in fact represented by the manifold M in models of GTR: one might, at least in principle, be a substantivalist and think that there is a manifold of points in reality, without thereby thinking that they are represented by M. What the substantivalist *does* think is that her manifold of points instantiates the same topological structure as M (this is the grain of truth in my initial characterization of GTR).

Second, it will be important in what follows to emphasize that substantivalism is not simply the claim that a manifold of points exists. Rather, it is that claim that the manifold exists *and is a fundamental part of reality.* To see why an additional clause of this sort is needed, recall the classic debate between Leibniz and Newton on the nature of space: Newton thought that space was a substance that existed independently of the matter it contained, and Leibniz disagreed. But what was the nature of their disagreement? Did Leibniz simply think that there is no such thing as space? Not really, for he

may have treated space as a "logical construction", in which case he would have agreed with Newton that there is such a thing as space and that all matter is situated within it. Instead, what distinguishes him from Newton is his insistence that all facts about location in space ultimately reduce to facts about how material bodies are spatially related to one another. That is, Leibniz thought that *fundamentally speaking* all facts about location are facts about how material bodies are spatially related to one another, and that all other facts about location—e.g. facts about where something is located "in space"—reduce to such relational facts. Newton's position can thus be characterized as the view that facts about the existence of space, its structure, and the positions of material bodies in it, are themselves fundamental.

Once the dispute between Newton and Leibniz is seen as a dispute about what is fundamental, it is clear that we should characterize substantivalism in the context of GTR in the same way. Any fan of desert landscapes may agree that the world contains a manifold of points, so long as she takes care to add that the manifold is a "logical construction" or something of that ilk. What she will deny is that the manifold is a *fundamental* part of reality, and it is precisely this claim that the substantivalist endorses.

Later we will see how important it is to recognize this sort of "fundamentality" clause, so it should not be surprising that one finds it implicit in most work on substantivalism. For example, Hoefer characterizes substantivalism as the claim that the manifold exists '*independently* of material things... and is properly described as having *its own* properties, *over and above* the properties of any material things that may occupy parts of it.'[3] Although he doesn't use the term 'fundamental', the terms I emphasized in this quotation suggest to me that he recognizes that substantiavlism is not simply the view that the manifold exists but includes some extra claim to the effect that it is fundamental.

Admittedly, there are many questions about how to understand claims about fundamentality. One important issue that will come up later is whether it is to be analyzed in terms of metaphysical possibility, or whether it is best taken as a primitive or analyzed in some other terms. Another issue concerns what sorts of entities are properly said to be fundamental. One view is that fundamentality is a feature of *facts*; on this view the substantivalist will claim that facts about the existence of the manifold and certain facts about its nature are fundamental. Another view is that fundamentality is a feature of *objects*; on this view the substantivalist will claim that the manifold itself is a fundamental object. For ease of prose I will often talk as if the first view was correct, though nothing hangs on this choice and the reader can translate everything I say into the second view if desired. Thus, I understand the substantivalist as claiming that amongst the fundamental

[3]Hoefer [11], p. 5; my emphasis.

facts of the world are facts about the structure of the manifold and how the fields are distributed over it.

Now, precisely what form these facts take will depend on one's views about the metaphysics of fields. If one thinks that fields are fundamentally mathematical objects—functions from points in the manifold to some mathematical space—then the fundamental facts will be of the form

p is mapped to m by the metric field

where p is a point in the manifold and m is a mathematical object. If instead one thinks of fields as instantiations of physical properties or relations at points, then the fundamental facts will be of the form

p instantiates F

or

p bears R to q

and so on, where F is a physical property and R is a physical relation, and p and q are points in the manifold. But ignoring details about how we think of fields, the substantivalist's defining claim is that facts such as these are amongst the fundamental facts of the world.

Actually, this is still not quite the best way to characterize the view. The way I just put things suggests that the substantivalist is committed to thinking that the fundamental facts of the world concern *points* in the manifold, but substantivalists may in fact disagree on this point. Thus, some substantivalists might think that the fundamental facts concern the manifold as a whole, and that facts about the individual points such as the above are derivative.[4] Other substantivalists might think that the manifold is "gunkish", i.e. that every part of the manifold has proper parts, in which case there will be no points (at least, not at the fundamental level).[5] But this complication won't matter for our purposes, so for ease of prose I shall ignore it and suppose for simplicity that our substantivalist thinks that the fundamental facts of the world include facts which concern points in the

[4]For example, see Schaffer [20].

[5]See Arntzenius [1] for a development of this sort of view about the mereological structure of space. In fact, if every region has proper parts it will be somewhat unclear what the substantivalist should think the fundamental facts are. One way to develop the view is to say that there are fundamental facts about the distribution of the fields over a given region, even though the region has proper parts. Another way to develop it is to claim that there are no fundamental facts, but that all facts about the distribution of the fields over any region in the manifold reduce to facts about the field distribution throughout their proper parts, and so on *ad infinitum*. If so, then the view would count as "substantival" not because of any claim about what the fundamental facts are, but because it claims that facts about the manifold always hold in virtue of other facts about the manifold and never in virtue of facts about the matter situated within it.

manifod, such as those listed above. Other substantivalists can translate my claims into their own favored framework if they wish.

So this is how I understand substantivalism, at least in the context of GTR. It is worth mentioning that as I understand the view it is entirely neutral on whether the manifold of points deserves to be called 'space-time'. I mention this because there is a debate in the literature about whether space-time is to be identified with the manifold of points on its own or with the manifold along with the metric field, and this is characterized by some participants as a debate about what 'substantivalism' should be taken to mean in the context of GTR.[6] Obviously they are free to use the term as they wish, and there may be virtues of retaining the historical tie between substantivalism and space-time. But for the sake of clarity I should emphasize that as I use it 'substantivalism' is stipulated to denote a claim about whether a manifold of points exists in reality, and is completely neutral on whether the manifold of points is to be identified as "space-time". (I would add parenthetically that I don't think the question of what counts as space-time in GTR has a determinate answer, since I doubt that the meaning we attach to 'space-time' is determinate enough to withstand such close scrutiny.)

I am now in a position to outline my aims in this paper. The central claim I wish to defend is:

> The Covariance Substantivalism Link: If GTR were our best-confirmed candidate for the true and complete physical theory of our world, the fact that it satisfies General Covariance would give us a good reason to think that substantivalism is false.

Now, why should we believe that this Link holds? General Covariance is a mathematical, model-theoretic property of GTR while substantivalism is a metaphysical thesis about the fundamental nature of the world, so why think that the former is of any relevance at all to the latter? As far as I know, the literature contains only one argument in favor of it: Earman and Norton's infamous Hole Argument.[7] Since it is uncontroversial that GTR satisfies general covariance, it is unsurprising that substantivalists responded to their work by arguing that the Hole Argument fails and that, contra the Covariance Substantivalism Link, general covariance gives us no reason whatsoever to doubt the truth of substantivalism.

Since my aim is to argue that the Link is true, I obviously have a disagreement with these substantivalists. But I also have a disagreement with Earman and Norton, for I think that the Hole Argument is unconvincing. Dialectically, then, my position is this: I agree with Earman and Norton

[6]For a defense of the former view, see Earman [5] chapter 9. For a defense of the latter view see Maudlin [14] and Hoefer [11].

[7]See Earman and Norton [6], and Earman [5] chapter 9.

that the Covariance Substantivalism Link is true but I disagree over why; and I agree with substantivalists that the Hole Argument is not convincing, but I disagree with them because I think that the Link is true for other reasons. My aim in this paper is to motivate this position.

I'll start in section 2 by discussing the Hole Argument. I will argue that it makes an assumption that the substantivalist is in no way committed to, and (moreover) that the assumption is contentious to say the least. In section 3 I will outline my own argument in favor of the Covariance Substantivalism Link, and I'll argue that it does not make the controversial assumption made by the Hole Argument. Finally, in section 4 I will make a few remarks about how we should go about developing a novel metaphysics for GTR if substantivalism is rejected.

Two caveats are needed. First, I suspect that some substantivalists will claim that the new metaphysics I end up developing counts as a version of substantivalism—different, perhaps, from typical substantivalist views, but substantival nonetheless. This claim isn't entirely unjustified, and indeed I will argue that whether we call the new view "substantival" is largely a verbal issue. The real issues are whether the Hole Argument gives us reason to reject substantivalism-as-typically-understood—I will argue that it does not—and whether there is a different argument from General Covariance that *does* give us reason to reject it—and I will claim that there is. Now, these issues can be discussed while remaining neutral on whether the view I end up endorsing should be called 'substantivalism', but for ease of presentation it helps to leave this issue till later and pretend for the moment that it is substantivalism itself on the line, as the Covariance Substantivalism Link states.

The second caveat is that by "the Hole Argument" I mean the argument that is typically attributed to Earman and Norton in the literature. While some of their presentations of the argument support this attribution, other presentations make me suspect that they had in mind something like the argument I will end up endorsing. If so, then my disagreement with Earman and Norton as to whether the Hole Argument is convincing is not really a disagreement with them, but rather a disagreement with one way that the argument has been presented in the literature. Nonetheless, for ease of prose I shall work on the assumption that the standard presentation of the argument is really the argument that Earman and Norton had in mind.

Enough preliminaries: let us turn to examining whether the Hole Argument gives us any reason to accept the Covariance Substantivalism Link.

2 The Hole Argument and its Discontents

The thrust behind the Hole Argument is that because of General Covariance substantivalism implies that GTR is indeterministic, a consequence which

is taken to be unacceptable. The argument has in fact been presented in at least three different ways, and unfortunately the differences matter. I will start with the worst version, which I will put aside, and then introduce the two versions that will be under discussion in what follows.

In the first version one picks a model $\mathbf{M} = (M, g, t)$ and a diffeomorphism f that is identity on all of M save for a bounded region, the 'hole'. One then argues as follows:

(1) If substantivalism is true then $\mathbf{M}$ and $f(\mathbf{M})$ represent different possible worlds in which GTR is true.

(2) If (1) is true then substantivalism implies that GTR is indeterministic.

