Points, particles, and structural realism*

Oliver Pooley
Oriel College, University of Oxford

First draft: 30 July 2004;
this draft: 12 December 2005

> Even if we are able to decide on a canonical formulation of our theory, there is the further problem of metaphysical underdetermination with respect to, for example, whether the entities postulated by a theory are individuals or not…
>
> We need to recognise the failure of our best theories to determine even the most fundamental ontological characteristic of the purported entities they feature…What is required is a shift to a different ontological basis altogether, one for which questions of individuality simply do not arise. Perhaps we should view the individuals and nonindividuals packages, like particle and field pictures, as different *representations* of the same structure. There is an analogy here with the debate about substantivalism in general relativity. (Ladyman 1998)

In his paper "What is Structural Realism?" (1998) James Ladyman drew a distinction between *epistemological* structural realism (ESR) and *metaphysical* (or ontic) structural realism (OSR). In recent years this distinction has set much of the agenda for philosophers of science interested in scientific realism. It has also led to the emergence of a related discussion in the philosophy of physics that concerns the alleged difficulties of interpreting general relativity that revolve around the question of the ontological status of spacetime points. Ladyman drew a suggestive analogy between the perennial debate between substantivalist and relationalist interpretations of spacetime on the one hand, and the debate about whether quantum mechanics treats identical particles as individuals or as 'non-individuals' on

*A version of this paper will appear in Rickles, French and Saatsi (eds), *The Structural Foundations of Quantum Gravity*, Oxford University Press. Please quote from and cite the published version.

the other. In both cases, Ladyman's suggestion is that a structural realist interpretation of the physics—in particular, an *ontic* structural realism—might be just what's needed to overcome the stalemate. The purported analogy between the physics of spacetime points and the physics of quantum particles has been further articulated and defended by Stachel (2002), by Saunders (2003) and by French and Rickles (2003).

The main thesis of this paper is that, whatever the interpretative difficulties of generally covariant spacetime physics are, they do not support or suggest structural realism. In particular, I hope to show that there is in fact no analogy that supports a similar interpretation of the metaphysics of spacetime points and of quantum particles. But the story is not simple, and a certain amount of stage setting is required.

1 What is Structural Realism?

The genesis of contemporary structural realism is well-known: John Worrall (1989) proposed a realist interpretation of science with the intention of doing justice to two opposing arguments. On the one hand there is the 'no miracles' argument: the empirical success of science is held to be miraculous unless it has succeeded in correctly describing the reality behind the phenomena it saves. On the other hand, there is the 'pessimistic meta-induction', most famously, and forcefully, put by Laudan (1981): theories that have been superseded in Kuhnian scientific revolutions are now judged to be radically mistaken in their claims about the reality behind the phenomena, and this *despite* their often having enjoyed exceptional empirical success. It is only rational, therefore, to expect our current theories to suffer a similar fate come the next scientific revolution. Worrall's structural realism is designed to be realist enough so as to do justice to the 'no miracles' intuition and yet agnostic enough to avoid falling prey to the pessimistic meta-induction. In particular, continuity through theory change at the level of form or structure is supposed to license the belief that science is succeeding in characterizing the *structure* of reality (and hence its empirical success is not miraculous) even if science is radically wrong in its description of the fundamental *nature* of reality.

Ladyman asked whether this is epistemology or metaphysics. If it is epistemology, he argued, then there are problems if the doctrine is cashed out in the most obvious way, in terms of a theory's Ramsey sentence. Epistemological structural realists, like John Worrall and Elie Zahar, have since explicitly endorsed the Ramsey sentence strategy, and have sought to defend it against Ladyman's objections.

Recall that the Ramsey sentence T^* of a theory T is formed from T by replacing all the theoretical predicates that occur in the theory with predicate variables and then existentially quantifying. The epistemological structural realist holds that the cognitive content of a theory is fully captured by its Ramsey sentence. Two, related, results appear to pose a problem for such a position.

The implications of the first for structural realism were raised against Russell's structuralist position by Newman (1928); they have been discussed in the context of the recent debate by Demopoulos and Friedman (1985). Essentially the same result lies at the heart of Putnam's paradox. The problem is that it is theorem of set theory and second-order logic that any consistent proposition to the effect that a certain set of properties and relations exist, no matter what structural constraints are placed upon this set, will be true of any domain, provided that the domain has the right cardinality (and the upward Lowenheim-Skolem theorem entails that the only cardinality constraint concerns whether or not the domain is finite). The second result is due to Jane English (1973): any two Ramsey sentences that are incompatible cannot have all of their observational consequences in common. The conclusion to be drawn appears to be that if a Ramsey sentence is empirically adequate, then it is true. It doesn't really make substantive claims about the structure of reality beyond the phenomena since such claims, because they are made only in terms of existential quantification, will always be true (providing that the realm of the unobservable has the right cardinality).

Worrall and Zahar's response (2001) is to stress that the Ramsey sentence is formed by replacing only the *theoretical* predicates with variables; the *observational* predicates remain in place. But it is far from clear that this solves the problem. It certainly means that a consistent Ramsey sentence doesn't simply make a statement about the cardinality of the set of individuals in the world; it does have empirical consequences. But, putting aside Newman's specific target in Russell, the charge was never that structural realism conceived in terms of a theory's Ramsey sentence made only trivial claims about the world. The charge is that it collapses into *empiricism*; that the realist claim that science is in some way latching onto reality beyond the phenomena so as to do justice to the 'no miracles' intuition has been lost.[1]

In fact, in further articulating their response, Worrall and Zahar face a dilemma. They rely on a distinction between observational and theoretical predicates, but let us ask how this might relate to a distinction between (directly) observable and unobservable, or inferred, *entities*. Grant such a partition of the domain of the world, and one has two possibilities: (1) an observational predicate is one that is satisfied only by observable entities or (2) observational predicates can be satisfied by unobservable entities.[2] If one opts for (1), then clearly the Newman problem applies in full force to the question of which *unobservable* entities fall in the extensions of the theoretical predicates. If one opts for (2), then one can ask with what right does the structural realist claim to be able to keep fixed the *extensions* of the observational

[1]For a careful defence of the claim that the truth of a theory's Ramsey sentence is equivalent to the combination of its empirical adequacy together with a cardinality constraint, see Ketland (2004).

[2]The phenomenalistic strain in Zahar's philosophy might suggest that he is committed to (1).

predicates, given that such predicates apply to unobserved entities whose existence is conjectural.

And even if we grant that the extensions of the observational predicates are to be held fixed, and that there is no subdomain of the world to which none of them apply, the structural realist is still not entitled to assume that holding true the Ramsey sentence will give a fix on the extensions of the theoretical properties and relations over which it quantifies that is sufficient to satisfy realist intuitions. A complete fix will be achieved if the original theory is categorical—if all of its models are isomorphic—but (formalizations of) real scientific theories, even together with reports of all relevant empirical data, will not, of course, be categorical.

In fact, Worrall and Zahar seem to be prepared to acknowledge that there is no more to the content of a Ramsey sentence than all of its consequences that contain only observational predicates. The reason why they do not consider that this means that their version of structural realism collapses into empiricism is that they distinguish between a theory's observational content—which they claim must be empirically *decidable*, if it is to count as genuinely observational—and the more inclusive set that also includes all of the empirical generalizations that the theory entails. They concede that the theory's Ramsey sentence itself might be amongst the latter. But even if to endorse empirical generalizations really is to "go against the canons of even the most liberal version of empiricism" (2003, p. 241), by itself it hardly amounts to scientific realism, structural or otherwise.[3]

Let me conclude this discussion of the problems facing epistemological structural realism by mentioning one further worry. It is clear that the Ramsifying structural realist has to tread a very fine line. The Ramsey sentence must be held to give us enough of a fix on the extensions of theoretical predicates so as to avoid the Newman problem. But to achieve too strong a fix would undermine the original motivation for structural realism. The reason is that the Ramsey sentence refers to, and quantifies over, exactly the same entities as the original theory. As French and Ladyman press:

> If the meta-induction is a problem about lack of continuity of reference then Ramsifying a theory does not address the problem at all. (2003, p. 33)

[3]Although it is perhaps close to the version of structural realism that I take Ladyman, in particular, to advocate. This has no use for a notion of reality 'behind' the phenomena and involves only the 'minimal metaphysical commitment' to "mind independent modal relations between phenomena (both possible and actual)" (2003, p. 46). Where Ladyman would differ from Worrall and Zahar, presumably, is in claiming that these modal relations are not "supervenient on the properties of unobservable objects and the external relations between them, rather this structure is ontologically basic" (*ibid.*, p. 46; *cf.* also Ladyman 1998, p. 418).

2 Underdetermination.

The idea of an *ontic* structural realism can provoke a definite sense of unease, for it hard not to worry that something akin to mystery mongering is taking place. It is one thing to claim that our knowledge of the unobservable realm is limited to structural knowledge, that all we can know about unobservable objects are their structural properties. It is quite another thing to claim that *all there is* beyond the phenomena is structure. What can this claim *mean*?

It would, however, be quite unfair to accuse the most prominent defenders of ontic structural realism of mystery mongering. In order to get beyond the slogans and metaphors and to arrive at a characterization of, if not ontic structural realism itself, then of one of its expected achievements, it is time to consider the problem of underdetermination. This problem is central to Ladyman's positive case for OSR: whatever else it is, OSR is intended to dissolve various metaphysical disputes centring on underdetermination. It is supposed to provide a new metaphysical perspective with respect to which certain troublesome, apparently irresolvable, choices simply do not arise.

Ladyman distinguishes between two quite distinct ways in which a *single* theory might be said to be empirically underdetermined.[4] The first type arises due to the existence of different formulations of a single theory, formulations that can suggest radically different ontologies if interpreted realistically. This particular problem for the realist was stressed by Jones (1991). Call the type of underdetermination involved *Jones underdetermination*. Ladyman calls the second type of underdetermination *metaphysical underdetermination*. Roughly speaking, this type of underdetermination arises because of the existence of alternative (realist) interpretations of a single *formulation* of a theory, interpretations that again are supposed to involve radically incompatible ontologies. According to Ladyman the underdetermination between individual and non-individual interpretations of the quantum mechanics of identical particles, and between substantivalism and relationalism is of this type. Another nice example is the underdetermination between two rival

[4]Traditional 'underdetermination of a theory by data' might be thought to involve two quite distinct theories which are nevertheless empirically equivalent, either with respect to all the observational evidence so far (weak underdetermination), or in principle (strong underdetermination). The types of underdetermination that Ladyman highlights, as will been seen, involve the interpretation of a *single* theory. In what follows I simply set aside the traditional problem, and use the term *empirical underdetermination* as a general term for the two *special* types of underdetermination that Ladyman discusses.

Note that the division is not quite as clear cut as this summary suggests. Someone might argue, on the basis of an instance of strong traditional underdetermination that in fact we have two formulations of a single theory, not two theories. Conversely, someone might argue that Jones underdetermination (defined below) really shows us that what we were inclined to treat as two formulations of a single theory should in fact be regarded as two theories.

conceptions of fields. Fields can be viewed either as substantival entities in their own right, with infinitely many degrees of freedom associated with their infinitely many point-like parts, or they can be viewed as consisting in the instantiation of a pattern of properties by spacetime points.

It is useful at this point to mention a parallel problem in the philosophy of mathematics. In 'What Numbers Could Not Be' (1965), Benacerraf highlighted the following difficulty facing anyone who would identify numbers with sets: there are many proposals, they are incompatible so at most one can be correct, but it seems that no cogent reason can be given for preferring one over another. If numbers are sets then which sets they are seems to be underdetermined. Defenders of structuralist views of mathematics view this difficulty that faces those who would take the reducibility of arithmetic to set theory as disclosing that numbers are really sets as one of the best supports for their philosophy. What all the set-theoretic constructions share is the same structure, and this structure is all that matters. Does this parallel reflect favourably on the ontic structural realist's attitude to empirical underdetermination?[5]

For cases of empirical underdetermination to support ontic structural realism, Ladyman had better be correct in saying that traditional realism goes beyond a commitment to structure precisely in terms of metaphysical commitments that are underdetermined by the evidence for them (Ladyman 1998, p. 418). On the face of it, and without being told what a commitment only to structure might involve, the claim does not look plausible. An obvious problem is that, on the most straightforward characterizations of structure (e.g. a set-theoretic one), most cases of different formulations of a theory will involve *different* structures.[6] Consider a model of a theory of Newtonian gravitation formulated using an action-at-a-distance force and an empirically equivalent model of the Newton-Cartan formulation of theory. There is no (primitive) element of the second model which is structurally isomorphic to the flat inertial connection of the first model, and there are no (primitive) elements of the first model which are structurally isomorphic to the gravitational potential field, or the non-flat inertial structure of the second. Clearly a more sophisticated notion of structure is needed if it is to be something common to models of both formulations of the theory.

The claim that the structural realist might have the resources to be able to identify something beyond the empirical that is in common to different formulations of a single theory, and that he thus might be able to dissolve various interpretative problems by transcending the root underdetermination, is supposed to get

[5]I return to the parallel in Section 3. A crucial question will be whether the parallel supports the elimination of objects altogether.

[6]A fact which undermines, I think, the parallel between empirical underdetermination (at least of the Jones variety) and the underdetermination that motivates structuralist positions in the philosophy of mathematics.

some of its plausibility from the case of Schrödinger's and Heisenberg's original rival formulations of quantum mechanics. Through the work of Weyl and others, these formulations were soon recognised to be *different representations* of a single, mathematical *structure* in which states of a system correspond to rays in a Hilbert space, and observables correspond to operators (of the appropriate sort) that act on this space.[7] But there are at least three reasons to be sceptical that this example alone lends support to the idea that underdetermination in general, and Ladyman's metaphysical underdetermination in particular, motivate a radical ontic structural realism.

First, it is not clear that the ontic structural realist has a story to tell about this example. As is well known, whether anyone has come up with a truly successful *realist* interpretation—structural or otherwise—of the standard formulation of non-relativistic quantum mechanics that subsumes both Schrödinger's and Heisenberg's original formalisms is a controversial issue. Two of the potential candidates—GRW collapse theory and de Broglie–Bohm pilot wave theory—break the unifying picture by preferring one basis with respect to which either genuine collapse occurs, or with respect to which the true beables are defined. And if more that one of the various interpretative options ultimately survives the many criticisms they all face, it would seem that quantum mechanics remains beset by underdetermination, albeit of a very different type to that involving wave versus matrix mechanics.

Second, this one example gives us little reason to suppose that whatever was achieved in this particular case will be, or even can be, repeated for other instances of Jones underdetermination. Consider the underdetermination that exists (relative to certain solutions) between Julian Barbour's Machian 3-space approach to general relativity (where the fundamental ontology consists of instantaneous 3-spaces and does not involve any primitive temporal notions), the traditional curved-spacetime formulation, and formulations involving spin-2 fields on a flat (or at least fixed) background spacetime. Here we have a clear case of different formulations of a theory that are associated with *prima facie* incompatible ontologies. A structural realist dissolution of this problem requires an explicit characterization of a mathematical framework that stands to each formalism as the abstract Hilbert space formalism of quantum mechanics stands to Schrödinger's wave mechanics and Heisenberg's matrix mechanics. Note that it is not enough that we have a good understanding (as we do) of the various mathematical relationships that exist between the formalisms. Ladyman's structural realist needs a single, unifying framework, which she can then interpret (in terms of an as-yet-to-be-articulated metaphysics of structure) as corresponding more faithfully to reality than do its various realist representations.

[7]In fact, Heisenberg's matrix mechanics and Schrödinger's wave mechanics were *not* strictly equivalent; see Muller (1997a, 1997b)!

I am not optimistic that any such development is in the offing. It seems more likely that theoretical advances will favour one formulation over the others. String theory's triumph would, in many senses, vindicate the spin-2 picture. The success of loop quantum gravity (or of a variant, provided with the right sort of interpretation) could vindicate Barbour's advocacy of 3-space concepts over spacetime concepts.[8] The idea that underdetermination associated with different formulations is to be transcended by a more general framework with respect to which the different formulations are seen as different representations of a single, underlying reality might look suspiciously like an unwarranted generalization from a single, special case.

Third, the Heisenberg–Schrödinger example involves Jones underdetermination. However, the type of underdetermination that is supposed to be involved in the debates between the substantivalist and the relationalist, and between advocates of the "individuals interpretation" of quantum particles and those who advocates a particles-as-non-individuals interpretation, is what Ladyman calls *metaphysical* underdetermination. Here it might seem more plausible that some interpretative stance according to which the rival viewpoints are merely different representations of the same reality will be possible. (But it is important to stress that those sympathetic to an ontic structural realism have yet to provide a positive characterization of any such position, so far they have only told us what the position is meant to achieve; see pp. 11–13 below.) But equally, one might wonder whether the underdetermination in question is one that should genuinely trouble the realist.

In fact, a genuine underdetermination between relationalism and substantivalism *is* one that should trouble the realist. This is because such an underdetermination *would* be an instance of Jones underdetermination. However, as things stand, there simply is no such underdetermination. The standard formulations of general relativity are straightforwardly substantivalist in that the metric field is (a) taken to represent a genuine and primitive element of reality and (b) most naturally interpreted as representing spacetime structure.[9] Now some believe that the hole argument calls this picture into question. But these same people also typically suggest that an alternative formulation, that would correspond to a genuinely relationalist world picture, should be sought (e.g., Earman 1989, Ch. 9). I agree that relationalism needs different physics (or at least a different formulation of the physics).[10] But I disagree that the formulation of GR that the realist naturally interprets along sub-

[8]Such discrimination between alternatives by theoretical advances gets some support from the history of physics, although the fact that underdetermination often reappears in a new guise means that the realist should not take too much comfort from this state of affairs.

[9]The second of these claims is contested by some (Earman and Norton 1987; Rovelli 1997). I say more about it below (pp. 20ff).

[10]Barbour's 3-space approach to GR constitutes a genuinely distinct interpretation, but not a relationalist one, since the fundamental ontology is substantival space (not spacetime).

stantivalist lines is in trouble because of the hole argument, for there are different interpretative options available *within* the substantivalist camp.

And it turns out that this is the true location of Ladyman's metaphysical underdetermination. The opposition that he, and Stachel, characterize as between relationalism and substantivalism is really an opposition between *haecceitist* and *anti-haecceitist* substantivalism.[11] I will explain later what I mean by these terms, and why I believe that there is no real contest: anti-haecceitism is the clear winner. But even if one thought that there was a genuine choice to be made, and that interpreting the physics realistically failed, by itself, to make the choice, it is not clear why this should trouble the scientific realist. For, as we will see, there is a sense in which haecceitist substantivalism is simply an extension of anti-haecceitist substantivalism. Anti-haecceitist substantivalism represents a realist core position which it may or may not be correct to supplement. If this is the only choice to be made, it hardly constitutes an interesting threat to the scientific realist's belief in the existence of spacetime points.

Things are otherwise with the quantum mechanics of identical particles. For the benefit of those already *au fait* with the terminology, it turns out that, while the physics of identical particles strongly suggests anti-haecceitism, anti-haecceitism by itself does not suffice to explain all the peculiarities of the physics of quantum particles. The difference between the two cases is traceable to a difference in the physics. Perhaps unsurprisingly, the way in which *classical* GR is diffeomorphism-invariant is rather different from the way in which the *quantum mechanics* of identical particles is permutation-invariant.

The disanalogy highlights a sense in which it is misleading to present the quantum mechanical case as an instance of underdetermination between two realist interpretations, if the intended implication is that the two interpretations are equally viable. While French (1989; French and Redhead 1988) may have clearly demonstrated that the individuals interpretation of QM particles exists in logical space, it is not really a serious contender.[12] Equally, the non-individuals interpretation, if it truly accommodates the phenomena, is a more radical position than anti-haecceitist substantivalism.[13]

[11]The positions have been labelled straightforward and sophisticated substantivalism by Belot and Earman (1999; 2001).

[12]Ladyman's claim that there is "much dispute about whether or not quantum particles…are individuals" (1998, p. 419) is unconvincing if to treat particles as individuals is to adopt the interpretative option delineated by French, and French and Redhead. I am not, of course, denying that exactly how one *should* conceive of quantum particles is a highly disputed question; rather I am only claiming that there is a fair degree of consensus that one should *not* conceive of quantum particles in certain ways. French (1998, p. 112, fn. 62) himself states that it's not true that "anything goes," but I understand him to take viewing the particle labels of the standard tensor product Hilbert space formalism as naming individuals as a genuine interpretative option. I don't believe it is.

[13]As Teller's discussion in his (2001) makes clear. As will become apparent, although I agree with

In the next section I briefly consider how far OSR's defenders have gone in attempting to characterize the position. Before doing so, I wish to raise two worries connected specifically to underdetermination. The first is that one might worry that dissolving the underdetermination is not desirable. Recall that I mentioned above the way in which the various alternative formulations of GR were linked to quite distinct attempts to advance beyond that theory. Having these alternatives in play might therefore serve a vital heuristic role in theoretical advance. Of course, from the perspective of OSR, the various formulations still exist; they are just now understood as different representations of the same structure. But perhaps, for advances to take place, it is important that the different formulations are considered to be genuinely distinct and exclusive alternatives. And perhaps, from the perspective of an advance, that the subsequently favoured alternative could be unified with the others in a single structure will seem like a happy accident; the overarching framework will appear to have a secondary status, rather than a fundamental one.

The second worry is that a radical structural metaphysics might make the underdetermination worse. Assuming such a metaphysics is possible (a big assumption), then *if* one adopts it, one will view the previous alternatives as different representations of the structure that one claims is fundamental. But is the fact that this is how it looks from the perspective of OSR, enough to commend adopting that perspective? It seems likely that every side of the original underdetermination will be able to explain the other side's worldview. For example, if one believes in a dynamical spacetime connection, one can explain why things are as if spacetime were flat and gravity were a universal force. But why not also expect that the red-blooded realists will be able similarly to explain away the structural realist's perspective, just as GRW theory and Bohm theory can (or are supposed to be able to) explain the success of orthodox quantum theory? The defender of OSR can perhaps wield Ockham's razor, but the dialectical problem here, for the structural realist, is that this is exactly the type of consideration that might also favour one traditional realist interpretation over another. And since we do not yet have the structural realist metaphysics, or any guarantee that such a thing is conceivable, if we will end up having to wield Ockham's razor in order to vindicate it, one might wonder why one should go looking for it in the first place.

Teller that the interpretative options that the substantivalist has in the face of the hole argument are not applicable to the case of identical QM particles, I differ with him both about how to characterize the substantivalist options, and about the difficulty of providing a successful realist interpretation of the quantum mechanics of identical particles.

3 What *is* Ontic Structural Realism?

If it has been made clear what the relation of ontic structural realism to the problem of underdetermination is supposed to be, what has been said by way of a *positive* characterization of a position that can do the job? The answer is, at this stage, not much, but in a recent paper French and Ladyman (2003) seek to further articulate their vision of OSR.

I have already mentioned Ladyman's claim that traditional realism goes beyond commitment to structure precisely in commitments that are underdetermined by the evidence. This suggests the tactic of attempting to identify exactly which elements of the realist's metaphysics are responsible for the underdetermination. French and Ladyman have a clear view:

> The locus of this metaphysical underdetermination is the notion of an object so one way of avoiding it would be to reconceptualise this notion entirely in structural terms. The metaphysical packages of individuality and non-individuality would then be viewed in a similar way to that of particle and field in QFT, namely as two different (metaphysical) representations of the same structure. (2003, p. 37)

The basic point in this quote, that OSR is to offer a perspective from which viewpoints previously taken to be alternatives are seen as representations of the same structure, has been well rehearsed above. What is new in this quote is the claim that an elimination (or, at least, a reconceptualization) of objects is the key. A little later, French and Ladyman are more specific:

> We regard the ontic form of SR as offering a reconceptualisation of ontology, at the most basic metaphysical level, which effects a shift from objects to structures. Now, in what terms does such a reconceptualisation proceed? This hinges on our prior understanding of the notion of an 'object' which has to do…with the metaphysics of individuality. Given the above metaphysical underdetermination, a form of realism adequate to the physics needs to be constructed on the basis of an alternative ontology which replaces the notion of object-as-individual/non-individual with that of structure in some form. (ibid.)

Before asking what we are to make of this talk of reconceptualizing objects in terms of structures, it is worth raising the following worry. In the previous section, a distinction was drawn between traditional and metaphysical underdetermination. The present proposal appears to be addressed only to the latter type (and then, rather specifically, to that involving the quantum physics of identical particles and, more controversially, the interpretation of spacetime points). Surely the notion of an object, and an object's individuality, is not the root cause of the underdetermination between, for example, spacetime formulations of GR and Machian geometrodynamics.

Putting aside this worry, let us ask what replacing the notion of objects with that of structure comes to. It is worth recalling the parallel with structuralist views in mathematics. Benacerraf's own view was that his argument to the effect that numbers could not be sets extended to support the conclusion that numbers could not be objects at all. But this point of view is not shared by many contemporary 'non-eliminative' mathematical structuralists, who hold that one can agree that a mathematical object is the very object that it is in virtue of its occupying its particular place in the relevant mathematical structure, without in any sense eliminating it as a genuine object.

One version of the non-eliminative view is Stewart Shapiro's *ante rem* structuralism, so-called after the analogous view concerning universals. This view takes "structures, and their places, to exist independently of whether there are any systems of objects that exemplify them" (1997, p. 9). The indifference to whether there exist objects exemplifying the structures should not be taken to suggest an indifference to the existence of *mathematical* objects. Rather the mathematical objects are to be understood in terms of the 'places' in the structures; although they enjoy a somewhat secondary ontological status to the structures in which they are places, their existence is not being denied.

The idea of the independent existence of structures suggests an obvious comparison, viz. with the view that physical objects are nothing but bundles of collocated properties ('bundle theory'). Most variants of such a view are held to face the decisive objection that they entail an intolerably strong version of the principle of the identity of indiscernibles. But if one takes relations, and the structures that they form, seriously, one has the resources to frame a sophisticated bundle theory that entails only a relatively weak form of the identity of indiscernibles. According to such a view, it must always be possible to make out numerical diversity in relational terms that do not presuppose identity and difference, but this allows that two objects may nevertheless satisfy exactly the same open sentences with just one free individual variable.[14] Is this all that French and Ladyman have in mind when they talk of a structural reconceptualization of objects?

Although it is suggested by their quoting, apparently with approval, Cassirer's talk of electrons as the " 'points of intersection' of certain relations", and of entities being 'constituted' in terms of relations, there are two reasons why I doubt that it is what they intend. First, the new structural metaphysics was supposed to transcend

[14]The requisite formal treatment of identity goes back to Hilbert and Bernays (1934). It is advocated by, e.g., Quine (1986, pp. 63-4), and recently has been systematically applied to issues in the philosophy of physics by Simon Saunders (2000; 2003).

questions of individuality/non-individuality. But the obvious interpretation of the bundle theoretic proposals is that they *do* yield individuals: determinately numerically distinct particulars, albeit ones whose ontological status, and individuality, is secondary to, and dependent upon, that of properties and relations. The second reason, which is related to the first, is that there is no reason for a bundle theorist to have a problem with standard logic and set theory. Standard logic and set theory presuppose the existence of individuals that are determinately numerically distinct, but they do not presuppose that the individuality of these individuals is independent of the properties and relations that predicates can express, or of the sets that the individuals can form. And yet French and Ladyman do see standard logic as a barrier to articulating their view:

> How can you have structure without (non-structural) objects? Here the structuralist finds herself hamstrung by the descriptive inadequacies of modern logic and set theory which retains the classical framework of individual objects represented by variables and which are the subject of predication or membership respectively (cf. Zahar (1994)). In lieu of a more appropriate framework for structuralist metaphysics, one has to resort to a kind of 'spatchcock' approach, reading the logical variables and constants as *mere placeholders which allow us to define and describe the relevant relations which bear all the ontological weight*. (2003, p. 41)

Talk of "mere placeholders" might suggest that their view is no more radical than the bundle theoretic suggestion, but in a later footnote they are more explicit about what they perceive to be the inadequacies of set theory, and the current unavailability of anything that serves the ontic structural realist's needs:

> [B]oth of these modes of representation – group theory and set theory – presuppose distinguishable elements, which is precisely what we take modern physics to urge us to do away with. If we are going to take our structuralism seriously, we should therefore be appropriately reflective and come up with thorough-going structural alternatives to group theory and set theory... [Krause's attempt to construct a 'quasi-set theory' (Krause 1992)] insofar as [it] is based on *objects* which do not have well defined identity conditions... represents a formalism of one side of our metaphysical underdetermination, rather than a structuralist attempt to avoid it altogether. What is needed is the construction of a fundamental formalisation that is entirely structural; we shall leave this to future works or future (cleverer) philosophers. (2003, p. 52; my emphasis)[15]

[15]Note that talk of 'distinguishable elements' is ambiguous. I stated above that standard logic and

It is time to examine more closely whether the realist's commitment to objects really does lead to an objectionable form of underdetermination, as French and Ladyman maintain.