(3) That implication is unacceptable: indeterminism should not follow straight from a theory such as substantivalism.

(4) Therefore, substantivalism is false.

The idea behind (1) is this: by hypothesis, $\mathbf{M}$ is a model of GTR; so by General Covariance $f(\mathbf{M})$ is a model too. So *if* the two models represent distinct possible worlds, it follows that they represent distinct possible worlds in which GTR is true. These worlds differ only in facts concerning the specific points in the hole at which the field values are located. Since the metric and mass-energy fields are both shifted around in the same way, they agree on all facts about space-time distances between bits of matter; and they even agree on the specific field values of all points outside the hole.

Notice that General Covariance by itself does not ensure that $f(\mathbf{M})$ represents a *possible* world at all, or even that it represents a *distinct* possibility from that represented by $\mathbf{M}$. What (1) assumes, then, is that the substantivalist is committed to the claim that $\mathbf{M}$ and $f(\mathbf{M})$ represent distinct possible worlds. But given this assumption, General Covariance then ensures that they are both worlds in which GTR is true. It is because General Covariance plays a crucial role in motivating (1) that this argument can be used to establish the Covariance Substantivalism Link.

The reasoning behind premise (2) is as follows. Consider a time t to the past of the hole. By construction, the possible world represented by $f(\mathbf{M})$ agrees with that represented by $\mathbf{M}$ on all facts at t and yet diverges thereafter, namely with respect to facts about the specific points at which the field values are located within the hole. Now, the standard definition of determinism is that a theory is deterministic if and only if any two possible worlds satisfying the theory that agree in all respects at one time agree in all respects at all later times. On this definition, it follows that GTR is indeterministic. Of course, this reasoning needs to be refined to get rid of my appeal to "times", but let us ignore those details for now.

Premises (1) and (2) concern possible worlds, but what sort of possibility is at issue here? It is universally assumed in the literature that it is

metaphysical possibility. This is the concept we use when we say that it is impossible for a cat to be a vegetable, or that it is impossible for me to have had different parents. Of course, one might worry about whether there is a coherent notion of metaphysical possibility in the offing here, but it is assumed in the literature on the Hole Argument that there is and I shall not question that here.

As for premise (3), their reasoning is a little unclear. They don't think that indeterminism is apriori false. Nor do they rule out discovering empirically that GTR is indeterministic. Instead, their idea seems to be that if GTR is indeterministic this should follow from facts about the physics, not from metaphysical theses such as substantivalism. After all, if the indeterminism follows straight from substantivalism and the General Covariance of GTR, then it will follow straight from substantivalism and *any* Generally Covariant theory! In their own words: 'If a metaphysics, which forces all our theories to be deterministic, is unacceptable, then equally a metaphysics, which automatically decides in favour of indeterminism, is also unacceptable. Determinism may fail, but if it fails it should fail for reasons of physics...'.[8] Of course, this raises the question of what makes something a reason of *physics* rather than a reason of *metaphysics*. But as unclear as (3) might be, it is granted in all discussions of the Hole Argument and I won't question it for now.

Improving the Argument

Instead, I'd like to argue that given the definition of determinism used in (2), premise (1) is ill-formulated. After all, what matters for (2) is the claim that substantivalism is committed to the existence of distinct possible worlds in which GTR is true and which differ only on facts that occur within the hole. If the substantivalist is committed to the existence of such worlds, and if the above definition of determinism is accepted, it follows that substantivalism implies that GTR is indeterministic. But if that's how the argument is supposed to go, it does not matter one bit whether the substantivalist must take these distinct worlds to be represented by $\mathbf{M}$ and $f(\mathbf{M})$. Facts about how models *represent* possible worlds are neither here nor there when it comes to (2).

This is a small point, but it is worth emphasizing because a large bulk of the substantivalist literature responding to the Hole Argument is concerned to argue that the substantivalist is free to interpret models related by a hole diffeomorphism, such as $\mathbf{M}$ and $f(\mathbf{M})$, as representing the same possible world. In defense of this claim, it is often pointed out that physicists typically interpret their models in this way, and that substantivalists are presumably free to follow them. An example of such a theorist is Brighouse,

[8] Earman and Norton [6], p. 524.

who writes that "It is clear that the argument loses its force if a substantivalist can hold that Leibniz equivalent models [i.e. models related by a diffeomorphism] represent the same situation".[9] But this line of argument is misguided. True, substantivalists can indeed hold that models related by a hole diffeomorphism represent the same situation; after all, they're free to adopt the convention of interpreting their models in that way if they so wish! And, true, if that's right then premise (1) is, strictly speaking, false. But to emphasize this point is to miss the force behind the argument. Even if the substantivalist interprets models related by a hole diffeomorphism as representing the same possible world, this is consistent with *the existence of* pairs of distinct possible worlds in which GTR is true and which differ only in a diffeomorphic shift of the fields within a bounded region. And so long as there exist two such worlds, the argument in (2) goes through to establish that GTR is indeterministic. All that follows from this method of interpreting models is that the two possible worlds won't be represented by diffeomorphically related models such as $\mathbf{M}$ and $f(\mathbf{M})$.

This suggests that the first premise should really be reformulated as follows (called '(1W)', for Worlds):

(1W) If substantivalism is true then there are distinct possible worlds in which GTR is true and which differ only in a diffeomorphic shift of the field values within a bounded region (the "Hole").

Here we recognize that the focus on models was a distraction, and re-focus the argument on possible worlds instead.

This is not to say that models are of no use in defending (1W), for they can be thought of as tools used to describe what the two possible worlds are like. The idea is this: just as the substantivalist is free to interpret $\mathbf{M}$ and $f(\mathbf{M})$ as representing the same world if they so wish, so too are *we* presumably free to interpret models in such a way they represent different worlds (remaining neutral for the moment as to whether both worlds are *possible*). Actually, this idea needs a slight qualification: it might be that our concept of metaphysical possibility doesn't allow us to make sense of differences between such worlds, in which case we obviously couldn't interpret our models in this way. But putting this qualification aside until later, the idea is that *if* it so much as makes sense to consider differences between worlds of the sort described by (1W), then we could presumably decide to adopt the convention of interpreting models in such a way that those differences are represented by models such as $\mathbf{M}$ and $f(\mathbf{M})$. If we interpret them this way, we get a convenient and precise way of describing two worlds that differ only in a diffeomorphic shift of their fields across the manifold—the shift being described by the function f. Having described the worlds and arguing that the substantivalist is committed to accepting both as possible,

[9] Brighouse [3], p. 119.

the argument then proceeds through premises (2) and (3) independently of any issue about how substantivalists do, can or should interpret their models, for that issue is entirely beside the point.

What is the role of General Covariance in this argument? It ensures that *if* both worlds are possible, then since the first world is a GTR world (by hypothesis), so too is the second. Of course, General Covariance does not ensure that the two worlds are in fact both possible, so the first premise is justified only on the additional assumption that the substantivalist must think that they are. Nonetheless, as in the last version of the argument, General Covariance plays a crucial role in justifying the first premise.

That's the second version of the Hole Argument. The third version differs from the second only by eliminating any reference to possible worlds and instead using modal operators such as "it is possible that...". Thus, on this reading the first premise reads as follows (called '(1O)' for Operators):

(1O) If substantivalism is true then it is possible for the field-values to have been diffeomorphically shifted only within a bounded region (the "Hole").

On this version of the argument, the definition of determinism will be re-expressed in a similar way, replacing quantification over worlds with modal operators. But for the most part the distinction between this argument and the last won't matter, so in what follows I will discuss the second version of the argument which uses (1W) unless otherwise stated.

Substantivalist Responses

In response, substantivalists have invariably taken issue with either the first or second premise. Those who deny the second premise agree (at least for the sake of argument) with (1W), i.e that substantivalism does indeed imply that there are distinct possible GTR worlds that differ only in a diffeomorphic shift. But they deny that the existence of the possible worlds mentioned in the (1W) implies indeterminism. Why? Because they think that the definition of determinism used by Earman and Norton is incorrect. That definition said that a theory is deterministic if and only if any two possible worlds satisfying the theory that agree in all respects at one time agree in all respects at all later times (as I said before, this definition must be refined to eliminate reference to times, but those details won't matter so I'll ignore them.) What deniers of the second premise say is that this is too strong: worlds that agree at one time needn't agree at later times in *all* respects but only in some. Precisely which respects are relevant is a delicate matter, but a standard move is something like this: a theory is said to be deterministic, say these substantivalists, if and only if any two worlds satisfying the theory that agree in all respects at a time t agree at all other times *in general, qualitative respects and in respects concerning the individuals that exist at* t.

The idea is that the only respects in which diffeomorphically shifted worlds differ are facts about which particular points in the hole the field-valeus are instantiated at, and these are non-qualitative respects that do not concern individuals that existed at the prior time. Therefore, on this conception of determinism the existence of the two worlds does not imply indeterminism. Defenders of this move bring to bear an impressive armory of intuitions supporting their conception of determinism.[10]

I will discuss this strategy of responding to the Hole Argument in more detail later on. For now, I would like to focus on the more popular substantivalist strategy: namely, to deny (1W) and argue that substantivalism does not, after all, imply that there exist two distinct possible worlds that differ only on facts concerning the field-values in the hole.

I think that this strategy is clearly right. Substantivalism, recall, is a claim about *what the fundamental facts of our world are*, namely that they include facts about how the fields are distributed over the manifold. But the first premise is about *what is metaphysically possible*. What is the relation between claims about what's fundamental and claims about what's metaphysically possible? Well, one might argue that all facts *supervene* on the fundamental facts, i.e. that there is no difference between two possible worlds without a difference in their fundamental facts. But beyond that, claims about metaphysical possibility are left entirely open by claims about what the fundamental facts are, and so are left entirely open by the substantivalist. For example, the substantivalist might be a radical modal fatalist and think that all truths are necessarily true, in which case she will immediately deny (1W)! Admittedly, fatalism is rather implausible and it would be something of a hollow victory for the substantivalist if by denying (1W) she had to endorse views such as this, so in a moment I will outline a number of respectable and popular views in modal metaphysics which all imply that (1W) is false. But until then, I just want to make the minimal point that, in and of itself, substantivalism is a claim about what the fundamental facts of the world are, and so (other than the above qualification) it has no modal consequences at all—at least, not when the modality is metaphysical possibility—*afortiori*, it does not imply that there exist two distinct possible worlds that differ in a hole diffeomorphism.