4 Objects.

Ladyman claims that "traditional realism *does* involve acceptance of more than the structural properties of theoretical entities" (1998, p. 418). The realist's additional, metaphysical commitments are underdetermined by the empirical evidence but are supposed to be of such interpretative importance to the realist that our best theories fail to determine "even the most fundamental ontological characteristics of the purported entities they feature" (ibid., pp. 419–20). What are the realist's additional commitments? Are they as central to a traditional realism as Ladyman claims? Should we be troubled that facts concerning them are underdetermined by the physics itself?

The problem is supposed to concern whether the fundamental entities (spacetime points or quantum particles) that the realist posits are *individuals*, but what is it for an object to be an individual? This is, of course, a question that has been much discussed in this context (van Fraassen 1991; French 1998; Teller 1998; French and Rickles 2003) but I have to confess myself unhappy with the course that these discussions sometimes take, and with many of their presuppositions. For example, the problem is sometimes approached by suggesting that if the objects are individuals, then there is a limited range of options for understanding what their individuality consists in:

(i) objects might be individuated in virtue of their possessing some sort of haecceity,

(ii) the individuality of the substance or matter of an object might be held to account for the object's individuality,

(iii) objects might be individuated in terms of their spatiotemporal location, or

(iv) objects might be individuated in terms of the properties they possess (and perhaps, also, the relations they stand in).[16]

set theory presuppose the *determinate distinctness* of its elements. I claim that logic and set theory presuppose that their elements are distinguishable in no stronger sense than this. The diffeomorphism invariance of GR in no way suggests that spacetime points fail to be distinguishable elements in this sense, but the quantum physics of identical particles does threaten the view that fundamental particles are numerically distinct.

[16] *Cf.* French and Rickles (2003, p. 223).

An object would then, presumably, be a 'non-individual' if none of these ways of understanding it as 'possessing individuality' were available.

Talk of individuality, and of individuation, in these contexts is highly obscure. A more promising approach to the question is to focus on questions of identity and non-identity. First, one can ask, within the context of a *single* situation, about the *numerical distinctness* of the objects featuring in the situation. Here there are two obvious questions: (ND1) are the objects determinately numerically distinct? And, if they are, (ND2) what (if anything) confers, or is the ground of, this numerical distinctness? Second, one can ask about the *trans-situation identity* and distinctness of objects in *different* situations (for example, situations that obtain at different times, or different counterfactual possibilities). Here some of the possible questions are:

(TT) if an object that exists at one time is the same object as a particular object that exists at another time, does anything account for, or ground, this identity, or is it a primitive fact?

(TW) If an object that exists in one possible situation is the same object as an object that exists in another possible situation does anything account for, or ground, this identity, or is it a primitive fact?

Related to these questions is a rather more specific question:

(P) are the objects such that there can be two, genuinely distinct, situations which differ solely in terms of a permutation of some of the objects involved?

When the situations in question are distinct possible worlds, (P) becomes a question concerning *haecceitism*. As I will use the term, haecceitism is the position that there are pairs of genuinely distinct possible worlds that differ solely in terms of a permutation of some of the objects that exist in[17] both possible worlds.[18] Anti-haecceitism (concerning a class of objects) is simply the denial that two possible worlds can differ solely in terms of a permutation of objects of that type. When the situations in question are understood as obtaining at different times within a single world, (P) is not a question about haecceitism.

Now clearly one's answers to the questions (TT) and (TW), which concern trans-temporal and trans-possibility identity respectively, will have a bearing on

[17]Read 'exist in' in such a way that it is compatible with counterpart-theoretic approaches to trans-world identity.

[18]In this I comply with a usage that is standard in much recent philosophical literature. It is due to David Kaplan (1975), and is, e.g., explicitly followed by Lewis (1986, Ch. 4). As will become clear, it should not be confused with a commitment to *haecceities*, however these are to be understood. Of course, belief in a certain robust type of haecceity might license haecceitism in the modal sense meant here.

how one answers the corresponding versions of (P). If one thinks that there can be primitive facts concerning trans-temporal, or trans-world, identity, then it seems that one will be committed to the view that a permutation of objects is alone sufficient to yield genuinely different situations of the type in question.[19] The converse, however, does not hold, at least in the case of trans-temporal identity. One can hold that the state of a system at two different times differs solely by a permutation of the system's constituent objects without holding that the objects' trans-temporal identities are brute matters of fact, for one might think that the trans-temporal identities are determined by various trans-temporal relations that do not supervene on the intrinsic states of the system at the two times in question. The most obvious possibility, of course, is that the identities are underwritten by the continuity of the objects' trajectories. It seems plausible, however, that no such relations are available to ground haecceitistic differences in the absence of primitive trans-world identity.

Now questions of individuality appear originally to have entered discussions of the interpretation of many-particle quantum mechanics in terms of question (P), though it is perhaps not totally clear whether the trans-world or trans-temporal version was in question. Crudely put, the assumption of equiprobability together with the answer "Yes" to (P)—states which differ solely over a permutation of the objects involved are genuinely distinct states—yields Maxwell-Boltzmann statistics (see page 37 below). Answering "No", therefore, looks like a way of accounting for quantum statistics. It is quite possible that all that some of the founding fathers of quantum mechanics—such as Born, Heisenberg and Pauli—meant by quantum particles lacking individuality was that (some version of) (P) should receive a negative answer. As we will see in Section 8, denying that a permutation yields a distinct situation is not by itself sufficient to explain the full peculiarities of the quantum mechanics of identical particles.

Let us return to the question of the metaphysical commitments of traditional, object-positing realism. We have listed four putative accounts of individuality, (i) to (iv), and we have reviewed a set of questions—(ND1), (ND2), (TT), (TW) and (P)—concerning object identity. I wish to urge the following point of view: it is sufficient for a certain class of objects to qualify as individuals that (ND1) gets answered "Yes"—that in a given situation there are facts of the matter about the objects' numerical distinctness. In particular, I claim that answering "No" to (P)—especially in its modal version, does not impugn the objects' status as genuine, substantial, individuals.

Here I disagree with Paul Teller. He is concerned to identify a minimalist sense of haecceity that is connected with the idea that a particular subject matter (e.g.,

[19]In fact things are not quite so straightforward: primitive trans-world identity can perhaps be combined with a denial of purely haecceitistic differences if it is coupled with a strong enough essentialism; *cf.* Maudlin's response to the hole argument (1989; 1990).

that of a particular physical theory) concerns *things*. I take it that an entity should be counted as an individual just if it is a thing and has a haecceity in some properly minimalist sense. Teller proposes three "tests" for whether a subject matter includes (minimalist) haecceities:

1. Strict identity:…there is a fact of the matter for two putatively distinct objects, either that they are distinct or, after all, that they are one and the same thing.

2. Labeling:…the subject matter comprises things that can be referred to with names directly attaching to the referents; that is…things can be named, or labeled, or referred to with constants where the names, labels, or constants each pick out a unique referent, always the same on different occurrences of use, and the names, labels, or constants do not function by relying on properties of their referents.

3. Counterfactual switching:…the subject matter comprises things which can be counterfactually switched, that is just in case *a* being *A* and *b* being *B* is a distinct possible case from *b* being *A* and *a* being *B*, where *A* and *B* are complete rosters of, respectively *a*'s and *b*'s properties in the actual world. (1998, p. 121)

"Strict identity" corresponds to a positive answer to (ND1); "counterfactual switching" corresponds to a positive answer to the possible worlds version of (P). Now although Teller does not claim that the three tests are necessarily "different ways of getting at the same idea" (1998, p. 122), he clearly thinks that the connections are close enough for the tests to be usefully grouped together.[20] But if—as I claim—determinate intra-situation distinctness has no implications for trans-situation identity, then such grouping can only lead to confusion. It is only the denial of determinate intra-situation identity and distinctness that threatens the individuality—the genuine objecthood—of the putative entities in question. However, if questions of intra-situation distinctness and trans-situation identity are not distinguished, one might erroneously infer non-individuality from the denial of (primitive) trans-situation identity. Keeping the two notions distinct is crucial if

[20]In fact, in his (2001, p. 377), Teller claims that (3) is a consequence of (1) and (2). It is true that (1) and (2) allow us to form descriptions which, if they both describe possibilities, describe possibilities that involve counterfactual switching (i.e., that differ merely haecceitistically). But it does not follow from this fact alone that such possibilities *exist*. Note that "labeling" will only fail to be possible even though strict identity applies when *no* reference to the individuals is possible. But even in such situations 'reference' via variables is possible. Compare our ability to talk about abstract symmetrical structures, e.g., Black's sphere world. In such cases labels or names can be used in a generic sense, but do not refer to one, rather than the other, of the objects related by the symmetry (*cf.* Teller 2001, p. 367).

one is to understand the difference between difficulties raised by the diffeomorphism invariance of classical GR, and by the (anti)symmetrization of the quantum states of identical particles. In the next section we will see that the former only has implications, via the hole argument, for trans-situation identity, whereas the latter, at least according to some, threatens determinate numerical distinctness.

I claim that the realist, in positing objects as individuals, is committed to the determinate intra-situation numerical distinctness of the entities posited. To posit 'non-individual' objects, if sense can be made of this, would be to posit a class of entities whose numerical distinctness was somehow not determinate.[21] These two positions—the positing of a class of determinately distinct objects and the positing of a class of non-individual objects respectively—represent *core* realist positions. They do not go beyond an acceptance of the 'purely structural' properties of the entities in question (what properties could be more structural that the determinateness or otherwise of numerical distinctness?). And to go so far, and no further, is hardly an '*ersatz* form of realism', but rather a realism worthy of the name.

If the phenomena covered by a particular theory really were indifferent between these two, radically different metaphysics, then one would have an interesting case of genuine metaphysical underdetermination. Perhaps this is what we face in the quantum mechanics of identical particles. It is inappropriate, however, to regard the spacetime points of diffeomorphism-invariant generally relativistic physics as non-individuals in this sense; the physics simply gives us no scope to do so. In any case, the alleged metaphysical underdetermination discussed in the previous sections was not solely an underdetermination between these two core realist positions. Rather it concerned *further* metaphysical commitments to which the realist may or may not sign up.

Of course, the realist *is* perfectly entitled to sign up to other metaphysical commitments, and he may well endorse particular answers to questions such as (TW) and (P). In particular, it is clear that the four ways of understanding 'individuality', (i)–(iv), mentioned above might well be held to underwrite particular answers to such questions. The haecceities of (i), for example, might be thought both to ground numerical distinctness and to underwrite haecceitistic differences; the properties and relations of (iv) might be thought to ground numerical distinctness in such a way as to rule out haecceitistic differences. A 'bare particular' view of objects is a specific example of a metaphysics in line with (i), or perhaps (ii); the sophisticated bundle theory of the last section is a specific example (though not the only example) of a metaphysics in line with (iv).

To sum up, there are two points that I wish to make at this juncture. First any 'underdetermination' at this level surely should not trouble the scientific realist. It looks as if a quite general, metaphysical debate is being played out in the context

[21] And some think that sense *can* be made of this: see Dalla Chiara *et al.* (1998).

of the entities of physics. Why should the fact that these two options arise in the context of the interpretation of, e.g., spacetime points trouble the scientific realist? In this context, realism is a commitment to the existence of spacetime points, as determinately distinct, substantial individuals. It is hardly a failure of our best theory of them that it fails to determine whether they are bundles of properties, or bare particulars. In fact, it is not clear that the theory *is* indifferent to the choice, for the hole argument shows precisely that conceiving of spacetime points as something akin to bare particulars has the unwelcome consequence of a thoroughgoing indeterminism.

The second point is that it has long been recognised that the choice between the bare particulars view (according to which relata, and their numerical diversity, are ontologically prior to their properties and relations) and a (sophisticated) bundle theory (according to which relations are ontologically 'prior' to, and 'constitute' their relata) presents us with a false dichotomy. A 'no priority view' seems to many to be far more plausible than either. One can endorse the structuralist claim that the numerical diversity of certain objects is grounded in their being situated in a relational structure without reducing these objects to the properties and relations themselves. Equally, and conversely, one can claim that facts of numerical diversity are not so grounded, without going on to claim that they nevertheless *are* grounded by mysterious haecceities or substrata. Instead, one can just take such facts as primitive and as in need of no further metaphysical 'explanation'. It is time to see how these issues play out in the context of the debate concerning substantivalism.

5 Sophisticated Substantivalism.

In debates concerning the nature of spacetime, *substantivalism* is the simply the view that spacetime and its pointlike parts exist as fundamental, substantial entities. This realist view would appear to be what follows from a fairly literal-minded reading of the mathematical formalism of the standard formulations of relativistic physics. For example, the models of general relativity are typically taken to be n-tuples of the form $\langle M, g, \phi_1, \ldots, \phi_{n-2} \rangle$ that satisfy Einstein's field equations. M is a 4-dimensional differential manifold and g is a pseudo-Riemannian metric tensor. M and g, taken together, are naturally understood as representing substantival spacetime: the elements of M represent spacetime points, and g encodes the spatiotemporal relations in which they stand. The fields ϕ_i represent the material content of spacetime.

This simple story is supposed to be threatened by Einstein's hole argument. In its modern guise, which it owes to Stachel and to Earman and Norton, it points out that if $\mathcal{M}_1 = \langle M, g, \phi_1, \ldots, \phi_{n-2} \rangle$ is a model of a generally relativistic theory, then the theory's diffeomorphism invariance entails that $\mathcal{M}_2 = \langle M, d^*g, d^*\phi_1, \ldots,$

$d^*\phi_{n-2}\rangle$ is also a model, for any diffeomorphism d (d^*g etc. are the pull-backs and push-forwards of the original fields under the action of the d). According to the argument, the substantivalist is committed to the view that $\mathcal{M}_1$ and $\mathcal{M}_2$ represent distinct possible worlds. It is then pointed out that this commits the substantivalist to a radical form of indeterminism. In the light of the previous section, it will be clear that what interests us is the claim that substantivalism—viewing spacetime points as genuine individuals—entails that $\mathcal{M}_1$ and $\mathcal{M}_2$ represent two distinct physically possible worlds. Before addressing this issue, however, I should briefly consider another threat to the substantivalist understanding of GR, namely that g should be understood, not as representing spacetime structure, but as a 'gravitational' field, much like any other material field.

Carlo Rovelli is someone who advocates such a view:

> In the physical, as well as philosophical literature, it is customary to denote the differential manifold *as well as* the metric/gravitational field ... as spacetime, and to denote all the other fields (and particles, if any) as matter. But... [i]n general relativity, the metric/gravitational field has acquired most, if not all, the attributes that have characterized matter (as opposed to spacetime) from Descartes to Feynman: it satisfies differential equations, it carries energy and momentum, and, in Leibnizian terms, it *can act and also be acted upon*, and so on.
>
> [...]
>
> Einstein's identification between gravitational field and geometry can be read in two alternative ways:
>
> i. as the discovery that the gravitational field is nothing but a local distortion of spacetime geometry; or
>
> ii. as the discovery that *spacetime geometry is nothing but a manifestation of a particular physical field*, the gravitational field.
>
> The choice between these two points of view is a matter of taste, at least as long as we remain within the realm of nonquantistic and nonthermal general relativity. I believe, however, that the first view, which is perhaps more traditional, tends to obscure, rather than enlighten, the profound shift in the view of spacetime produced by general relativity. (1997, pp. 193-4)

When seeking to decide between these two views it should also be borne in mind that the "metric/gravitational field" has also retained *all* of the attributes that lead us to view the analogous structures in pre-GR theories as codifications of spacetime structure. This point cannot be emphasized enough: there is a sense in which the variable, dynamic metric field g of generally relativistic theories plays *precisely* the same role as the flat, non-dynamic metric η of special relativistic theories. The sole difference between the two types of theory is that in one case spacetime is dynamical, and is governed by Einstein's field equations; in the other it is not. So the *sole* attribute that g has lost is the flip-side of one of the attributes that Rovelli claims it has gained, namely it is no longer *immutable* but is affected by matter. And the substantivalist will, of course, see this as making his realism about spacetime all the more plausible: as Rovelli says, spacetime now obeys the action–reaction principle (Anandan and Brown 1995).

What of Rovelli's contention that the metric of GR satisfies differential equations? The metric and affine structures of pre-GR theories also satisfy differential equations, albeit equations (such as the vanishing of the Riemann tensor) that are not of a great deal of physical interest.

What of the claim that the metric has acquired "most, if not all" of the attributes that might lead us to regard it as matter? Rovelli elaborates the point as follows:

> Let me put it pictorially. A strong burst of gravitational waves could come from the sky and knock down the rock of Gibraltar, precisely as a strong burst of electromagnetic radiation could. Why is the [second] "matter" and the [first] "space"? Why should we regard the second burst as ontologically different from the [first]? (1997, p. 193)

The attributes in question all arise from the fact that the metric is dynamical. Now this certainly supports the view that the metric represents a genuine entity, which does not enjoy an inferior ontological status to matter. But why go further and seek to assimilate it to matter? The phenomenon of gravitational waves certainly pushes one to regard whatever is represented by the metric field as a concrete, substantival entity. But why can't we interpret the potentially devastating effect of gravitational radiation as due to ripples in the fabric of spacetime itself?

Clearly we *can* give very different accounts of the rock's destruction by electromagnetic radiation and by gravitational radiation. According to the substantivalist, the parts of the rock of Gibraltar, as part of an extended rigid body, are being continually and absolutely accelerated away from their natural free-fall motions towards their common centre. The accelerating forces are the electromagnetic forces that account for the rock's rigidity. When the rock is hit by a strong burst of electromagnetic radiation, the natural motions of the parts of the rock do not (significantly) change. Rather the parts of the rock are *differently* accelerated by forces that overcome the counteracting forces between the parts of the rock. When the rock is hit by gravitational radiation, however, no additional accelerative forces are applied. Rather the natural motions are no longer towards the rock's centre but are radically divergent. So divergent, in fact, that the electromagnetic binding forces of the rock are no longer sufficient to accelerate the parts of the rock away from their natural trajectories.

The extent to which the metric should be assimilated to other fields is connected to the controversial question of whether gravitational waves, and more generally the metric, carry energy and momentum. The status of gravitational stress-energy is an intricate topic, but the case for drawing a distinction between it and the stress-energy of matter seems compelling (see Hoefer 2000a, for an extended discussion). Scenarios where gravitational stress-energy seems most well-defined typically involve island matter distributions in asymptotically flat spacetime. Exactly those cases, in other words, where the metric field can analyzed into the flat metric of SR plus a perturbation (see Norton 2000, §3 and the references therein). It is the perturbation, if anything, that corresponds to the "gravitational field" and carries energy-momentum. But Rovelli appears not to wish to view just the perturbation as representing a material field but instead wishes to brand the *entire* metric as 'material'. Gravitational stress-energy and gravitational waves do not force such an interpretation.

Let us grant, then, that a compelling case for regarding the metric field as just another material field has not been made, and that a straightforward, spacetime realist reading of generally relativistic theories remains viable. Our question now is whether such a reading also supports the view that two models related by a non-trivial diffeomorphism represent distinct possibilities. It is clear that this question is a version of the previous section's question (P). If our two diffeomorphic models, $\mathcal{M}_1$ and $\mathcal{M}_2$, are taken to represent two distinct physically possible worlds, then they are worlds which differ solely over a permutation of the spacetime points.

There is close to a consensus in the philosophical literature on Earman and Norton's hole argument that there is nothing anti-substantival about denying that there can be such distinct possible worlds (Butterfield 1989a; Brighouse 1994; Rynasiewicz 1994; Hoefer 1996; an exception is Maudlin 1990). Following Belot and Earman, call any substantivalist position that denies haecceitistic differences, and regards $\mathcal{M}_1$ and $\mathcal{M}_2$ as two representations of the same possible world, *sophisticated substantivalism*.

Belot and Earman refuse to be convinced that the substantivalist can have it so easy. They see sophisticated substantivalists' responses either as lacking a "coherent and plausible motivation", or as indicative of the "insularity of contemporary philosophy of space and time" (2000, p. 167).

I would argue that it is Belot and Earman's refusal to take seriously the responses that they criticize which lacks a coherent and plausible motivation. In particular, their concern is that philosophers of physics appear to have cut themselves off from what physicists working in the area view as genuine conceptual problems. But the philosophers' defence of substantivalism, and their rejection of the dilemma posed by the hole argument, is not incompatible with taking seriously the concerns of working physicists. On closer scrutiny, and despite the lip service some physicists pay to it, the hole argument, and the debate between substantivalism and relationalism, turns out to have rather little to do with the issues of concern to physicists. To insist on reading the issue of substantivalism into these interpretative questions can only lead to confusion.

Two prime examples of physicists' concerns are (i) the notion of a background independent theory and (ii) whether a theory's observables should be diffeomorphism invariant.

Concerning (i), note that background independence has to do with whether a theory posits non-dynamical, absolute (background) fields, not with whether it sanctions haecceitistic differences. This feature of a theory has nothing to do with the hole argument, as Earman and Norton's application of that argument to 'local spacetime' formulations of background-dependent pre-GR theories illustrates (1987, pp. 517–8).[22] However, it may well be connected to whether the appropriate Hamiltonian formulation of the theory is a constrained Hamiltonian theory for which the diffeomorphism group is a gauge group in a technical sense.[23]

Concerning (ii), the question of whether a theory's observables are diffeomorphism-invariant needs further explication.[24] If it is taken to entail that no physical magnitude can take different values at different times (a version of the so-called 'problem of time'), then it is a *stronger* claim than the anti-haecceitistic claim that all diffeomorphic models represent the same physical situation.

Sophisticated substantivalism may be compatible with taking seriously physicists concerns, but does it have a coherent motivation? The obvious thing to be said for the position is that one thereby avoids the indeterminism of the hole argument. This motivation is, of course, rather *ad hoc*. A less *ad hoc* motivation would involve a metaphysics of individual substances that does not sanction haecceitistic differences, perhaps because the individuals are individuated by—their numerical distinctness is grounded by—their positions in a structure. In the next section we will see that Stachel has recently sought to embed his response to the hole argument in exactly this type of more general framework. I hope enough has been said in this section and the previous one to indicate the coherence of such a point of view; it is perhaps a modest structuralism about spacetime points, but it is a far cry from the objectless ontology of the ontic structural realist.

There is one final line of defence of sophisticated substantivalism that needs to be undertaken. One might concede that in principle anti-haecceitism is com-

[22]Of course, it becomes a moot question whether the fields representing spacetime structure should count as 'non-dynamical' background fields in the context of local spacetime formulations of pre-GR theories. In one sense they are dynamical, since they are held to obey field equations such as $R^a_{bcd} = 0$. In another sense, they are non-dynamical, since they do not vary (except globally) from model to model.

[23]See Earman (2003, pp. 151–3); I hope to return to this topic on a future occasion.

[24]Consider, for example, Smolin's distinction between "causal observables" and "Hamiltonian constraint observables" in (2000b).

patible with spacetime points being substances, but nevertheless believe that the theoretical treatment of spacetime, read literally, strongly supports haecceitism. Is it not the case that the *natural* reading of $\mathcal{M}_1 = \langle M, g, \phi_1, \ldots \phi_{n-2} \rangle$ and $\mathcal{M}_2 = \langle M, d^*g, d^*\phi_1, \ldots d^*\phi_{n-2} \rangle$ interprets each point of M as representing the very same point in each model, and therefore interprets the two models as attributing different properties to each point? Moreover, doesn't the mathematics of GR presuppose that the numerical distinctness of the points of M is independent of the properties and relations assigned to them by the fields g, ϕ_i? So if we're being literalistic realists about our theories, shouldn't we take a similar stance towards the individuality of spacetime points?

Something like this line of thought might well be responsible for what resistance there remains to sophisticated substantivalism's combination of anti-haecceitism and realism about spacetime points. But it does not stand up to scrutiny. For a start, it is not obvious that the numerical distinctness of the points of the mathematical object M *is* independent of their properties and relations. We have already had reason to consider structuralist approaches to mathematical objects. According to the mathematical structuralist, the individuality of the points of M *does* depend on their positions in the mathematical structure of which they are part. The mathematical structuralist, of course, needs to be able to give an account of the difference between the two models $\mathcal{M}_1$ and $\mathcal{M}_2$. Here the most obvious strategy is to point out that such a difference can be made out if the two are considered as distinct substructures embedded in a larger structure (*cf.* Parsons 2004, pp. 68–9). (And if this line is taken, there is no reason, of course, to think that the substantivalist should postulate a counterpart in concrete reality of this larger (unspecified!) structure.)

If one remains attracted to these particular lines of haecceitist argument, a useful question to ask oneself is the following. Suppose, for the sake of argument, that the sophisticated substantivalist is right: individual spacetime points exist as basic objects, but possible spacetimes correspond to equivalence classes of diffeomorphic models of GR. How should the formalism of GR be modified to take account of the anti-haecceitism? (Note that this is not the demand for a *relationalist* reformulation that *does away with* spacetime points.) It should be clear that no such thing is needed. As soon as one is committed to the existence of a set of points with various geometrical properties, even if one is avowedly anti-haecceitistic, the most obvious way of representing such a set will be open to a haecceitistic misinterpretation.

In fact, the haecceitist substantivalist's mistake is a specific instance of a common type. Back in 1967, Kaplan identified the occurrence of essentially the same error in a rather different context:

> the use of models as representatives of possible worlds has become so natural for logicians that they sometimes take seriously what are really only artifacts of the model. In particular, they are led almost un-

> consciously to adopt a *bare particular metaphysics*. Why? Because the model so nicely separates the bare particular from its clothing. The elements of the universe of discourse of a model have an existence which is quite independent of whatever properties the model happens to tack onto them. (p. 97)

It seems that the use of mathematical models as representatives of possible worlds has become so natural for some philosophers of physics that they too have been led almost unconsciously to endorse haecceitistic distinctions that are really only artifacts of the model.

6 Stachel's Generalized Hole Argument for Sets.

According to John Stachel, the moral of the hole argument is that diffeomorphically-related mathematical solutions to the field equations of GR (hereafter: *diffeomorphs*) do not represent physically distinct solutions. Although, in the past, he has referred to his position as a relationalist one, it is really, as I use the terms, a version of sophisticated substantivalism. He does not believe that to count diffeomorphs as representing the same physical solution one has to eliminate spacetime points.[25] According to Stachel, one would be forced to view diffeomorphs as representing distinct physical solutions if one took spacetime points to be "individuated independently of the metric field". One can maintain the existence of points and count diffeomorphs as representing the same physical solution if one assumes that "the points of the manifold are *not* individuated independently of the $g_{\mu\nu}$ field; i.e., that these points inherit *all* their chronogeometrical (and inertiogravitational) properties and relations from that field" (Stachel 2002, p. 233).[26]

[25]He does refer to fibre bundle formulations of theories in which one 'eliminates' an independently specified base space, replacing it with the quotient space of the total space by the fibres. With the theory so formulated, one cannot permute the fibres of the total space without thereby permuting the points of the base space. But this does not mean, as Stachel claims, that one cannot even generate the models that provide the basis of the original hole dilemma. Two models the cross sections of which are related by a fibre-preserving diffeomorphism of the total space still represent *mathematically* distinct objects. Further, since such cross sections are 'differently placed' relative to the fibres, and since the base space is simply defined as the quotient space of the total space by the fibres, these two models represent two distributions of structurally identical fields 'differently placed' on the base space (i.e. spacetime). Of course, Stachel deems all cross sections related by fibre-preserving diffeomorphisms of the total space to be physically equivalent. I agree. But we can, with as much right, make the analogous claim of the models of the traditional formulation of the theory. Consideration of formulations in terms of fibred manifolds without an independently specified base space gains us nothing.