This, I believe, is the main weakness of the Hole Argument. Its stated aim is to argue that a consequence of substantivalism—i.e. the possibility of diffeomorphic shifts—leads to indeterminism, but unfortunately the putative consequence it mentions is not a consequence after all.

This point is relatively straightforward but, surprisingly, it has been systematically overlooked by protagonists of the Hole Argument. For example, Earman accuses the substantivalist who denies that there are diffeomorphi-

[10]Melia [16] and Skow [23] advocate responding to the Hole Argument along these lines.

cally shifted worlds of endorsing a 'cheap instrumentalist rip-off'.[11] And Belot concurs, arguing that the substantivalist's defining claim is that shifted worlds are possible.[12] Given how I have set up the issues, these attitudes are surprising. What explains them? One explanation is that these theorists simply mean something different by 'substantivalism' than I do. But there is another explanation of the disagreement that is perhaps more illuminating. As I emphasized when characterizing substantivalism, everyone agrees that it is not *just* the view that a manifold of points exists, since that leaves it open that facts about the manifold reduce to facts about the relative positions of bits of matter. Rather, it is the claim that facts about the manifold are fundamental and not reducible to other kinds of facts. But I also mentioned that there is an issue about how to understand claims of the form 'X is fundamental' and 'X reduces to Y'. Until recently, there has been a tendency to understand these claims in terms of metaphysical possibility. Thus, 'X reduces to Y' will be understood, on this account, as meaning that X supervenes on Y. Understood like this, the anti-substantivalist who claims that the manifold is a logical construction out of matter will be understood as claiming that facts about the manifold supervene on facts about the relative positions of bits of matter; or, to put it otherwise, that possible worlds which agree on facts about the relative positions of bits of matter will also agree on all facts about the manifold. Therefore, on this understanding of reduction it is definitional of the anti-substantivalist's view that there do not exist possible worlds that differ only in a hole diffeomophism. But if that's how the *anti*-substantivalist's view is understood, then obviously it will be definitional of the substantivalist's view that there *are* such differences between possible worlds! Thus, my best guess is that Earman and Belot are both understanding claims about fundamentality and reduction in terms of metaphysical possility.

But this is not the only way to understand fundamentality and reduction. In fact, there is a growing recognition in the literature that we should instead understand these notions as primitive.[13] Understood like this, claims about fundamentality have few modal consequences. True, as mentioned above we might concede that what's fundamental places limits on what's possible; in particular, we might concede that everything supervenes on the fundamental, so that there are no differences between metaphysically possible worlds without a difference in their fundamental facts. But even if what's fundamental places *limits* on possibility, it does not imply what *is* possible: one may reasonably think that the manifold is a fundamental part

[11]Earman [5], p. 127. See also Earman and Norton [6] for an expression of a similar attitude.

[12]See Belot, [2]. Strictly speaking, in this paper he discusses the case of Leibnizian shifts in classical space-time structures, but his points easily generalize to the case of the Hole Argument.

[13]See Fine [8], Schaffer [19] and Sider [21] for starters.

of reality without thereby being committed to any claims about what is possible regarding the distribution of fields across it. So my diagnosis of why protagonists of the Hole Argument have missed the relatively simple point that substantivalism does not in and of itself have any modal consequences is just that they have been assuming one, and only one, way of understanding what claims about fundamentality and reduction amount to.

Digression: The Varieties of Anti-Haecceitism

As I said, this would be a hollow victory for the substantivalist if the only way to deny the possibility of diffeomorphic shifts is to become a radical modal fatalist. So it is worth outlining a number of respectable and popular views in modal metaphysics which all imply that (1W) is false. I don't intend to argue that any of these views are true, but rather to emphasize that anyone wanting to run the Hole Argument must first establish that they are false. In the next section I will outline my own argument for the Covariance Substantivalism Link which makes no assumptions about metaphysical possibility. So, while these views give substantivalists a way to respond to the Hole Argument, *none* of them will be of any use when responding to my argument.

The substantivalist who denies that there are worlds which differ only in a hole diffeomorphism thinks that there is a limit to metaphysical possibility. But why should metaphysical possibility be limited in this way? One view is that this limit is a brute fact about metaphysical possibility, but there are in fact a variety of different factors in terms of which this limit might be explained.

One view locates the source of the limit in the essences of the individuals involved. Thus, if one thought that it is impossible for me to have been a fried egg, one might explain why this is impossible by appealing to the fact that it is part of my essence that I am human. Similarly, the substantivalist might claim that points in the manifold have essences, and that this is why diffeomorphic shifts are impossible. Maudlin takes this tack when he claims that the metric properties of points in the manifold—the properties and relations specified by the metric tensor field g—are essential to points in the manifold that possess them.[14] Since hole diffeormophic shifts alter the metric properties of points, it follows from this view that they are metaphysically impossible.

While Maudlin locates the source of the impossibility in the nature of the points in the manifold, others locate it in the nature of metaphysical possibility itself. These views have often fallen under the heading of "anti-haecceitism", though as we will see there are many distinctions to be drawn under this heading. To discuss the issue, it will help to introduce a little

[14]See Maudlin [14].

terminology. Call facts that concern particular individuals, such as

a is F, b is G, a bears R to b, $a \neq b$

individualistic facts. And call facts that do not concern any particular individual *general facts*. These include facts that can be expressed in predicate logic without constants (but with identity), such as

$(\exists x)Fx$, $(\exists y)Gy$, $(\exists x)(\exists y)Rxy$, $(\exists x)(\exists y)x \neq y$

(I assume that the predicates just used do not include substitutes for constants, such as 'x Socratizes'.) These facts may be taken to imply the existence of individuals, but they are not individualistic facts because they do not concern any particular individual.

Notice that the two putative worlds that differ in a diffeomorphic shift agree on all general facts and differ only in their individualist facts. Now, Lewis famously held the view that there are no mere individualistic differences between possible worlds: worlds that agree in all general respects, he claimed, are the same simpliciter.[15] He called this view "anti-haecceitism", but since we will be considering other views that have been discussed under the same title let us call it "L-anti-haecceitism" instead, for Lewis. Note that this is a general thesis about the extent of the space of possible worlds, not a particular view about the manifold of points. So if one had independent reasons to endorse L-anti-haecceitism, then if one were also a substantivalist one would have independent (and non-ad hoc) reasons to think that there do not exist two distinct possible worlds that differ only in a diffeomorphic shift. Therefore, on this view substantivalism would in no way imply indeterminism—even on Earman and Norton's definition of indeterminism. Pooley advocates this response to the hole argument. On considering the question of whether there are distinct possible worlds that differ only in their individualistic facts, he writes that "there is nothing anti-substantival about denying that there can be such distinct possible worlds."[16]

Now, the obvious question facing the L-anti-haecceitist is *why* there are no worlds that agree in all general respects and differ only in individualistic respects. One answer we have already considered is that things have essences in such a way as to rule out differences between worlds that agree in all general respects, but here I am interested in answers that locate the source of the limit in the nature of metaphysical possibility itself.

One class of answers to this question come from views about the relation between individuals in different possible worlds. Consider two metaphysically possible worlds w and $w*$, and individuals i and $i*$ in each world respectively. Is it the case that $i = i*$? If so, what does such a fact consist

[15] See Lewis [12].
[16] Pooley [17], p. 22.

in? One view is that facts about the identity and distinctness of individuals in different possible worlds hold in virtue of the qualitative similarity of the individuals. Let's call this view Q-anti-haecceitism, since it takes Qualitative similarity to be what makes for trans-world identity. Another view is that it does not make sense when describing two possible worlds to ask whether an individual in one world is identical to an individual in the other world. This is *not* the view that individuals from distinct worlds are distinct, but rather that there is no fact of the matter as to their identity or distinctness. Kaplan called this view "anti-haecceitism", so let us call it "K-anti-haecceitism" in his honor.

Importantly, Q-anti-haeccietiism and K-anti-haecceitism both imply L-anti-haecceitism: if two worlds were exactly the same in all qualitative (and therefore general) respects and yet differed in individualistic respects, then there would have to be facts about the identity and distinctness of individuals in the two worlds, and these would have to be fixed by some means other than qualitative similarity. Therefore, if there are no facts about identity and distinctness, or if they are fixed qualitatively, there can be no mere individualistic differences between worlds. So if one had independent reasons to endorse Q-anti-haecceitism or K-anti-haecceitism as a thesis about the nature of metaphysical possibility, then if one were also a substantivalist one would have independent and non-ad hoc reasons to think that there do not exist two distinct possible worlds that differ only in a diffeomorphic shift.[17]

Unsurprisingly, then, a number of substantivalists have endorsed these views when responding to the Hole Argument. For example, Brighouse endorses Q-anti-haecceitism when she recommends that "what a substantivalist should say about the way we individuate space-time points or regions across possible worlds is that we individuate according to qualitative similarity".[18]

Hoefer appears to recommend endorsing K-anti-haecceitism when responding to the Hole Argument. He says that the argument makes the unjustified assumption that space-time points 'have a primitive identity wholly independent of the properties [they] actually posses', and he says that the substantivalist should respond by denying 'the ascription of *primitive identity* to space-time points'.[19] Admittedly, this all sounds dangerously anti-substantival: if the identities of space-time points are not primitive, this makes it sound like space-time points are logical constructions out of

[17]It is worth noting that L-anti-haecceitism does not imply K-anti-haecceitism or Q-anti-haecceitism: on Lewis' own view of possibilia there are facts of the matter about the identity and distinctness of individuals in two possible worlds, and these facts are brute and do not hold in virtue of qualitative similarity, but nonetheless he thinks that the space of possible worlds does not contain two distinct worlds that agree generally but differ individualistically.

[18]Brighouse [3], p. 122.