[26]In my view there is something at least misleading about talk of the points 'inheriting' their properties from this 'field'. It suggests that the field is some entity in its own right (the 'gravitational

So far I am in agreement with Stachel. The talk of the points being 'individuated' is a little obscure. At one point Stachel talks equivalently of an entity being 'distinguishable' from entities of the same kind (p. 236). "Distinguishable" can be understood in an epistemological or an ontological sense. It must be the latter that is in question, and I suggest that, minimally, it is the entities' determinate numerical distinctness that is at stake (recall questions (ND1) and (ND2) from section 4). In claiming that points are not individuated independently of the metric field, Stachel can be understood as claiming that their determinate distinctness from one another is grounded in their standing in the spatio-temporal relations to one another that they do. This in turn is held to prevent our interpreting diffeomorphically-related models as representing two situations involving the very same points as occupying different positions in the very same network of spatio-temporal relations.

Now Stachel wishes to situate the hole argument, and its moral concerning the individuation of spacetime points, in a more general framework. Rather than considering only sets of spacetime points, he considers an arbitrary set **S** of n entities. And rather than considering the spatio-temporal properties encoded by $g_{\mu\nu}$, he considers an arbitrary ensemble **R** of n-place relations.[27][28] Stachel calls the relational structure $\langle \mathbf{S}, \mathbf{R} \rangle$ a "world". What can be said about the 'individuation' of the elements of **S**? Stachel suggests that there are two possibilities (within which he draws further divisions, which I will ignore for now):

1. the entities are individuated (that is, [are] distinguishable from other entities of the same kind) prior to and without reference to the relations **R**...

2. the entities are not individuated (that is are indistinguishable among themselves) without reference to the relations **R**. (2002, p. 236)

Later Stachel goes on to adapt terminology he borrows from Marx, and dubs entities of kind (2) *reflexively defined entities*.

Now diffeomorphisms are simply a special class of permutations of the set of manifold points, those that preserve the manifold's topological and differential

field'). My preferred view, as explained in the previous section, is that the mathematical metric field is simply a *specification* of the points' (spatio-temporal) properties. It does not represent some*thing* that bestows or engenders the points' properties. It stands to the points as, e.g., red stands to red things (however that is!).

[27]In what follows I adopt as far as possible Stachel's convention of using boldface type when referring to sets, italic type when referring to n-tuples (and to numbers), and roman type when referring to elements of sets or of n-tuples.

[28]One might consider relations of any adicity. For any adicity $N < n$, Stachel claims that the restriction to n-place relations is no real restriction because there is a natural way of associating an n-place relation with any N-place relation. Assuming I have understood his description correctly, the procedure that Stachel outlines will associate a single n-place relation with some pairs of distinct relations (of different adicity). I do not claim that this conflation causes problems for Stachel.

structure. And just as one can start with a given model $\langle M, g, \phi_i \rangle$ and consider the mathematically distinct model $\langle M, d^*g, d^*\phi_i \rangle$ generated by the action on the fields induced by a particular diffeomorphism d, one can start with $\langle \mathbf{S}, \mathbf{R} \rangle$ and consider a new structure $\langle \mathbf{S}, \mathbf{PR} \rangle$ generated by the action on the M relations $\mathrm{R}_1, \ldots \mathrm{R}_M$ in $\mathbf{R}$ induced by an *arbitrary permutation* $\mathrm{P} : \mathbf{S} \to \mathbf{S}$. The definition of $\mathrm{PR}_i \in \mathbf{PR}$ is obvious: (taking an n-place relation to be defined extensively as a set of n-tuples of elements of $\mathbf{S}$, i.e. as a subset of $\mathbf{S}^n$) for any $s = \langle \mathrm{s}_1, \ldots, \mathrm{s}_n \rangle \in \mathbf{S}^n, s \in \mathrm{PR}_i$ just if $\mathrm{P}^{-1}s = \langle \mathrm{P}^{-1}(\mathrm{s}_1), \ldots, \mathrm{P}^{-1}(\mathrm{s}_n) \rangle \in \mathrm{R}_i$.

By analogy with the interpretative questions that arise in connection with $\langle M, g, \phi_i \rangle$ and $\langle M, d^*g, d^*\phi_i \rangle$, one might consider whether the structure $\langle \mathbf{S}, \mathbf{R} \rangle$'s being a possible world entailed that $\langle \mathbf{S}, \mathbf{PR} \rangle$ is also a possible world. And, if both are held to be possible worlds, one might consider whether $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{PR} \rangle$ should be interpreted as the same, or as distinct possible worlds. Of course, $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{PR} \rangle$ will, in general,[29] not be identical, just as $\langle M, g, \phi_i \rangle$ and $\langle M, d^*g, d^*\phi_i \rangle$ are, in general, distinct mathematical entities. If we wish nonetheless to talk of their 'being' the same possible world, then a distinction needs to be drawn between the structures $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{PR} \rangle$, considered as mathematical objects, and the possible worlds they represent. Our second question then becomes whether $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{PR} \rangle$ *represent* the same, or different, possible worlds.[30]

Stachel talks of *permutable* and *generally permutable* worlds, theories, and even entities. Although this is not how Stachel defines the terms, I propose the following definitions. The structure $\langle \mathbf{S}, \mathbf{R} \rangle$ is a *permutable world* just if, if it represents a possible world then $\langle \mathbf{S}, \mathbf{PR} \rangle$ also represents a possible world, for every permutation P. The structure $\langle \mathbf{S}, \mathbf{R} \rangle$ is a *generally permutable world* just if, if it represents a possible world then $\langle \mathbf{S}, \mathbf{PR} \rangle$ represents the *same* possible world, for every permutation P. With one minor qualification, I can follow Stachel in his definition of a *theory* and of a *permutable theory*:

> A *theory* is a "rule that picks out a class of worlds: in other words, a class of ensembles of n-place relations: $\mathbf{R}, \mathbf{R}', \mathbf{R}''$, etc., whose places are filled by the members of the same set $\mathbf{S}$ of n entities $\mathbf{a}$; further, let it be a permutable theory … [just in case], if $\mathbf{R}$ is in the selected class of worlds, so is P$\mathbf{R}$ for all P. (2002, p. 244)[31]

[29] The qualification concerns the case when *all* of the relations in $\mathbf{R}$ are symmetric with respect to all permutations of $\mathbf{S}$ (i.e., if for any $\mathrm{R}_i, s \in \mathrm{R}_i$ just if $\mathrm{P}s \in \mathrm{R}_i, \forall \mathrm{P}$), a case that will be important in the context of quantum particles.

[30] A way of talking that Stachel quickly slips into.

[31] My reservation concerns talk of a relation's 'places' being 'filled' by the members of some set. If we are conceiving of relations on a given domain in purely extensional terms (*cf.* Stachel 2002, p. 237), then strictly it makes no sense to talk of a fixed relation as having places that might be variously filled.

Despite their status as the analogues of diffeomorphic models of GR, Stachel nowhere explicitly considers two distinct *structures* such as $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{PR} \rangle$. Instead he seeks to define things in terms of expressions of the form "$\mathrm{R}_i(a)$ holds" and "$\mathbf{R}(a)$ holds".[32]

Now the claim that $\mathrm{R}_i(a)$ holds is simply the claim that $a \in \mathrm{R}_i$ (recall that a is a particular n-tuple and that we are considering only n-place relations). But what does "$\mathbf{R}(a)$ holds" mean? The obvious interpretation is that $a \in \mathrm{R}_i$ *for all* $\mathrm{R}_i \in \mathbf{R}$.[33] But note that, for generic sets $\mathbf{S}$ and ensembles of relations $\mathbf{R}$, there will be *no* sequence $a \in \mathbf{S}^n$ such that $\mathbf{R}(a)$ holds![34]

This might lead one to suspect that Stachel's attempts to define the ensemble of relations **PR**, and his notion of a permutable world, in terms of the expression $\mathbf{R}(a)$ are doomed, and indeed they are. **PR** is said to hold for a if and only if $\mathbf{R}(\mathrm{P}^{-1}a)$ holds. Stachel is not explicit about whether this is required for every a, but either way the definition is not equivalent to the (more standard) definition of **PR** given above. For consider, as an example, $\mathbf{S} = \{\mathrm{s}_1, \mathrm{s}_2\}$, $\mathbf{R} = \{\mathrm{R}_1 = \{\langle \mathrm{s}_1, \mathrm{s}_1 \rangle\}, \mathrm{R}_2 = \{\langle \mathrm{s}_2, \mathrm{s}_2 \rangle\}\}$ and $\mathbf{R}' = \{\mathrm{R}'_1 = \{\langle \mathrm{s}_1, \mathrm{s}_1 \rangle\}, \mathrm{R}'_2 = \{\langle \mathrm{s}_1, \mathrm{s}_2 \rangle\}\}$. Now $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{R}' \rangle$ are non-isomorphic structures, but because neither $\mathbf{R}(a)$ nor $\mathbf{R}'(a)$ holds for any ordered pair a of elements of $\mathbf{S}$, $\mathbf{R}(a)$ holds iff $\mathbf{R}(\mathrm{P}^{-1}a)$ holds.[35]

Stachel's definition of a permutable world is equally problematic. $\langle \mathbf{S}, \mathbf{R} \rangle$ is said to be permutable if, whenever $\mathbf{R}(a)$ is a possible state of the world, then $\mathbf{PR}(a)$ is also a possible state of the world, for every $n!$ permutations P of **a**.[36] In the light of

[32]Note that, as is standard practice, Stachel here makes the letter R_i do double duty as a predicate letter (in the expression "$\mathrm{R}_i(a)$") and as the name for the relation that is this predicate's extension.

[33]For the record, Stachel's own elucidation is that '$\mathbf{R}(a)$' stands for "the entire ensemble of relations filled by that sequence [i.e., by the particular n-tuple a]" (2002, p. 237)

[34]It will fail to hold for all a whenever $\mathbf{R}$ contains two disjoint relations. Rather unfortunately for Stachel, the pair of binary relations (promoted to continuously infinite relations as per Stachel's recipe) on a set of spacetime points expressed by the predicates "x is timelike related to y" and "x is spacelike related to y" is just such an example of disjoint relations. It will also fail even if no relations are disjoint, just so long as, e.g., the intersection of two relations does not intersect with a third etc. There is another oddity worth noting. Even for structures $\langle \mathbf{S}, \mathbf{R} \rangle$ such that $\mathbf{R}(a)$ (interpreted in this way) holds, $\mathbf{R}(a)$ constitutes an incomplete, and very arbitrary, specification of the structure. We are told that for this particular a, $a \in \mathrm{R}_i$ for all $\mathrm{R}_i \in \mathbf{R}$, but this tells us next to nothing about $\mathbf{R}$. To be given the *complete* state of the world, we need to be told, for *every* n-tuple, and for every R_i individually, whether or not the n-tuple is in the relation.

[35]This will be true for every pair of ensembles of relations on the same domain such that, for each ensemble, no n-tuple of elements of the domain is a member of every relation; the ensembles do not even have to be equinumerous! The condition "$\mathbf{PR}(a)$ holds iff $\mathbf{R}(\mathrm{P}^{-1}a)$ holds" therefore clearly fails to define **PR** in terms of **R**.

[36]If there is a possible world in which $\mathbf{R}(a)$ holds, Stachel calls $\mathbf{R}(a)$ a possible state of the world, and if $\mathbf{R}(a)$ holds he calls it a state of the world (2002, pp. 237–8). As Stachel notes, there will only be as many as $n!$ distinct permutations if a is a *nonduplicating* n-tuple. With what right does Stachel consider only *nonduplicating* n-tuples? This is an indication that he is not really interpreting "**R**" in "$\mathbf{R}(a)$" as expressing an ensemble of relations, extensively defined or otherwise.

the troubles noted in the last paragraph, let us simply take $\langle \mathbf{S}, \mathbf{PR} \rangle$ to be the isomorphic structure generated by P, in the way outlined above on page 27, and consider again the example $\mathbf{S} = \{s_1, s_2\}$, $\mathbf{R} = \{R_1 = \{\langle s_1, s_1 \rangle\}, R_2 = \{\langle s_2, s_2 \rangle\}\}$. Intuitively, if $\langle \mathbf{S}, \mathbf{R} \rangle$ is to be permutable, $\langle \mathbf{S}, \mathbf{PR} \rangle$ should be (should represent) a possible world. But since $\mathbf{R}(a)$ and $\mathbf{PR}(a)$ are *not* states of the worlds $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{PR} \rangle$ for any a, whether $\langle \mathbf{S}, \mathbf{R} \rangle$ counts as permutable will be independent of whether or not $\langle \mathbf{S}, \mathbf{PR} \rangle$ also counts a possible world. One might hope that by taking the relations to be defined intensively, rather than extensively, sense can be made of these definitions. A little reflection shows that this will not work either.[37]

The ensemble **PR**, then, simply cannot be defined in terms of $\mathbf{R}(a)$ in the way Stachel suggests, however one understands $\mathbf{R}(a)$. In fact, the obvious definition of **PR** in terms of expressions such as $R_i(a)$ is the following:

$$\forall \mathrm{P}, \forall R_i \in \mathbf{R} \text{ and } \forall a \in \mathbf{S}^n, \mathrm{PR}_i(a) \text{ iff } R_i(\mathrm{P}^{-1}a)$$

That is, it must be given in terms of the individual expressions $R_i(a)$, not via the single expression $\mathbf{R}(a)$. Why does Stachel introduce the expression $\mathbf{R}(a)$ at all? It figures prominently in his set-theoretic version of the hole argument. We are asked to consider the class of ensembles of relations $\mathbf{R}, \mathbf{R}', \mathbf{R}''$ picked out by some permutable theory, and then the question is put "Could such a permutable theory pick out a *unique* state of the world by first specifying a unique world, i.e., one **R**; and then specifying how any number m of its places *less than* $(n - 1)$ are filled?" (2004, p. 244). Not if the entities in question are not reflexively defined, claims Stachel, for then, "$\mathbf{R}(a)$ and $\mathrm{P}\mathbf{R}(a)$ represent different states of the world" (*ibid.*). But this is just false. Whenever all the relations in **R** are symmetric with respect to all permutations, we get back the identical structure ($\mathbf{R}(a)$ iff $\mathrm{P}\mathbf{R}(a)$), so there can be no question of their representing different worlds. And even setting aside the special case of symmetric ensembles of relations, we again have the unwarranted restriction to nonduplicating n-tuples, for otherwise there is the possibility that $\mathbf{R}(a)$ and $\mathrm{P}\mathbf{R}(a)$ are both states of the same structure because, although $\mathbf{R} \neq \mathrm{P}\mathbf{R}$, $a = \mathrm{P}a$.