[19]Hoefer [11], p. 14.

the matter and metric fields, contra the substantivalist's position. However, when he goes on to develop the view it becomes clear that the denial of "primitive identity" is not in fact the denial that space-time points are fundamental entities, but rather a claim about the nature of the space of possible worlds that is reminisent of K-anti-haecceitism. He writes

> Suppose I have two dice and name them A and B by pointing to them. I now ask: Does it make sense to ask whether die A could have been located where die B is with all B's actual properties, and die B located where A is with all of A's properties? Note that the question is not: Is this *possible*?... I want to suggest that the question, though grammatically sound, makes no sense—or rather, rests on a confusion.

But why does the question not make sense? The natural answer is that there are no facts of the matter as to whether an individual in the one world is identical to an individual in the other world, which is exactly what K-anti-haecceitism states.[20]

Up till now I have focused on views about the nature of the space of possible worlds, which are views that are most relevant to responding to the second version of the hole argument. To respond to the third version which uses premise (1O), one must focus on modal claims using possibility and necessity operators. Now, many of the above views are normally taken to imply that claims such as "it is possible for the fields to have been diffeomorphically shifted across the manifold" are false.[21] Obviously the details here will depend on the specific way in which the framework of possible worlds is used to give a semantics for modal claims such as this. But still, the above "anti-haecceitisms" are typically views that a theorist holds about the nature of possible worlds in part because they hold views about the status of modal claims of the sort made in (1O) on which they are false (or, perhaps, meaningless).

But other views about the semantics of modal operators imply that diffeomorphic shifts are impossible too. For example, Butterfield outlines a response based on Lewis' counterpart theory.[22] On his view, points in dis-

[20] One might also interpret his quotation in another way: that it is only when considering isomorphic worlds that it does not make sense ask about the identities of individuals across the two worlds. It is consistent with his quote that it does make sense in other cases. But this "hybrid" view seems ad hoc, so it is more natural to think of him as a K-anti-haecceitist.

[21] I should stress that this isn't strictly speaking required. Lewis, for example, was an L-anti-haecceitist, but he held a liberal enough theory of the counterpart relation which allowed him to make sense of the possibility of hole diffeomorphisms. On his view, counterparts can be drawn from *within* a world, so although there is no *distinct* possible world that can be used to make true the claim that hole diffeomorphs are possible, the actual world can make this true on its own. "One possible world, two possibilities", as he once said.

[22] See Butterfield [4].

tinct worlds are distinct points, so he is no K-anti-haecceitist. To give a semantics for modal discourse, he uses a counterpart relation that relates points when they agree on their field values. Now, unlike Lewis, he does not explicitly endorse L-anti-haecceitism, so he leaves it open that there are distinct though isomorphic possible worlds; and, consequently, he leaves it open that he rejects Q-anti-haecceitism too. But given the counterpart relation he uses, counterparts will always have the same field values. Therefore, claims such as "it is possible for the fields to be diffeomorphically shifted across the manifold" will turn out false.

I have outlined a number of views in modal metaphysics all of which imply that diffeomorphic shifts are impossible. I do not claim that any of these views are correct. Rather, my point is that if the Hole Argument shows that any view is false, it shows that the conjunction of substantivalism *along with a certain conception of metaphysical possibility* is false. But it does not show that substantivalism itself is false, so it misses its target.

3 The Argument from Redundancy

Let us recap. The central claim I wish to defend in this paper, remember, is this:

> The Covariance Substantivalism Link: If GTR were our best confirmed candidate for the true and complete physical theory of our world, the fact that it satisfies General Covariance would give us a good reason to think that substantivalism is false.

So far, I have argued that the Hole Argument gives us little reason to accept this claim. I will now outline what I believe is a better argument which makes no assumptions about metaphysical possibility or determinism whatsoever. Therefore, all the substantivalist manouevers we just explored involving determinism, essence, anti-haecceitism, counterpart theory and so on will be of no use in evading my argument.

Learning from the Case of Velocity

The style of argument I want to run will not be new to readers familiar with the philosophy of physics literature on the metaphysics of motion, so it will be useful to discuss that issue first. Let's distinguish between *absolute* and *relative* velocity. A car cruising down the highway alongside a train might have a velocity of 55 mph to the north *relative to the highway* and 10 mph to the north *relative to the train*. These are both velocities relative to another material body. But how fast is it *really* going, independent of any reference point? If there is an answer to this question, then this is its *absolute* velocity.

Now, most philosophers of physics these days think there is no such thing as absolute velocity, only velocity relative to other reference points, none of which is privileged over the others as defining its *real* velocity.[23] Moreover, they tend to think this not for apriori reasons but rather for broadly empirical reasons. Of particular interest to us is that there is an almost universal agreement in the following:

> The Galilean Velocity Link: If Newtonian Gravitation Theory (NGT) were the best-confirmed candidate for the true and complete physical theory of our world, the fact that it is Galilean Invariant would give us a good reason to think that there is no such thing as absolute velocity.

where by NGT I mean Newton's three laws of motion and his inverse-square gravitational force law.[24] I am going to argue that, when properly understood, the universally accepted argument in favor of the Galilean Velocity Link makes no assumptions about determinism or metaphysical possibility (sections 3.2 and 3.3), and that exactly the same style of argument can be used in favor of the Covariance Substantivalism Link (section 3.4). If this is right, then we will have an argument for the latter that is entirely immune to all the standard manouevers that substantivalists have used to evade the Hole Argument.

The Standard Argument

Admittedly, arguments for the Galilean Velocity Link often use the language of metaphysical possibility—either in the form of quantification over possible worlds or else modal operators. I will first present the argument in this way, and then show that this sort of presentation is a mere illustrative convenience and quite inessential to the main thrust behind the argument.

First, what does it mean to say that NGT is Galilean Invariant? Unfortunately there are as many definitions of the term as there are theorists who

[23] I allow "reference points" to include unoccupied inertial trajectories in substantival space-time. Incidentally, it might be worth noting that my characterization of absolute velocity leaves open many views about its nature. On most views, absolute velocity is not itself a fundamental property but is defined in other terms, usually as the limit of the spatial distance traveled by the body over time. But at least in principle, one might treat it as a fundamental property of material bodies.

[24] I am slurring over a subtlety here. The laws of NGT *as written down by Newton himself* make reference to absolute velocity, so what Newton wrote down would not be true if there were no such thing as absolute velocity. But I am using the phrase 'laws of NGT' to refer to laws that can be expressed in different ways depending on what one takes to be the underlying metaphysics of the world they govern. We use 'the Schrodinger equation' in quantum mechanics the same way, to refer to a law that will be formulated in very different ways depending on one's view about the fundamental ontology of a quantum mechanical world. What the Galilean Velocity Link should really be understood as stating, then, is that orthodoxy considers a theory that dispenses with absolute velocity and formulates Newton's laws without reference to it better than Newton's own theory.

discuss the topic. Here I will just use one definition in terms of metaphysically possible worlds. Consider the set of possible worlds in which there is such a thing as absolute velocity. A "Galilean Boost" is a function on this set of worlds that maps each world to the one that differs only in the fact that everything's velocity is uniformly boosted, at all times, by some specified amount in some specified direction. The claim that NGT is Galilean Invariant is the claim that Galilean boosts always map worlds in which NGT is true to worlds in which NGT is also true.

Now, the orthodox argument for the Galilean Velocity Link is that the fact that NGT is Galilean Invariant means that absolute velocity is *redundant* to NGT, and should be dispensed with by Occam's razor. But it is worth distinguishing two versions of this argument.

The first version is that because NGT is Galilean Invariant, the particular absolute velocity of each particle is of no relevance to whether or not NGT is true; after all, Galilean Invariance means that the result of systematically changing all their absolute velocities preserves the truth of NGT. So the feature is said to be "superfluous" to the truth of the laws, and Occam's razor has us dispense with it.[25]

But one might object to this argument on the grounds that Occam's razor does not tell us to dispense with structure that is superfluous to the truth of the *laws*; rather, it tells us to dispense with structure that we cannot detect and which is is redundant to explaining all available data. Thus, the second version of the argument from redundancy is that it follows from the Galilean Invariance of NGT that absolute velocity is indeed undetectable and redundant to explaining all available data, and therefore we have good reason (by Occam's razor) to prefer theories about the structure of the material world that dispense with it.

To make this argument clear it is crucial to distinguish between causal and metaphysical explanations. Consider the fact that particle a is 2 meters from particle b at time t. Very roughly, a *metaphysical* explanation of this fact is an answer to the question of what the fact consists in, or what it holds in virtue of, in terms of facts that are more fundamental. For example, one might say that it holds because a and b are each located in a region of space, and that those regions are themselves 2 meters apart (incidentally, this would be a substantivalist explanation because it appeals to regions of space). In contrast, a *causal* explanation is (roughly speaking) an answer to the question of how this fact came to hold. For example, one might explain it, in this sense, by saying that five minutes ago the two particles were created 10 meters apart, and then mention that each was subjected to certain forces which, along with the laws governing the system, entailed that they will end up five minutes later being 2 meters apart. Of course,

[25]This is the justification that Earman gives in his [5], p. 45. In fact, the principle he defends in this way is a good deal more stronger than the Galilean Velocity Link.

there is much to say to clarify each kind of explanation, but I will take it that we understand them well enough to be getting on with.

With this distinction in mind, the argument for the Galilean Velocity Link that I have in mind says that it follows from the Galilean Invariance of NGT that absolute velocity is undetectable and also redundant to *causal* explanations of the data, and therefore Occam's razor gives us good reason to doubt that it is real.

Let's spell this out in a little detail. Suppose we live in a world in which the laws of NGT are the true and complete laws of motion, and consider once again the fact that a is 2 meters from b at t. Suppose we are to causally explain it in the above way, i.e. by stating various facts about the state of the system at t minus 5 minutes. What sorts of facts would be relevant to the explanation? Well, what we know from Galilean Invariance is that facts about the absolute velocities of each particle would be completely irrelevant. After all, consider the result of subjecting the world to a Galilean boost. Clearly, at t minus 5 minutes the boosted world differs only in facts about the absolute velocities of things. And since NGT is Galilean Invariant, it we know that the boosted world is also a world in which NGT is true. It follows that the boosted world represents how the particles would behave, given the laws of NGT, were the particles to have those different velocities at t minus 5 minutes. But since Galilean boosts only alter facts about absolute velocity, the boosted world is also a world in which a is 2 meters from b! What this means is that if NGT are the laws governing our world, prior facts about absolute velocity made no difference to whether or not a ended up 2 meters from b at t, and are in this sense redundant to causal explanations of this fact.