[37] Stachel provides the following example to illustrate his definitions. Our set is $\{\mathrm{Cat}, \mathrm{Cherry}\}$ and $R(x_1, x_2)$ holds iff x_1 is on a mat and x_2 is on a tree. With the relation specified in this way, one can perhaps make sense of the same relation 'having its places filled' by different n-tuples, even when the relation does not in fact hold of the n-tuples in question (*cf.* 2002, p. 237). But this does not enable us to make sense of $\mathbf{PR}(a)$. For consider extending the example with another relation: $R'(x_1, x_2)$ holds iff x_1 is red and x_2 has claws. And consider the world in which the cat is on a mat (and the cherry is not), the cherry is on a tree (and the cat is not), the cherry is red (and the cat is not) and the cat has claws (and the cherry does not). The obvious interpretation of "$\mathbf{R}(\mathrm{Cat}, \mathrm{Cherry})$ holds" is that the cat is on a mat and the cherry is on a tree *and* that the cat is red and that the cherry has claws, i.e., the claim that $\mathbf{R}(\mathrm{Cat}, \mathrm{Cherry})$ holds is *false* at the specified world, and, let us suppose, at every possible world. And, more generally, it seems that we again we have an ensemble **R** of relations for which $\mathbf{R}(a)$ never holds for any ordered pair a of our domain.

I think enough has been said to show that, taken literally, Stachel's description of a set-theoretic hole argument is in terminal trouble. But it is equally clear how to make precise what he must have had in mind. We need the notion of a *complete description* of the structure $\langle \mathbf{S}, \mathbf{R} \rangle$, and also of a *partial* description, for the latter will be the analogue of a *specification* of a metric field on all of a differentiable manifold save for a compact region (the 'hole' of the hole argument).[38] One way of giving such a complete description is the following. We suppose that for every relation $\mathrm{R}_i \in \mathbf{R}$ we have a predicate symbol R_i (I follow Stachel in using the same letter for the predicate and the relation it expresses). And similarly, we suppose that we have a name $\mathrm{a_k}$ for every element of **S**. A complete description will be a conjunction of the following three formulas: (i) an $(n^2 \times M)$-place conjunction that includes, for every predicate letter R_i and for every n-place sequence a of the names for elements of **S**, either the formula $\mathrm{R}_i(a)$ or the formula $\neg\mathrm{R}_i(a)$; (ii) an $n!$-place conjunction $(\bigwedge_{i \neq j} \mathrm{a}_i \neq \mathrm{a}_j)$ stating that different names name different elements of **S**; and (iii) an n-place disjunction $(\forall y (\bigvee_{i=1}^{n} y = \mathrm{a}_i))$ stating that there are no elements of **S** other than $a_1, \ldots, a_k$.

Now we may introduce $\mathcal{R}(a)$ as an abbreviation for this conjunction, where a is some particular *non-repeating* n-place sequence of the n names a_k. $\mathcal{R}(\mathrm{P}^{-1}a)$ will then be (an abbreviation of) a complete description of the (mathematically) distinct but isomorphic structure $\langle \mathbf{S}, \mathbf{PR} \rangle$, and $\exists x_1 \ldots \exists x_n \mathcal{R}(x_1, \ldots, x_n)$ will be a structural description true of any structure isomorphic to $\langle \mathbf{S}, \mathbf{R} \rangle$.[39]

I suggest that it is expressions such as $\mathcal{R}(a)$ that Stachel needs for his hole argument for sets. Recall that he considers first specifying a unique world **R**, and then specifying how many any number m of its places *less than* $(n-1)$ are filled. Given the extensive conception of relation that he began with, to specify **R** is *already* to specify, for every n-tuple, whether or not it is a member of each relation in **R**.[40] Instead, what he should have considered was first specifying an isomorphism class of structures via $\exists x_1 \ldots \exists x_n \mathcal{R}(x_1, \ldots, x_n)$, and then specifying how m 'places' in this structure are filled via a formula such as $\exists x_1 \ldots \exists x_{(n-m)} \mathcal{R}(b, x_1, \ldots, x_{(n-m)})$, where b is some particular non-repeating m-place sequence of the names for the elements of **S**.

[38] In what follows I am influenced by the notational conventions of §3 of Belot (2001), which provides an admirably clear and uncluttered account of the symmetries and permutations of abstract structures and our means of describing them in a first-order language. Stachel is at pains to stress the coordinate independence of the hole argument, which might partly account for his reluctance to engage in the necessary, though limited, semantic ascent.

[39] As Belot notes, no first-order theory (i.e., no set of sentences involving only variables and no names) can determine a structure up to isomorphism if n is infinite. Correspondingly, the prescription we are considering, which *does* determine the structure up to isomorphism, yields an infinitary formula, i.e. something that is not well-formed according to standard first-order logic.

[40] And to give only an intensive definition of an ensemble of relations will fail to pick out a particular isomorphism class of structures at all.

Let $a = (b, c)$ (a is an n-place sequence of names, b is an m-place sequence of names and c is an $(n - m)$-place sequence of names). If the theory we are considering is a permutable theory, then if the structure that corresponds to the description $\mathcal{R}(a) = \mathcal{R}(b, c)$ is a possible world, so too is the structure corresponding to the description $\mathcal{R}(\mathrm{P}a) = \mathcal{R}(b, \mathrm{P}^{(n-m)}c)$, i.e., we consider a permutation P that acts non-trivially only on the last $(n - m)$ members of the sequence a. If the elements of **S** are not reflexively defined (i.e., if these entities are individuated independently of the ensemble of relations **R**, i.e., if $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{P}^{-1}\mathbf{R} \rangle$ represent distinct possible worlds), then specifying only $\exists x_1 \ldots \exists x_{(n-m)} \mathcal{R}(b, x_1, \ldots, x_{(n-m)})$ fails to pick out a unique world, for it is compatible with both $\mathcal{R}(a)$ and $\mathcal{R}(\mathrm{P}a)$, which both describe (distinct) structures allowed by the theory (namely, $\langle \mathbf{S}, \mathbf{R} \rangle$ and $\langle \mathbf{S}, \mathbf{P}^{-1}\mathbf{R} \rangle$ respectively), structures that represent distinct possible worlds. On the other hand, if the elements of **S** *are* reflexively defined (if they are 'generally permutable entities' that are *not* individuated independently of the ensemble of relations), then specifying only $\exists x_1 \ldots \exists x_{(n-m)} \mathcal{R}(b, x_1, \ldots, x_{(n-m)})$ *is* sufficient to pick out a unique world because, although compatible with both $\mathcal{R}(a)$ and $\mathcal{R}(\mathrm{P}a)$, these describe only formally distinct, isomorphic structures which represent the same world. It is surely in this latter case that Stachel intends to apply the label "generally permutable" to the theory, and to the worlds it picks out. This, I submit, is the proper set-theoretic analogue of the hole argument.[41]

7 Stachel on Identical Particles.

It is finally time to consider Stachel's application of all of this to quantum particles. I quote at length:

> One can immediate apply this result to the current discussion about the individuality of elementary particles…One group maintains that each elementary particle retains its individuality, and that quantum statistics are merely the result of the fact that certain states…that are accessible to systems of elementary particles that are not of the same kind, are for some reason inaccessible to systems of particles that are all of the same kind. The other group maintains that quantum statistics has its origin in the lack of individuality of elementary particles. As far as I know, no one has …mentioned the possibility of extending the hole argument from the discussion of the individuality of space-time

[41] A specification of the metric on all of a manifold except a 'hole' (once one has solved the equations and determined the metric in the hole up to isomorphism) corresponds, in effect, to the formula $\exists x_1 \ldots \exists x_{(n-m)} \mathcal{R}(b, x_1, \ldots, x_{(n-m)})$. I.e., one has specified an equivalence class of diffeomorphic solutions to Einstein's field equations and has further specified, for all but the hole, which manifold points have which spatio-temporal properties.

> points to the discussion of the individuality of elementary particles, as I shall now do.
>
> If we take the points of our set to represent n elementary particles of the same kind, then quantum-mechanical statistics imposes the requirement that all physical relations between them be permutable. Our set theoretical hole argument shows that, if we ascribe an individuality to the particles that is independent of the ensemble of permutable relations, then no model can be uniquely specified by giving all the n-place relations $\mathbf{R}$ between them unless we further specify *which* particle occupies *each* place in these relations... (2002, p. 245)

In a footnote he expands on what he means by "the requirement that all physical relations between [the particles] be permutable":

> The relations will represent values of physical properties of the system of identical particles, which must remain invariant under all permutations of the particle labels. Since these physical properties are represented by bilinear functions of the state vector of the system, they will remain invariant *whether the state vector remains invariant under a permutation (bosons) or changes sign (fermions)*... (ibid., p. 261; my emphasis)

The obvious interpretation of the claim that the *relations* between the particles are themselves permutable is that all permutations of the set $\mathbf{S}$ are *symmetries* of these relations. That is, for every R_i, $a \in \mathbf{S}^n$ and permutation P, $a \in R_i$ iff $Pa \in R_i$. In this special case we have $\langle \mathbf{S}, \mathbf{R} \rangle = \langle \mathbf{S}, \mathbf{PR} \rangle$; there just can be no question of the two structures representing distinct worlds because we do not have two structures: there is only one. Such a situation does indeed correspond to what we find in the case of the quantum mechanics of identical bosons and fermions. The quantum states of such systems are required to be either symmetrized or antisymmetrized. That is, for an arbitrary permutation of the particle labels, one gets back the very same state (up to a phase factor of -1 in the case of fermions, if the permutation is odd). Stachel's footnote suggests that he is indeed considering such symmetrized and antisymmetrized states. His characterization of the group who maintain that elementary particles retain their individuality but that "certain states... that are accessible to systems of elementary particles that are not of the same kind, are for some reason inaccessible to systems of particles that are all of the same kind" also strongly suggests that he has (anti)symmetrized states in mind. Moreover, when one talks of the "permutation invariance" of the quantum mechanics of identical particles, one is typically referring to the fact that the *physically allowed states them-*

selves are permutation invariant (up to a phase).[42]

But now consider how this situation plays out in the context of the set-theoretic hole argument. First consider Stachel's 'states of the world' $\mathbf{R}(a)$ and $\mathbf{R}(\mathrm{P}a)$. These are indeed distinct states of the world, but for any world $\langle \mathbf{S}, \mathbf{R} \rangle$ allowed by quantum mechanics, if $\mathbf{R}(a)$ is a state of this world, then $\mathbf{R}(\mathrm{P}a)$ is also a state of the very same world. If the relations themselves are permutable then $\mathbf{R}(a)$ and $\mathbf{R}(\mathrm{P}a)$ are either both true, or both false. Consider, instead, the complete description $\mathcal{R}(a)$ introduced above. If $\mathcal{R}(a)$ is a description of a world allowed by quantum mechanics, then $\mathcal{R}(\mathrm{P}a)$ will be a description of *exactly the same world* (it will be a logically equivalent formula). Thus stipulating that $\exists x_1 \ldots \exists x_{(n-m)} \mathcal{R}(b, x_1, \ldots, x_{(n-m)})$ *does* suffice to pick out a unique world, *even if one believes that the individuality of quantum particles transcends the relational structure in which they are embedded.*

In fact, Stachel does not always talk as if the ensemble of relations in question is itself permutable (i.e. that it corresponds to a symmetric or antisymmetric state). Before applying the hole argument to quantum particles, he writes:

> Some suggest that, if elementary particles are not individuated, then any attempt to label them is misguided. On the contrary, it is just an example of the usual method of coordinatization, introduced when treating any set of entities that are numerous, yet indistinguishable...What is important to realize is that, in all such cases, no one coordinatization (labelling in this case) is preferred over another; and that it is precisely invariance of all relations under all permutations of the labels that guarantees this. It is entirely indifferent which six electrons out of the universe make up a particular carbon atom[.] They are individuated, as K-shell or L-shell electrons of the atom, for example, entirely by the ensemble of their relations to the carbon nucleus of the atom and to each other. Indeed, the notation for the electronic structure of an atom is based on this type of individuation. (2002, p. 243)

Here, again, Stachel talks of the "invariance of all relations under all permutations of the labels." However, he also claims that one can think of, e.g., a K-shell electron as being individuated by the ensemble of the relations that hold between the electrons of the atom, and between the electrons and their nucleus. In a footnote he elaborates: "as a result of the Pauli exclusion principle for fermions, each electron

[42]Strictly, this corresponds to the quantum mechanics of identical particles satisfying the *symmetrization postulate*. Sometimes 'permutation invariance' is used to refer the strictly weaker requirement that the expectation values of all *physical* observables be permutation invariant (see French and Rickles 2003, §2-3). Note that (i) the more general possibilities allowed by the weaker interpretation do not appear to be realized in nature and that (ii) to require the permutation invariance of *states* (up to a possible phase factor of −1) is to impose the symmetrization postulate.

in an atom can be fully individuated by the set of its quantum numbers" (2002, p. 260). And later he claims that his set-theoretic hole argument shows that:

> if we ascribe an individuality to the particles that is independent of the ensemble of permutable relations, then no model can be uniquely specified by giving all the n-place relations **R** between them unless we further specify just *which* particles occupies *each* place in the relations. For example, the rules for filling atomic shells in the ground state of an atom with electrons would have to be regarded as radically incomplete, since they do not tell us *which* electron has the different quantum numbers that characterize that state. (2002, pp. 245-6)

At this point, we should distinguish two, quite distinct, ways of describing, e.g., the electrons in a particular atom. First we might give its state as an anti-symmetrized vector in a tensor product Hilbert space. It is this description that corresponds to an ensemble of *permutable* relations. And, despite what Stachel appears to suggest, one *cannot* think of the labels that feature in this description as (arbitrary coordinate) labels for electrons with particular quantum numbers, individuated by the ensemble of relations. First, precisely because they are permutable, the relations fail to individuate in the way Stachel appears to imply: every place in the ensemble of relations is exactly like every other place.[43] Secondly, and relatedly, in this formalism, each electron, i.e., the entities associated with *each* label, enters equally into the state associated with a particular set of quantum numbers, e.g., those corresponding to a K-shell electron with a particular spin.[44] It is for this reason that it is held by some that "attempts to label the electrons are misguided." It is just hard to believe that the labels in the tensor product formalism are genuine labels (*cf.* Teller 1998, §5); they are certainly not an example of the "usual method of coordinatization".