This argument generalizes to a well-defined class of facts. Say that a fact is "invariant under Galilean boosts" just in case if it holds in a world w, it also holds in $g(w)$ for any Galilean boost g. Facts about relative velocities, relative positions, accelerations and so on are all invariant under Galilean boosts. Clearly, the argument above generalizes to show that facts about absolute velocity will be redundant to causal explanations of *any* fact that is invariant under Galilean boosts.

Now, some theorists are tempted at this point to conclude that absolute velocities are redundant to causal explanations simpliciter. But this is not true: consider the fact that a is in a state of absolute rest at t. To causally explain this fact I will have to appeal to facts about a's absolute velocity at prior times![26] So facts about absolute velocity are indeed relevant to causal explanations of facts about absolute velocity. What this means is that *if* facts about absolute velocities were part of our data, then facts about absolute velocities would be necessary components of causal explanations of the data. But—and here's the trick—it follows from Galilean Invariance

[26]This point was made by Sklar in his [22], p. 180.

that facts about absolute velocities are *undetectable*, i.e. are never part of our data. Therefore, it follows from Galilean Invariance that absolute velocity is redundant to explaining the data after all.

Why does it follow from Galilean Invariance that facts about absolute velocity are undetectable? The key point is that if w is a possible world and g is a Galilean Boost, then $g(w)$ is a world that is *indistinguishable* from w: at any particular time, everything looks and feels and smells exactly the same in $g(w)$ as it does in w. So, let us suppose that in w I undergo some sort of activity in order to discover my absolute velocity at a time t. We need make no assumptions about what this sort of activity is—it could be as simple as opening my eyes and having a look, or as complicated as building a multi-billion dollar particle accelerator underneath the Alps—other than it being an activity that is completely governed by the laws of NGT and that it concludes in some finite amount of time in an outcome: a needle pointing in some direction, perhaps, or some sort of computer printout. Of course, this activity has its counterpart in $g(w)$, a world in which my absolute velocity at t differs. Now, the crucial point once again is that since NGT is Galilean Invariant, $g(w)$ is also a world that is completely governed by NGT. Therefore, the outcome of the activity in $g(w)$ represents what the outcome of my actual experiment would have been were my absolute velocity to have differed in the way specified by g. But the outcome in $g(w)$ is indistinguishable from the outcome in w! Since g is an arbitrary Galilean boost, what this shows is that no matter what experiment or measurement I undertook, it would have produced an indistinguishable outcome no matter what my absolute velocity was; which is to say that it is not an activity that counts as my having detected my absolute velocity.[27]

So that's the idea behind the argument for the Galileo Velocity Link: the fact that NGT is Galilean Invariant is being used to argue that if the laws of NGT were the true and complete laws governing our world, then (i) absolute velocity would be redundant to causal explanations of all facts save those concerning absolute velocities, and also (ii) this latter class of facts would be undetectable. It follows that absolute velocity would be undetectable and redundant to all causal explanations of the data, as required. More fully, the argument from redundancy goes like this:

(1) If the laws of NGT were the true and complete laws of motion governing our world, then the fact that they are Galilean Invariant means that absolute velocity would be undetectable and redundant to all causal explanations of the data.

(2) Occam's razor gives us good reason to dispense with structure that is undetectable and redundant to causal explanations of the data.

[27] I discuss this sort of argument in more detail in the first paper of this dissertation, *Symmetry and the Undetectable*.

(3) Therefore, the Galilean Velocity Link is true.

Note that just because we have a *good* reason to dispense with absolute velocity, it does not follow that we have an *all-things-considered* reason. To take one example that will be useful later on, suppose one thought that facts about relative velocity are to be *metaphysical* explained in terms of facts about absolute velocity, the idea being the relative velocity between two bodies is defined as the difference between their absolute velocities. Well, since NGT implies that facts about relative velocity *are* causally relevant to explanations of the data, it would follow on this view that there is no way to dispense with absolute velocity without thereby rejecting NGT itself. In a case like this, dispensing with absolute velocity runs the risk of leaving us with an empirically inadequate theory, so our best overall theory of the world might contain absolute velocity even though we are fully aware of its causal redundancy. Indeed, I believe that this was Newton's own view of absolute velocity: he knew it was redundant to *causal* explanations of the data, though he also believed it was a necessary component of *metaphysical* explanations of other features (e.g. relative velocity, acceleration) that are themselves relevant to of *causal* explanations of the data. Nonetheless, what the argument from redundancy shows us is that in this sort of case, the *only* barrier to dispensing with absolute velocity is that one thinks it is necessary for metaphysically explaining other causally relevant features: if we were to come up with a reasonable alternative metaphysics of those causally relevant features, we could happily dispense with dispense with absolute velocity. And indeed this is precisely what we did when we learnt how to metaphysically explain relative velocity and acceleration in the context of a 4-dimensional "Galillean" space-time structure. So the Occamist reason to dispense with causally redundant features is certainly not an all-things-considered reason to think that they are not real, but instead should be seen as a good (albeit defeasible) reason.

Removing the Modality

Now, when I just argued for premise (1) I made an assumption about metaphysical possibility, namely that there are possible worlds that differ only in facts about absolute velocity. Put otherwise, I assumed that for any world w in which there is such a thing as absolute velocity, $g(w)$ is a distinct possible world for any Galilean boost g.

I made this assumption when I defined what it is for NGT to be Galilean Invariant. It was also made when I argued that facts about absolute velocities are redundant to causal explanations of other facts such as those concerning distances. Remember, what I initially argued was that NGT worlds like w and $g(w)$ agree on all facts about distances at any given time but differ on temporally prior facts about absolute velocity, and I took that to

imply that the earlier facts about absolute velocity are redundant to causal explanations of the later facts about distances. Finally, I made the assumption when I argued that absolute velocity is undetectable: what I initially argued for was the claim that no activity of measurement or experimentation could possibly allow us to distinguish between w and $g(w)$, and I took that to imply that absolute velocity is undetectable.

But was it necessary to make this assumption? Perhaps not. After all, what was crucial to premise (1) was the idea that it follows from NGT being the true and complete laws of motion governing our world that absolute velocity is *undetectable and redundant to causal explanations of the data*. True, the above argument for (1) established this result by *first* establishing results about the space of metaphysically possible worlds, and *then* taking those results to imply the redundancy and undetectability of absolute velocity. And since the strategy of argument proceeded in modal terms, we started off by defining Galilean Invariance in those terms also: that's why what we ended up establishing was the Galilean Velocity Link. But for all I have said, there may be other ways to establish that absolute velocity is causally redundant and undetectable which do not take the detour through metaphysical possibility but instead proceed in other terms. And if so, it may be that we could define Galilean Invariance in those other terms too. If so, we might still establish the Galilean Velocity Link, but this time the statement of the link and its defense would be understood in terms that make no assumptions about the extent of the space of metaphysically possible worlds.

Very well; but is it plausible that one can run the argument non-modally? To see that it is, consider again a radical modal fatalist, Al, who thinks that if p is true then p is necessarily true. We might imagine that Al justifies his fatalism on the basis of his conviction that God is a necessary being, that God necessary created the best of all putatively possible worlds, and that there is therefore only one way the world could possibly have turned out to be. But whatever justifies his fatalism, suppose he also believes that there is such a thing as absolute velocity and that he lives in a world in which NGT is the true and complete physical theory. Now, suppose we try to convince Al that there is in fact no such thing as absolute velocity by running our argument above. Since our argument for (1) made the assumption about the extent of the space of possible worlds, one response available to Al is to say 'You made assumptions about metaphysical possibility which I reject. Therefore your argument is unconvincing and I retain my belief in absolute velocity'.

As a response to the *letter* of our argument, this is perfectly fair. But should we admit defeat? I think not: intuitively, Al's curious beliefs about the nature of metaphysical possibility has no bearing on the real thrust behind our argument for (1), which was simply *that absolute velocity is undetectable and redundant to causal explanations of the data*. This is some-

thing that Al could, in principle, agree with even if he denies that there are possible worlds that differ only in Galilean boosts.

To see this, let's start small and build up. First, Al is likely to agree that the fact that a is 2 meters from b does not causally depend on facts about absolute velocities at priori times. This would appear to be consistent with his radical modal fatalism: the notion of causal dependence is one he is free to recognize, so long as he does not analyze it in terms of metaphysical possibility—he may take causal dependence to be a primitive, or he may analyze it in yet other terms. Second, he is free to think that absolute velocity lacks this causal role in virtue of the fact that the laws of NGT are the true and complete laws of motion. Again, this is consistent with his fatalism: the notion of one fact holding in virtue of another is one he is free to recognize, so long as he does not analyze it in terms of metaphysical possibility.

He could even say that absolute velocity has this causal role in virtue of the fact that the laws of motion are Galilean Invariant. Here he will of course need to understand Galilean Invariance in non-modal terms, but this is not an insurmountable problem. One attractive approach here is to define Galilean Invariance as I did above but with reference to *models* instead of possible worlds. For Al, the space of models will of course not represent the space of possible worlds, but he could instead use it to represent the structure of causal dependence between different states of affairs. For example, if there are a whole range of models in which a is 2 meters from b at t but which all differ on the absolute velocities of things at earlier times, Al might take this to represent the fact that a's being 2 meters from b at t does not causally depend on facts about absolute velocity

Al is also likely to agree that the Galilean Invariance of the laws means that absolute velocity is undetectable. Here one might complain that the "able" in "undetectable" is involves metaphysical possibility and is therefore out of bounds to Al given his fatalism, but this is not clear. One might, for example, understand it as the claim that the quality of our experiences do not causally depend on facts about absolute velocity. Finally, if the counterfactual formulation of (1) offended his fatalistic sensibilities, he could simply reformulate it as the claim that absolute velocity is undetectable and redundant to causal explanations of the data, and that this is true in virtue of the fact that the laws of motion, NGT, are Galilean Invariant.