The alternative description involves Fock space and the occupation number formalism (see, e.g., van Fraassen 1991, pp. 438-48). Here no particle labels are employed at all. Rather the formalism uses *occupation numbers*, "numbers describing how many times each maximal property is instantiated, with no regard to "which"

[43] This is not say that such relations cannot be held to individuate; although the relations are permutable, they may represent, e.g., symmetric yet irreflexive relations, which is enough to force the numerosity of the domain if one adopts the Hilbert and Bernays' definition of identity (see Saunders 2003, esp., 294-5). We can even think of *bosons* as being individuated by such relations when their symmetrized entangled state does not involve more that one copy of any single-particle state in each element of the superposition.

[44] These points are related to the claim that all identical particles, fermions just as much as bosons, violate the identity of indiscernibles, in every physically possible state (French and Redhead 1988). The claim is only true if the labels of the tensor product Hilbert space formalism are interpreted as genuine labels.

particle has which of the properties" (Teller 1998, p. 128). Such descriptions correspond to formulae of the type $\exists x_1 \ldots \exists x_n \mathcal{R}(x_1, \ldots, x_n)$ involving no names and only variables. Note that the $\mathcal{R}$ occurring in *this* description will not abbreviate the same complex relation as the $\mathcal{R}$ that occurs in a description of the permutable structure corresponding to a symmetrized state vector. The $\mathcal{R}$s of the latter type of description have every possible symmetry: $\mathcal{R}(a)$ and $\mathcal{R}(\mathrm{P}a)$ are logically equivalent formulae for every P. In contrast, the $\mathcal{R}$ of an occupation number description will have *no* symmetries: each place must correspond to a *different*, maximal, property.

If we now try to run the hole argument for an occupation number description we run into trouble, for there simply is nothing in the formalism that corresponds to further specifying which entities occupy which occupied states, e.g., the various atomic shells in an atom. As soon as we do introduce things that formally look like such names—the labels of the tensor product Hilbert space formalism—we symmetrize, so that every label is associated with every occupied atomic shell. Either way, there appears to be no analogue of the hole argument for quantum particles. Stachel appears to conflate the two descriptions: one the one hand he talks about labels, and permutable relations (suggesting the labelled tensor product Hilbert space formalism), on the other hand he talks about different electrons being individuated by different sets of quantum numbers, in conformity with the Pauli exclusion principle (suggesting an occupation number description).

I hope to have made it clear that the diffeomorphism invariance of GR, and the permutation invariance of quantum mechanics are very different. In the first case, the diffeomorphism invariance is a symmetry of the *theory*.[45] If $\langle M, g, \phi_i \rangle$ is a solution to the theory, then so is (the mathematically distinct) $\langle M, d^*g, d^*\phi_i \rangle$. But whether or not we take the interpretative step of regarding these two models as representing the same world, arbitrary diffeomorphisms are not symmetries of the *worlds* they represent (except in special cases where the metric has Killing vectors, in which case a small *subgroup* of the diffeomorphism group will be a symmetry group of the solution). Contrast this with case of quantum mechanics. Starting with the tensor product Hilbert space formalism, permutations of particle labels *are* symmetries of the theory. For example, if Ψ_{12} is a physically possible state of two identical particles, then so is Ψ_{21}. But this is *because* $\Psi_{12} = (-)\Psi_{21}$. Permutations are symmetries of every *solution* of the theory, and that is how and why they are also symmetries of the theory. And if we consider instead the Fock space formalism, there simply are no particle labels to be permuted.

[45] The distinction between symmetries of theories, and of worlds, is discussed by Belot (2003, §4.2) and Ismael and van Fraassen (2003, p. 378); I am grateful to Paul Mainwood for emphasizing it to me.

I conclude that the diffeomorphism invariance of GR and the permutation invariance of quantum mechanics are not formally analogous, and do not generate the same interpretative problems concerning the individuation of the putative subject matter of the theories. This is not, of course, to say that there are no similarities. To treat the points of the manifold of a solution of Einstein's field equations as akin to variables rather than names (a stance Maudlin (1989) has dubbed 'Ramseyfying substantivalism'), would be to regard the models of GR as akin to a Fock space description; it involves regarding distinct mathematical models as strictly logically equivalent and as only syntactically distinct. Conversely, if one can really think of the states associated with the occupation numbers of the Fock space formalism as *genuinely occupied by objects*, objects that are individuated by the properties attributed to them by these states (*cf.* Stachel's talk of electrons in an atom being individuated by their different quantum numbers), then what is to stop us naming them? Such names would not correspond to the labels of any quantum formalism, but they would correspond to the informal talk of physicists, who are happy to talk of the *the* K-shell electron, for example, (assuming there is only one) or of *the* particle in the left-hand wing of the EPR apparatus.

In the next section I consider briefly whether this talk is really permissible. Before turning to that final topic, let us briefly consider French and Rickles' assessment of Stachel's analogy between points and particles. They write:

> Stachel…understands the non-individuality of particles as their being individuated 'entirely in terms of the relational structures in which they are embedded'…But then it is not clear what metaphysical work the notion of 'non-individuality' is doing, when we still have 'objects' which are represented by standard set theory (and this is precisely the criticism that can be levelled against attempts to import non-individuality into the spacetime context)…
>
> Again the alternative, 'middle way' is to drop objects out of the ontology entirely, regarding both spacetime and particles in structural terms. Indeed, this appears to be the more appropriate way of understanding both Stachel's talk of individuating objects 'entirely in terms of relational structures in which they are embedded'…However, rather than thinking of objects being individuated, we suggest they should be thought of as being structurally constituted in the first place. In other words, it is relational structures which are regarded as metaphysically primary and the objects as secondary or 'emergent'. (2003, p. 235)

In light of this section and the previous ones, the response I advocate to these suggestions should be clear. Why should non-individuality do any more work than (be anything more than) the denial of *primitive* individuality and haecceitistic differences? We have yet to be given a reason to think that standard set theory should *not* apply, at least to spacetime points.

8 Identical Particles and Identity Over Time.

A simple story is often retold in elementary discussion of quantum statistics. Suppose that we have two identical particles, a and b, and just two possible single-particle states L and R. We are told that if one 'thinks classically', one should expect four distinct states for the joint system:

1. $L(a)L(b)$
2. $L(a)R(b)$
3. $R(a)L(b)$
4. $R(a)R(b)$

And if we further suppose that each of these states is equally probable, then we get an instance of Maxwell–Boltzmann statistics: the possibility according to which one particle is in state L while one is in state R is twice as likely as each of the two possibilities according to which the particles are in the same state.

But quantum particles obey either Fermi–Dirac or Bose–Einstein rather than Maxwell–Boltzmann statistics. For example, in the Bose–Einstein case, the possibilities:

1. $L(2)R(0)$: two particles are in state L
2. $L(1)R(1)$: one particle is in state L and one particle is in state R
3. $L(0)R(2)$: two particles are in state R

are all equally likely. It seems that it is the supposition that $L(a)R(b)$ and $R(a)L(b)$ are distinct possibilities that led us to the wrong, classical, statistics. But isn't equating the possibilities $L(a)R(b)$ and $R(a)L(b)$ simply anti-haecceitism? It seems that the non-existence of haecceitistic differences between states involving identical quantum particles suffices to explain quantum statistics. Perhaps quantum statistics recommends exactly the same interpretative move as the hole argument in general relativity after all.

Unfortunately, things are not so simple. First, anti-haecceitism is compatible with Maxwell–Boltzmann statistics, as is shown by Huggett (1999b). What is required, if one is to obtain such statistics while denying haecceitistic differences, is, crudely put, a continuum of possible microstates relative to a countable number of particles (*ibid.*, §IV). It might be held that this result is just as well, for some see the Gibbs Paradox as motivating the denial of haecceitistic differences even in classical statistical mechanics (see, e.g., Saunders 2003, p. 302), although this remains a matter of controversy. The difference between quantum statistical systems, and classical statistical systems, is then to be seen as arising precisely because, in the quantum case, one does not have a continuum of microstates.

I agree with this as far as it goes, but it seems that a puzzle remains. Why is the symmetrization postulate imposed; i.e., why, rather than simply stipulating that $L(a)R(b)$ and $R(a)L(b)$ represent the same state, do we take the appropriate quantum state to be $(1/\sqrt{2})\,(L(a)R(b) + R(a)L(b))$?

The answer has to do with a distinction drawn in section 4. Statistics are manifest *over time* in frequencies. To regard the states $L(a)R(b)$ and $R(a)L(b)$ as distinct in the context of the *persisting* particles a and b is not necessarily to sign up to haecceitistic differences. A single solution might involve the instantaneous state $L(1)R(1)$ at two different times t_1 and t_2. If it makes sense to ask whether the particle that occupies state L at t_1 is the same as the particle that occupies the state L at t_2, then we have a legitimate, *non-haecceitistic* reason for distinguishing between the instantaneous states $L(a)R(b)$ and $R(a)L(b)$. According to a point of view that goes back at least to Reichenbach (1956), the difference between classical and quantum statistics bears on such questions concerning *identity over time*, rather than haecceitism.

As a very simple example, consider our two one-particle states L and R, which, let us suppose, at regular time intervals, $t_1, t_2, \ldots$, are instantiated by two 'particles'.[46] Consider the two cases where the frequencies exhibited over time correspond to (i) Maxwell–Boltzmann statistics and (ii) Bose–Einstein statistics. The frequencies of scenario (i) can be explained in terms of a very simple dynamical model. It involves two persisting particles such that (a) at each time the probability of each particle occupying each state is $\frac{1}{2}$ and (b) the likelihood of their occupying each state at each time is independent of which state the other particle occupies at that time. Scenario (ii), on the other hand, is most simply explained by postulating that at each time the three possible instantaneous states $L(2)R(0), L(0)R(2)$ and $L(1)R(1)$ are equally likely; there is no additional fact of the matter concerning whether the 'particle stage' that instantiates L at one time constitutes a stage of the same persisting particle as the particle stage that instantiates L at some other time. Of course, one *can* combine such additional facts about persistence with the Bose–Einstein statistics of scenario (ii). But if one does, the dynamics of the two particles can no longer be independent of each other but must involve 'causal anomalies' if the correct frequencies are to be recovered (Reichenbach 1956, pp. 69-71 (in Castel-

[46]I am indebted, at this point, to a conversation with Nick Huggett.

lani ed.)).[47]

The previous paragraph suggests the following possibility. Perhaps one can view the instantaneous *temporal stages* of quantum particles as genuine individuals, individuated by their sets of quantum numbers. The only problem with introducing particle labels as the names of such objects is that they illicitly introduce primitive trans-temporal identities between the particles that exist at one moment and those that exist at another. We are forced to (anti)symmetrize the state to 'rub out' out these illegitimate trans-temporal identities. The entities that exist at any given instant are not really to be thought of as each in an identical mixed state (it's not the case that every electron is currently equally a part of me, part of you, and part of this page of the paper you are reading). We would then, again, have a strong analogy between spacetime points and the fundamental ontology of identical particle quantum mechanics; both would be instantaneous entities fully individuated by their properties and relations.

The problem with this suggestion is that the non-commutative algebra of observables prevents our interpreting even the instantaneous ontology of quantum mechanics as a determinate set of reflexively defined individuals. The difficulty is that the choice of particular maximal properties to characterize the quantum particles is to a large extent arbitrary. Stachel states that "each electron in an atom can be fully individuated by the set of its quantum numbers" (2002, p. 260). If these quantum numbers are to fully individuate, then the electron's component of spin in some direction must be included, conventionally the 'z'-direction. If such a set of properties really did individuate, we should be able to talk about, e.g., *the* S-shell electron whose spin in the z-direction is $+\hbar/2$. But of course, we cannot, for if we could, symmetry would require that we could also talk about the S-shell electron whose spin in the x-direction is $+\hbar/2$. And we could then ask whether the S-shell electron whose spin in the z-direction is $+\hbar/2$ was the same electron as the S-shell electron whose spin in the x-direction is $+\hbar/2$. But this last question is illegitimate. There being a fact of the matter about its answer would contravene quantum mechanics' violation of Bell's inequalities. Unlike spacetime points in classical general

[47] There a rather interesting application of this type of example to the debate between perdurantists and endurantists. In response to the scenarios involving rotating homogeneous matter that endurantists press against perdurantists, Sider (2001, pp. 224-36) has suggested that one might exploit the Mill–Ramsey–Lewis account of laws to pick out a preferred genidentity relation in favourable cases. (For a comprehensive discussion of arguments concerning rotating homogeneous matter in the context of the perdurantists–endurantist debate, see Butterfield 2004.) Now we can envisage two spatio–temporal Humean mosaics for which the Mill–Ramsey–Lewis prescription yields probabilistic laws involving Maxwell–Boltzman and Bose–Einstein statistics respectively. As described in the paragraph above, the former favours a law formulated in terms of persisting particulars. However, there will be many ways of drawing the lines of persistence consistent with the statistics. Here is a case, then, where the simplest law favours introducing a genidentity relation, but fails to determine which particle stages should be regarded as genidentical.

relativity, quantum particles cannot be thought of as individuated by the relational structures in which they are imbricated. They are not even reflexively defined entities.