So even given his radical modal fatalism, Al is in a position to recognize an argument for (1) when reformulated like this, understanding all these terms in an appropriately non-modal way.

Similar remarks apply to a theorist with less extreme modal views than Al. Consider Bob who, like Al, believes in absolute velocity and thinks that he lives in a world in which the laws of NGT are the true and complete laws of motion. Unlike Al, though, Bob is no modal fatalist and thinks that many non-actual states of affairs are possible. Nonetheless, he thinks that

facts about absolute velocity supervene on facts about relative velocity, in the sense that any two worlds agreeing on all facts about relative velocities also agree on all facts about absolute velocity. Now Bob, like Al, will deny the central assumption of our original argument for (1)—i.e. that there are possible worlds that differ only in facts about absolute velocity. But does that mean that we should admit defeat? Again, I think not: intuitively, Bob's strange modal beliefs are neither here nor there when it comes to the thrust behind our argument, which was simply *that absolute velocity is redundant to causal explanations of the data.* Like Al, Bob could in principle accept premise (1)—or, at least, a reformulation of it such as that just discussed—even if he doubts that there are possible worlds that differ only in Galilean boosts.

This is not to say that the original argument for (1) presented in terms of metaphysically possible worlds should be abandoned. For one thing, most of us are happy to accept its central assumption. And for another thing, there is a certain vividness to running the argument in modal terms, describing the possible worlds and seeing in one's mind that there is no distinguishable difference between them. Thus, it makes for a particularly compelling argument to suppose for the sake of argument that its modal assumption is true. But when properly understood the assumption can be seen as a mere convenience to elicit emotive force, and is harmless since one is assured that the argument goes through without it.

Extending the Argument

Having outlined the argument from redundancy for the Galilean Velocity Link, I propose to run an exactly analogous argument for the Covariance Substantivalism Link. This, remember, was the following claim:

> The Covariance Substantivalism Link: If GTR were our best-confirmed candidate for the true and complete physical theory of our world, the fact that it satisfies General Covariance would give us a strong reason to think that substantivalism is false.

Now, the key to the argument I have in mind is to recall the substantivalist's defining claim, namely that amongst the fundamental facts of the world are facts about the structure of the manifold and how the fields are distributed over it. As I said in Section 1, the precise form of these facts depends on one's view about the metaphysics of fields. Thus, if one thinks of fields as fundamentally mathematical objects they will be facts of the form

> p is mapped to m by the metric field

where p is a point in the manifold and m is a mathematical object, but if one thinks of fields as instantiations of physical properties or relations then they'll be facts of the form

p instantiates F

or

p bears R to q

where F is a physical property and R a physical relation. I also mentioned that some substantivalists will deny that the fundamental facts concern *points* in the manifold but will instead concern the manifold as a whole, while others will claim that all regions of the manifold have proper parts and therefore that the fundamental facts concern not points but regions with proper parts (though see footnote 5 for a more about how this view might be spelt out). Nonetheless, *all* these theorists will agree that the fundamental facts of the world are *individualistic* in the sense introduced earlier: they concern particular individuals.

I believe that it is precisely *this* component of substantival thought that gives rise to a tension with General Covariance. I will argue that General Covariance gives us reason to believe that there are no such things as individualistic facts concerning the manifold—at least, not at the fundamental level of description. The argument will not depend on anything specifically to do with space-time; indeed, as I said in Section 1, I've defined substantivalism in such a way that it is an open question whether the view has anything to do with space-time in the first place! Instead, the argument zeros in on the substantivalist's claim that the fundamental facts of the world are individualistic: my idea is that General Covariance gives us reason to think that this is false. Why? The thought is that because of General Covariance, individualistic facts have the same status in GTR as facts about absolute velocity do in NGT. So, mimicking our argument against absolute velocity, the idea is to argue from the General Covariance of GTR that individualistic facts about the manifold are redundant to causal explanations of all the data.[28]

In fact, we can argue something stronger. Not only are individualistic facts about the manifold redundant to causal explanations of the data, but they are also causally redundant to all facts whatsoever and are therefore even more "causally isolated" than facts about absolute velocity in NGT. As a result, our argument from redundancy for the Covariance Substantivalism Link should be at least as convincing as our argument from redundancy for the Galilean Velocity Link.

As above, it will be convenient to first argue for this by making substantial assumptions about metaphysical possibility, and then showing that the argument can be restated in non-modal terms. Thus, let us assume that the space of metaphysically possible worlds includes worlds that are

[28]The second paper in this dissertation, *Individuals: An Essay in Revisionary Metaphysics*, expands on this sort of argument.

diffeomorphic shifts of each other. Now, consider a fact F about the distribution of the fields in one region R of the manifold, and let us suppose that part of the causal explanation of this fact appeals to facts about the distribution of the fields in a non-overlapping bounded region, R*.[29] What sorts of facts about the distribution of the fields in R* could these be? Well, what we know from General Covariance is that individualistic facts concerning the points within R* are completely irrelevant. After all, consider any hole diffeomorphism within R* (of which there are many!), and consider the corresponding diffeomorphically shifted possible world. This world differs only with respect to individualistic facts within R*, and since GTR satisfies General Covariance it is also a world in which GTR is true. It follows that the shifted world represents how the fields could evolve, given the laws of GTR, if the individualistic facts within R* differed in this way.[30] But since the shifted world agrees on *all* facts outside R*, the shifted world is also a world in which F holds! What this means is that if GTR were the true and complete laws governing our world, individualistic facts within R* would make no difference to whether or not the fact F ends up holding, and are in this sense redundant to causal explanations of it.

The argument is similar, of course, to that which we ran in the case of absolute velocity. However, note that here we needed to make no assumption about what sort of fact F is: it could be an individualistic fact or a general fact. This is because diffeomorphic shifts can change the individualistic facts within R* but leave *all* the facts outside of R* exactly the same. Galilean boosts don't do that: if a boost changes the absolute velocity of any particle at any time, it changes the absolute velocity of all particles at all times.

What about the claim that individualistic facts are undetectable? The argument proceeds in precisely the same way as it did in the case of velocity. The key point, once again, is that if w is a possible world and d is a diffeomorphic shift, then $d(w)$ is a world that is *indistinguishable* from w: in any particular region, everything looks and feels and smells exactly the same in $d(w)$ as it does in w. So, let us suppose that in w I undergo some sort of physically governed activity in order to discover the individualistic facts within a certain region R of the manifold. Of course, this activity has its counterpart in $d(w)$, a world in which the individualistic facts within R differ. And, once again, the crucial point is that since GTR satisfies General Covariance, $d(w)$ is also a world that is completely governed by GTR. Therefore, the outcome of the activity in $d(w)$ represents what the outcome of my actual experiment would have been had the individualistic facts within R

[29]Since I am making no assumptions about the size of R, this supposition introduces no loss of generality. After all, the only possibility this supposition excludes is that F is wholly causally explained by facts spatio-temporally coincident with F, in which case we would not take F to have been *causally* explained.

[30]I say that it represents how the fields *could* evolve, rather than how they *would* evolve, because of the worries about indeterminism brought out by the original Hole Argument.

differed in the way specified by d. But the outcome in $d(w)$ is indistinguishable from the outcome in w! Since d is an arbitrary diffeomorphism, what this shows is that no matter what experiment or measurement I undertook, it would have produced an indistinguishable outcome no matter what the individualistic facts were within R; which is to say that it is not an activity that counts as my having detected those facts.

So we are in a position to offer the following argument in favor of the Covariance Substantivalism Link:

(1) If GTR was the true and complete physical theory governing our world, the fact that it satisfies General Covariance means that individualistic facts about the distribution of the fields through the manifold would be undetectable and redundant to causal explanations of all facts whatsoever.

(2) Occam's razor gives us good reason to dispense with undetectable structure that is redundant to causal explanations of all facts whatsoever.

(3) Therefore, the Covariance Substantivalism Link is true.

I just outlined my argument in favor of premise (1), and premise (2) is just the same as it was in the argument for the Galilean Velocity Link so I won't say anything more about it here.

Now, the argument for (1) just presented proceeds in modal terms, and in doing so it assumes that there are possible worlds that differ only in a diffeomorphic shift. But, as in the case of our argument from redundancy in the case of velocity, this assumption is entirely dispensable. Consider again Al, our modal fatalist, and suppose he is a substantivalist who thinks he lives in a world in which the laws of GTR are true and complete, and suppose we presented him with the above argument for (1). Of course, he will reject its central assumption and so his belief in substantivalism will remain unfazed. But does this mean we should admit defeat? I think not, for as before Al's curious modal beliefs are completely irrelevant to the real thrust behind our argument, which was simply *that individualistic facts are undetectable and redundant to causal explanations.* This is something Al can agree with; all he will insist on is that the notions in play must not be understood in terms of metaphysical possibility.

Thus, Al can agree that no facts about the field distribution causally depend on individualistic facts about the field distribution, so long as he does not analyze causal dependence in modal terms. He can even agree that this is true in virtue of the fact that the laws governing his world satisfy General Covariance, so long as he does not analyze the notion of being "true in virtue of", or the property of General Covariance, in modal terms. Indeed, recall that we initially defined General Covariance in terms of models, not worlds,

so Al may naturally adopt our initial definition of the concept! Of course, he will not interpret the models and representing possible worlds. Instead, he may interpret them (as before) as representing the structure of causal dependence that holds between different facts about the field distribution across the manifold, as before. Al is also likely to agree that the fact that the laws satisfy General Covariance means that individualistic facts about the distribution of the manifold are undetectable; though, as before, he will have to understand the "able" in 'undetectable' in non-modal terms. Finally, if he took issue with the counterfactual formulation of (1) he could simply reformulate it as the claim that individualistic facts about the manifold are undetectable and redundant to causal explanations, and that this is true in virtue of the fact that the laws of motion, GTR, satisfy General Covariance.

So, once again, even given his radical modal fatalism, Al is in a position to recognize an argument for (1) when reformulated like this, understanding all these terms in an appropriately non-modal way.

Similarly, consider a substantivalist who is also an "anti-haecceitist" of one of the stripes that we encountered earlier on. This theorist thinks that the manifold of points is a fundamental part of reality, it's just that she also has views about the nature of metaphysical possibility which imply that individualistic facts supervene on general facts, and therefore that there do not exist possible worlds that differ only in diffeomorphic shifts. Thus, she will deny the central assumption we made in our original argument for (1) above. But does that mean we should admit defeat? Not at all, for her modal convictions are entirely irrelevant to the real thrust behind our argument, which was that individualistic facts are undetectable and not relevant to causal explanations. Like Al, the anti-haecceitist could in principle accept (1)—or, at least, a reformulation of it such as that just discussed—even if he doubts that there are possible worlds that differ only in their individualistic facts.

The Hole Argument Resuscitated?

Earlier I argued that the Hole Argument made assumptions about metaphysical possibility that the substantivalist was free to deny. But now that we've seen that my argument from redundancy can be run without modal assumptions, the natural thought is that we might run the Hole Argument without modal assumptions too. It is not hard to see how. Above I focused on the idea individualistic facts are redundant to causal explanations of other facts: in order to causally explain a fact F concerning the field distribution in a region R, we saw that individualistic facts within any non-overlapping region R* were irrelevant (where R* is naturally thought of as "temporally prior" to R). But similar reasoning will convince us that individualistic facts are themselves never causally explained. In modal terms, the idea will be this: pick any individualistic fact F about the manifold within a region R

(the "hole"), and consider a diffeormophism d that is identity outside R and non-identity within R. Since the d shifted possible world agrees on all facts outside R but is a world in which F does not obtain, the idea is that F does not causally depend on the facts outside R—after all, it may vary freely without any variation in the facts outside R. And one might naturally conclude from this that GTR is indeterministic. This is, of course, just the Hole Argument once again.

Now, the way I just said this made assumptions about metaphysical possibility. But if we could run my argument from redundancy without appeal to such assumptions, we should be able to run this argument without those assumptions too! After all, the thrust behind this argument is simply *that individualistic facts about the manifold are not causally determined or explained*. This is something that Al, our radical modal fatalist, can agree with; all he will insist on is that the notions of causal dependence and the like must not be understood in terms of metaphysical possibility. Similarly, it is something that any of the "anti-haecceitists" could agree with, so long as (once again) they phrase the argument without reference to diffeomorphically shifted possible worlds. So we would appear to be in a position to run the following sort of argument:

(1) Individualistic facts are not causally determined by the past, and this is true in virtue of the fact that the laws of motion, GTR, satisfy General Covariance.

(2) If individualistic facts are not causally determined by the past, GTR is indeterministic.

(3) This consequence is unacceptable: indeterminism should not follow straight from a theory such as substantivalism.

(4) Therefore, the Covariance Substantivalism Link is true.

So it would appear that at least *vis a viz* assumptions about metaphysical possibility, this reconstruction of the Hole Argument is at least as defendable as my argument from redundancy.[31]

Is this version of the Hole Argument to be recommended? I certainly think it's better than the argument as originally presented in Section 2, since it doesn't make the controversial modal assumptions and is therefore immune from the "anti-haecceitist" responses outlined back there. Nonetheless, I would still recommend my argument from redundancy over this argument, for two reasons.

First, the strategy of (all versions of) the Hole Argument is to argue that substantivalism entails that GTR is *indeterministic*, and unfortunately this

[31]Of course, given our relativistic framework this argument should ultimately be reformulated so as to eliminate reference to the "past", but this complication is not relevant to our concerns here.

leads to unwanted complications for the strategy. To see this, recall from Section 2.3 that some theorists have argued that once we understand what determinism amounts to we'll see that premise (2) is false.[32] More specifically, what these theorists argue is that in order satisfy our intuitive concept of determinism, a theory need not imply that *all* facts are determined by the past but only that all *non-individualistic* facts are, and that this is satisfied by GTR even if substantivalism were true.[33] Now, one question that arises here is whether these theorists are right about their analysis of the word 'determinism': defenders of this view offer an impressive array of examples that support their analysis, while defenders of the Hole Argument attempt to explain those intuitions away. But I don't want to engage in that dispute because I don't think that an analysis of the word 'determinism' is what's important here. What's important is whether the property of GTR that premise (2) claims is an implication of substantivalism—i.e. the property of being a theory according to which individualistic facts are not causally determined by the past, never mind whether we call this property "being indeterministic"—is a property that should not be allowed to follow straight from a "metaphysical" thesis such as substantivalism, as premise (3) claims it shouldn't. To my mind, it is extremely unclear whether this is the case. I will not try to settle the issue here: it suffices to point out that the Hole Argument faces a very delicate question of why a metaphysical thesis such as substantivalism should not be allowed to imply that individualistic facts aren't determined by the past.

The nice thing about the argument from redundancy is that it avoids this sort of issue. For it will be agreed by all that individualistic facts are undetectable and redundant to causal explanations of the data, just as it is agreed by all who discuss the Hole Argument that individualistic facts aren't determined by the past. But unlike the proponent of the Hole Argument, I do not now face the delicate question of why the existence of such facts should not be allowed to follow a thesis such as substantivalism. This is because arguments from redundancy are as old as the hills: everyone agrees that we have at least *some* reason to be skeptical of facts that are undetectable and causally redundant according to our best confirmed physical theory of the world! So my argument has the virtue of sidestepping the extremely delicate question that the Hole Argument plows straight into.

4 How To Be An Anti-Substantivalist

Whose Burden?

That completes my argument for:

[32] These theorists include Melia [16] and Skow [23].

[33] The details here differ from theorist to theorist, but that is the main idea.

> The Covariance Substantivalism Link: If GTR were our best-confirmed candidate for the true and complete physical theory of our world, the fact that it satisfies General Covariance would give us a good reason to think that substantivalism is false.

Now, substantivalism was the claim that a manifold of points exists and is a fundamental part of reality. So if we are to reject this and become anti-substantivalists, what should we say instead about the fundamental structure of the world?

Many theorists have remarked that this question falls squarely within physics. The idea is that since the Einstein field equations of GTR are understood as governing fields defined on a manifold, rejecting substantivalism amounts to rejecting GTR itself; therefore, a new physics is needed which does not talk about such fields. Earman expressed this sort of attitude when he said that 'replacing substantivalist theories of physics with a radically different alternative obviously requires an act of scientific creativity'.[34]

I disagree: the project of developing an anti-substantivalist metaphysics falls squarely within logic and metaphysics. For let us be clear about the structure of my argument against substantivalism. The argument was simply that General Covariance gives us reason to think that *individualistic* facts about the manifold, such as

> p is mapped to m by the metric field
>
> p instantiates F
>
> p bears R to q

are not part of the fundamental fabric of reality.

But nothing in the above argument had any consequences for general facts, such as

> $(\exists x)(x$ is mapped to m by the metric field$)$
>
> $(\exists x)(x$ instantiates F$)$
>
> $(\exists x)(\exists y)(x$ bears R to $y)$

It is entirely consistent with the above arguments that these sorts of facts are detectable, and indeed necessary for causal explanations. Thus, it is consistent with my arguments that these sorts of facts are real.

Of course, we normally think that general facts are to be *metaphysically* explained in terms of individualistic facts: e.g. that

> $(\exists x)(x$ instantiates F$)$

holds in virtue of some individualistic fact of the form

[34] Earman [5], p. 170.

p instantiates F

and so on. On this way of thinking about the relationship between individualistic and general facts, dispensing with the latter means *ipso facto* dispensing with the former. And once we are being asked to dispense with the existential generalizations that are true of a substantival GTR world, it is easy to think that what is needed is a radically new physical theory.

But wait—we've seen this sort of situation before! Recall Newton: he knew that absolute velocity was redundant to *causal* explanations of the data, but he also believed it was necessary for *metaphysical* explanations of other features (e.g. relative velocity, acceleration) that are themselves necessary components of *causal* explanations of the data. Therefore, he couldn't dispense with absolute velocity without *ipso facto* dispensing with those other features. What was the way out of his dilemma? We now know that the crucial advance wasn't in physics but in mathematics. Once we developed the theory of four-dimensional affine spaces, we learnt how to make good metaphysical sense of relative velocity (and acceleration, and everything else we needed) without the need of absolute velocity. Thus, by drawing on this new mathematics, we could develop a metaphysics of motion that contained everything needed for the physics without the offending structure.

I suggest that we see our anti-substantivalist project in an analogous light. Our task should be to attempt to make good sense of general facts in such a way that we may deny that they hold in virtue of individualistic facts. If we can do this, then we will have an empirically adequate theory of the world—i.e. GTR—with none of the redundant structure enforced on it by substantivalism. This is a project that falls squarely within logic and metaphysics, not physics.

Algorithmic Generalism

If we are to deny that general facts hold in virtue of individualistic facts, then we must either say that they are themselves fundamental, or that they hold in virtue of some other kind of fact altogether.

The former option is not attractive—at least, not at first blush. For it is all very well to be told that

$(\exists x)(x$ instantiates F$)$

is itself a fundamental fact. But we have all grown up being taught to interpret quantifiers as ranging over a domain *of individuals*—in this case points in a substantival manifold. If we are not told how else to interpret them, then one might worry that the only grasp we have of what this fact can possibly *be* is a grasp according to which it holds in virtue of individualistic facts after all.

So I suggest that we explore the second option. And the best way to develop this second option, I think, is to give up our ordinary conception of the fundamental facts as concerning a domain of *individuals* and how they are propertied and related to one another, and instead see the fundamental facts as concerning the nature of a domain of *properties* and how they are concatenated together.