9 Two Morals for Quantum Gravity.

In the previous two sections two conclusions were reached concerning identical particles in quantum mechanics. The first was that, since the *states* of identical particle quantum mechanics were permutation invariant, there could be no analogue of the hole argument that involved them. To run an analogue of the hole argument one needs solutions of a permutation invariant *theory* that are not themselves permutation invariant and which are thus interpretable, at least in principle, as representing physically distinct (although only haecceitistically distinct) states of affairs. The second, more tentative, conclusion, was that if a theory involves a non-commutative algebra of observables, then there is at least a *prima facie* problem facing those who would interpret the ontology of the theory as involving a single, determinate set of reflexively defined entities.

Both of these conclusions would appear to be applicable to loop quantum gravity (LQG), the first straightforwardly so. The *states* of loop quantum gravity satisfy the so-called diffeomorphism constraint. This means that they are (3-)diffeomorphism invariant: the states do not distinguish the points of the 3-manifold in terms of which they are, notionally, defined. In LQG the points of the spatial 3-manifold have a status exactly analogous to particle labels in identical particle quantum mechanics.[48]

The space of states that satisfy both the Gauss constraint and the diffeomorphism constraint is spanned by a basis of states that are labelled by *abstract* spin networks, or knots, where a knot is an equivalence class of graphs embedded in a manifold under diffeomorphisms. It is to these states (often, and somewhat confusingly, also referred to simply as spin network states) that popular accounts of LQG typically refer (see Rovelli 2001, pp. 110–1). The nodes of the graph can be thought of as quanta of volume—as elementary chunks of space—and the links as quanta of area separating these volumes. This picture suggests the following thought: might we regard the links and nodes of abstract spin networks as representing genuine

[48]Dean Rickles (2005) suggests that an analogue of the hole argument can be constructed in the context of LQG. The states involved in his construction are spin network states that solve the Gauss constraint but which do not solve the diffeomorphism constraint (or the Hamiltonian constraint). They are therefore *not* solutions of the (quantum) Einstein equations. To claim, therefore, that "Einstein's equation cannot determine where spin-networks are in the manifold" is misleading. Einstein's equation (or rather its quantum version) determines exactly where on the manifold a spin network is: it is smeared all over the manifold in a diffeomorphic-invariant fashion. See Pooley (forthcoming).

entities (i.e., elementary volumes and surfaces of space), entities that are reflexively defined by the network of relations in which they stand?

The obvious worry with this proposal concerns the second conclusion mentioned above. The spin network basis is just one basis for the space of states that satisfy the Gauss constraint. Other possible bases will provide us with a set of states that are not interpretable as networks of volumes and areas (the volume and area operators will not be diagonalized by these other bases). If non-commuting observables do not allow quantum particles to be straightforwardly interpreted as reflexively defined objects, the same will be true of the elementary quanta of loop quantum gravity.

Acknowledgements.

Previous incarnations of this paper were presented at the University of Leeds, the Universidad Autonoma de Barcelona, the LSE, the University of Oxford, and the University of Illinois at Chicago. I am grateful to members of all those audiences, in particular to Joe Melia, Carl Hoefer, Mauricio Suarez, James Ladyman, Michael Redhead, Nick Huggett and Steve Savitt, for comments and criticism. I am also grateful to Brian Leftow, Justin Pniower, Harvey Brown, Simon Saunders, Jan Westerhoff and Paul Mainwood for discussions about related matters, to an anonymous referee, and to the editors of this volume, not least for their patience. I am especially grateful to Jeremy Butterfield for extensive and careful comments on a draft of this paper, which led to numerous improvements and which should have led to more. This paper was composed during the tenure of a British Academy Postdoctoral Fellowship; I gratefully acknowledge the support of the British Academy.

References.

Anandan, J. and H. R. Brown (1995) "On the Reality of Space-time Geometry and the Wavefunction." *Foundations of Physics* **25**: 349–60.
Belot, G. (2001) "The Principle of Sufficient Reason." *Journal of Philosophy* **118**: 55–74.
—————— (2003) "Notes on Symmetries." In K. Brading E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, Cambridge (pp. 393–412).
Belot, G. and J. Earman (1999) "From Metaphysics to Physics." In J. Butterfield and C. Pagonis (eds.), *From Physics to Philosophy*. Cambridge University Press, Cambridge (pp. 166–86).
—————— —————— (2001) "Pre-Socratic Quantum Gravity." In C. Callender

and N. Huggett (eds.) *Physics Meets Philosophy at the Planck Scale*. Cambridge University Press, Cambridge (pp. 213–55).
Benacerraf, P. (1965) "What Numbers Could Not Be." *Philosophical Review* **74**: 47–73.
Brighouse, C. (1994) "Spacetime and Holes." In D. Hull, M. Forbes R. Burian (eds.) *Proceedings of the 1994 Biennial Meeting of the Philosophy of Science Association, Vol. 1*. Philosophy of Science Association, East Lansing, MI (pp. 117–25).
Butterfield, J. N. (1989) "The Hole Truth." *British Journal for the Philosophy of Science* **40**: 1–28.
——— (2004) "On the Persistence of Homogeneous Matter." Available at `http://philsci-archive.pitt.edu/archive/00002381/`.
Dalla Chiara, M. L., R. Giuntini, and D. Krause (1998) "Quasiset Theories for Microobjects: A Comparison." In E. Castellani (ed.), *Interpreting Bodies*. Princeton University Press, Princeton NJ (pp. 142–52).
Demopoulos, W. and M. Friedman (1985) "Bertrand Russell's *The Analysis of Matter*: Its Historical Context and Contemporary Interest." *Philosophy of Science* **52**: 621–39.
Earman, J. (1989) *World Enough and Space-Time: Absolute versus Relational Theories of Space and Time*, MIT Press, Cambridge MA.
——— (2003) "Tracking Down Gauge: An Ode to the Constrained Hamiltonian Formalism." In K. Brading E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, Cambridge (pp. 140–62).
Earman, J. and J. Norton (1987) "What Price Substantivalism? The Hole Story." *British Journal for the Philosophy of Science* **38**: 515–25.
English, J. (1973) "Underdetermination: Craig and Ramsey." *Journal of Philosophy* **70**: 453–62.
French, S. (1989) "Identity and Individuality in Classical and Quantum Physics." *Australasian Journal of Philosophy* **67**: 432–46.
——— (1998) "On the Whithering Away of Physical Objects." In E. Castellani (ed.), *Interpreting Bodies*. Princeton University Press, Princeton NJ (pp. 93–113).
French, S. and J. Ladyman (2003) "Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure." *Synthese* **136**: 31–56.
French, S. and M.L.G. Redhead (1988) "Quantum Physics and the Identity of Indiscernibles." *British Journal for the Philosophy of Science* **39**: 233–46.
French, S. and D. Rickles (2003) "Understanding Permutation Symmetry." In K. Brading and E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, Cambridge (pp. 212–38).
Hilbert, D. and P. Bernays (1934) *Grundlagen der Mathematik, Vol. 1*. Springer-Verlag: Berlin.
Hoefer, C. (1996) "The Metaphysics of Space-time Substantivalism." *Journal of Philosophy* **93**: 5–27.

——— (2000) "Energy conservation in GTR." *Studies in History and Philosophy of Modern Physics* **31**: 187–99.
Huggett, N. (1999) "Atomic Metaphysics." *Journal of Philosophy* **96**: 5–24.
Ismael, J. and B. C. van Fraassen (2003) "Symmetry as a Guide to Superfluous Theoretical Structure." In K. Brading and E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, Cambridge (pp. 371–92).
Jones, R. (1991) "Realism about What?" *Philosophy of Science* **58**, 185–202.
Kaplan, D. (1975) "How to Russell a Frege–Church." *Journal of Philosophy* **72**: 716–29.
——— (1979) "Transworld Heir Lines." In M. J. Loux (ed.), *The Possible and the Actual*. Cornell University Press.
Ketland, J. (2004) "Empirical Adequacy and Ramsification." *British Journal for the Philosophy of Science* **55**: 287–300.
Krause, D. (1992) "On a Quasi-set Theory'. *Notre Dame Journal of Formal Logic* **33**: 402–11.
Ladyman, J. (1998) "What is Structural Realism?" *Studies in the History and Philosophy of Science* **29**: 409–24.
Laudan, L. (1981) "A Confutation of Convergent Realism", *Philosophy of Science* **48**: 19–48.
Lewis, D. K. (1986) *On the Plurality of Worlds*. Oxford: Basil Blackwell.
Maudlin, T. (1989) "The Essence of Space-time." In A. Fine J. Leplin (eds.), *Proceedings of the 1988 Biennial Meeting of the Philosophy of Science Association, Vol. 2*. Philosophy of Science Association, East Lansing, Michigan (pp. 82–91).
——— (1990) "Substances and Space-time: What Aristotle would have said to Einstein." *Studies in History and Philosophy of Science* **21**: 531–61.
Muller, F. A. (1997a) "The Equivalence Myth of Quantum Mechanics—Part I." *Studies In History and Philosophy of Modern Physics* **28**: 35–61.
——— (1997b) "The Equivalence Myth of Quantum Mechanics—Part II." *Studies In History and Philosophy of Modern Physics* **28**: 219–47.
Newman, M. H. A. (1928) "Mr Russell's Causal Theory of Perception." *Mind* **37**: 137–48.
Norton, J. D. (2000) "What can we Learn about the Ontology of Space and Time from the Theory of Relativity?" `http://philsci-archive.pitt.edu/archive/00000138/`. Paper presented at the International Ontology Congress, San Sebastian, Spain, October 2000.
Parsons, C. (2004) "Structuralism and Metaphysics." *Philosophical Quarterly* **54**: 56–77.
Pooley, O. (forthcoming) "A hole revolution, or are we back where we started?" To appear in *Studies In History and Philosophy of Modern Physics*.
Pooley, O. (in preparation) "What's so Special about General Relativity? Some Skeptical Remarks about the Status of General Covariance." Paper presented at

the ESF Conference on Philosophical and Foundational Issues in Spacetime Theories, Oxford, March 2004.
Quine, W. V. (1986) *Philosophy of Logic* (2nd edition). Harvard University Press, Cambridge, Mass.
Reichenbach, H. (1956) *The Direction of Time*. University of California Press.
Rickles, D. P. (forthcoming) "A New Spin on the Hole Argument." To appear in *Studies in the History and Philosophy of Modern Physics* (September).
Rovelli, C. (1997) "Halfway through the Woods: Contemporary Research on Space and Time." In J. Earman J. Norton (eds.), *The Cosmos of Science*. University of Pittsburgh Press, Pittsburgh (pp. 180–223).
——— (2001). "Quantum Spacetime: What do we Know?" In C. Callender and N. Huggett (eds.), *Physics Meets Philosophy at the Planck Scale*. Cambridge: Cambridge University Press (pp. 101-122).
Rynasiewicz, R. A. (1994) "The lessons of the Hole Argument." *British Journal for the Philosophy of Science* **45**: 407–36.
Saunders, S. W. (2002) "Indiscernibles, General Covariance, and Other Symmetries." In A. Ashtekar, D. Howard, J. Renn, S. Sarkar A. Shimony (eds.), *Revisiting the Foundations of Relativistic Physics. Festschrift in Honour of John Stachel.* Kluwer, Dordrecht. Preprint available: `http://philsci-archive.pitt.edu/archive/00000459`.
——— (2003) "Physics and Leibniz's Principles." In K. Brading E. Castellani (eds.), *Symmetries in Physics: Philosophical Reflections*. Cambridge University Press, Cambridge (pp. 289–307).
Shapiro, S. (1997) *Philosophy of Mathematics: Structure and Ontology*. Oxford University Press, Oxford.
Sider, T. (2001) *Four Dimensionalism: An Ontology of Persistence and Time* Oxford University Press, Oxford and New York.
Smolin, L. (2000) "The Present Moment in Quantum Cosmology: Challenges to the Arguments for the Elimination of Time." In R. Durie (ed.), *Time and the Instant*. Clinamen Press, Manchester (pp. 112–43).
Stachel, J. (2002) "'The Relations Between Things' Versus 'The Things Between Relations': The Deeper Meaning of the Hole Argument." In D. B. Malament (ed.), *Reading Natural Philosophy. Essays in the History and Philosophy of Science and Mathematics*. Open Court, Chicago (pp. 231–66).
Teller, P. (1998) "Quantum Mechanics and Haecceities." In E. Castellani (ed.), *Interpreting Bodies*. Princeton University Press, Princeton NJ (pp. 114–141).
——— (2001) "The Ins and Outs of Counterfactual Switching." *Noûs* **35**: 365–93.
van Fraassen, B. C. (1991) *Quantum Mechanics: An Empiricist View*. Oxford University Press, Oxford.

Worall, J. (1989) “Structural realism: The Best of Both Worlds?” *Dialectica* **43**: 99–124.
Worrall, J. and E. Zahar (2001) “Ramseyfication and Structural Realism.” In E. Zahar, *Poincaré’s Philosophy: From Conventionalism to Phenomenology*. Open Court, Chicago and La Salle (IL). (Chapter Appendix IV: pp. 236–51).
Zahar, E. (1994) “Poincaré’s Structural Realism and his Logic of Discovery.” In G. Heinzmann *et al.* (eds.), *Henri Poincaré: Akten Des Internationale Kongresses in Nancy*.