How might we develop a view of this sort? Well, consider the language PL: predicate logic with identity but without constants. Sentences in this language express general facts. Now, consider the space of equivalence classes of formulas of PL. It is well known that this space instantiates a well-defined algebraic structure—the analogue of the familiar Boolean structure of sentential logic. For example, there is a conjunctive operation that takes two equivalence classes—e.g. those containing the formulas Rxy and Sxyz respectively—and delivers their "conjunct"—in this case the class containing the formula (Rxy & Sxyz). A "negation" operation can be defined similarly, and of course with those two operations in hand all other truth-functional operations can then be defined. We might also take an operator which switches the last two variables around—e.g. which takes the equivalence class containing the conjunctive formula above and delivers the class containing (Rxy & Sxzy)—and another operator which moves the last variable to the beginning—e.g. which takes the class containing (Rxy & Sxyz) and delivers the class containing (Rzx & Syxy). All other permutative operations can be defined as successive applications of these two. We can even define operators that identify variables in a formula or ensure that they are distinct. Finally, we can then define an "existential quantification" operator: it would take the class containing (Rxy & Sxyz) and deliver the class containing $(\exists x)$(Rxy & Sxyz). In general, we can think of the set of equivalence classes of formulas of PL as a set of elements with operations defined on them in such a way that every element can be expressed as the result of applying the operators to elements in a base set (where the base set is a proper sub-set of the set of all formulas).

This invites the following view about how the world is structured at the fundamental level. The idea is that for each equivalence class of n-ary formulas there is a corresponding n-place property. Equivalence classes of 0-place formulas correspond to 0-place properties, i.e. states of affairs. Just as the space of equivalence classes of formulas instantiates an algebraic structure, so too do the properties. And just as all the equivalence classes can be expressed as the result of applying the operators to equivalence classes in a base set, we can similarly think of there being a base set of properties from which the others are derived by successive applications of the operators. In particular, the states of affairs might be thought of as being derived from a base set of n-place properties, for $n \geq 1$. Thus, while we normally think of states of affairs as being constructed out of properties *and individuals*, on this view they are constructed solely out of properties. Finally, the fundamental

facts of the world will concern which of the states of affairs obtain. With regards the general facts with which we started, the natural view will be that they hold in virtue of facts of this sort.

In the case of GTR, what would the base properties be, out of which the states of affairs are constructed? This is a tricky question, and depends in part on what one thinks the correct metaphysics of fields is. Thus, if one initially thought, as a substantivalist, that the fundamental facts consisted of facts such as

p is 2 meters from q

where p and q are points in the manifold, then by formulating an anti-substantivalist view in the above way one would naturally be led to the view that one of the basic properties is the relation of *x being 2 meters from y*.

That, at least, is the picture. There remains the project of engineering: can one introduce a (finite?) set of operators that are sufficient to describe a property space with sufficient structure? This project breaks down into two main parts. First, we must determine what structure it is necessary for us to construct. It is tempting to approach this linguistically, by asking what sort of language is required to express all general facts and then asking what the algebraic structure of that language is. If one takes this route, one must first determine whether the required language is PL or perhaps a second-order or infinitary language. And once one has determined what the required language is, one must clearly articulate its algebraic structure. This has been widely explored in the case of PL, but less so with other languages.[35] The second part of this project is to take the resulting algebraic structure and ask whether one can formulate a defendable account of the fundamental structure of the world in its image.[36]

Clearly, there is much work to be done in this area. And if a defendable metaphysics of this sort is not forthcoming, one might conclude that we should be substantivalists after all. In that case, our attitude towards individualistic facts about the manifold should be the same as Newton's attitude to facts about absolute velocity: redundant to *causal* explanations of the data, but necessary for *metaphysical* explanations of other facts that *are* relevant to causal explanations of the data, and therefore part of our all-things-considered best theory of the world. Nonetheless, the take-home message of this paper is that the project of articulating a defendable metaphysics of this sort should be of considerable importance in the philosophy of physics.

[35]See Halmos [9] for one approach to this project in the case of PL. See Henkin, Monk and Tarski [10] for another approach, as well as a discussion of extending their approach to infinitary languages.

[36]The second paper in this dissertation, *Individuals: An Essay in Revisionary Metaphysics*, makes a start at this with regards the algebraic structure of PL.

Substantivalism Revisited?

At this point, it might be objected that the resulting metaphysics is, when properly understood, really a *substantivalist* metaphysics! The objection would be that for any facts that the substantivalist would recognize as fundamental, such as

p is mapped to m by the metric field

or

p is 2 meters from q

where p and q are points in the manifold, the metaphysical view arrived at by the above method will make good sense of the corresponding general facts

$(\exists x)(x$ is mapped to m by the metric field$)$

and

$(\exists x)(\exists y)(x$ is 2 meters from $y)$

And, moreover, the view we've arrived at does not claim that these facts hold in virtue of facts about the distribution of matter, as the relationalist does: rather, the view claims that they hold in virtue of facts about how base properties such as the relation of *x being 2 meters from y* are concatenated together. Thus, the objection is that insofar as the view makes good sense of these sorts of facts without reducing them to facts about the distribution of matter, it is properly called a substantival view.

Of course, if this objection were right, something would be wrong with my characterization of substantivalism. For I stressed from the outset that the substantivalist's defining claim was that the fundamental facts about the world include *individualistic* facts about the manifold. So to respond to the objection I'd have to argue that my characterization of substantivalism is correct.

However, I have a good deal of sympathy for this objection and do not wish to argue about how to characterize substantivalism. Instead, I believe that the question of whether the view outlined above is properly called "substantivalism" is largely a verbal issue. So, to clarify what non-verbal claims I'd like to defend, let us use "individualistic substantivalism" to name the view I called substantivalism at the beginning—i.e. the view that the fundamental facts of the world include individualistic facts about the manifold—leaving it open whether individualistic substantivalism is substantivalism proper. Then the central claim I take myself to have argued for in this paper is the claim that General Covariance gives us good reason to think that *individualistic substantivalism* is false (this is just the Covariance

Substantivalism Link, with 'individualistic substantivalism' substituted for 'substantivalism'). And my dialectical position can now be stated as follows: I disagree with Earman and Norton because I think that the Hole Argument gives us no reason to reject individualistic substantivalism, and I disagree with most substantialists because the view they typically defend is individualistic substantivalism. Whether or not individualistic substantivalism just is substantivalism—and hence whether or not the metaphysical view I've just described should properly be described as substantival—can then be seen as a verbal issue.

5 Conclusion

Earman and Norton famously argued that the fact that GTR satisfies general covariance gives us good reason to think that substantivalism is false. But their argument made a controversial assumption about the nature of metaphysical possibility, and substantivalists—rightly, in my opinion—responded by question the assumption. However, I have outlined a different argument for Earman and Norton's conclusion, one that does not make any assumptions about the nature of metaphysical possibility. I claim that it is at least as convincing as an analogous argument against absolute velocity that is normally considered decisive. And I have gestured at the sort of anti-substantival view that is naturally suggested in the light of this argument.

References

[1] Arntzenius, F. 2004. Gunk, Topology and Measure. *Phil-Sci Archive*.

[2] Belot, G. 2000. Geometry and Motion. *British Journal for the Philosophy of Science* 51: 561–95.

[3] Brighouse, C. 1994. Spacetime and Holes. *PSA: The Proceedings of the Biennial Meeting of the Philosophy of Science Association* Volume One: 117–125.

[4] Butterfield, J. 1988. Albert Einstein Meets David Lewis. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* II: 65–81.

[5] Earman, J. 1989. *World Enough And Space-Time*. Cambridge, MA: MIT Press.

[6] Earman, J. and J. Norton. 1987. What Price Substantivalism? The Hole Story. *British Journal for the Philosophy of Science* 38: 515–525.

[7] Field, H. 1989. Can We Dispense With Space-time? In his *Realism, Mathematics and Modality*. London, UK: Blackwell.

[8] Fine, K. 2001. The Question of Realism. *Philosophers' Imprint* 1: 1–30.

[9] Halmos, P. 1962. *Algebraic Logic*. New York: Chelsea Publishing Company.

[10] Henkin, L., J. D. Monk and A. Tarski. 1971. *Cylindric Algebras, Part I*. Amsterdam: North-Holland Publishing Company.

[11] Hoefer, C. 1996. The Metaphysics of Space-Time Substantivalism. *The Journal of Philosophy* 93: 5–27.

[12] Lewis, D. 1986. *On the Plurality of Worlds*. Cambridge, UK: Blackwell.

[13] Maidens, A. 1992. Review of John Earman's 'World Enough and Space-Time'. *British Journal for the Philosophy of Science* 43: 129–136.

[14] Maudlin, T. 1988. The Essence of Space-Time. *PSA*: 82–91.

[15] Maudlin, T. 1993. Buckets of Water and Waves of Space: Why Spacetime is Probably a Substance. *Philosophy of Science* 60, No. 2: 183–203.

[16] Melia, J. 1999. Holes, Haecceitism and Two Conceptions of Determinism. *British Journal for the Philosophy of Science* 50: 639–664.

[17] Pooley, O. 2005. Points, Particles and Structural Realism. In *The Structural Foundations of Quantum Gravity*, edited by Rickles, French and Saatsi. Oxford, UK: Oxford University Press.

[18] Rynasiewicz, R. 1999. Kretschmann's Analysis of Covariance and Relativity Principles. In *The Expanding Worlds of General Relativity* pp. 431–462, edited by H. Goenner, J. Renn, J. Ritter and T. Sauer. The Center for Einstein Studies.

[19] Schaffer, J. 2009. On What Grounds What. In *Metametaphysics: New Essay's on the Foundations of Ontology*, edited by D. Manley, D. J. Chalmers and R. Wasserman, pp. 347–84. Oxford, UK: OUP.

[20] Schaffer, J. *Forthcoming*. Spacetime the One Substance. *Philosophical Studies*.

[21] Sider, T. *Writing the Book of the World*. Manuscript.

[22] Sklar, L. 1974. *Space, Time and Space-Time*. Los Angeles, CA: University of California Press.

[23] Skow, B. 2005. *Once Upon a Space-Time*. Unpublished PhD dissertation, New York University.

[24] van Fraassen, C and Ismael. 2003. Symmetry as a Guide to Superfluous Structure. In *Symmetries in Physics: Philosophical Perspectives*, edited by K. Brading and E. Castellani. Cambridge, UK: CUP.