Challenging the Spacetime Structuralist

Christian Wüthrich[*]

30 January 2008

Word count: 4989

Abstract

Structural realist interpretations of generally relativistic spacetimes have recently come to enjoy a remarkable degree of popularity among philosophers. I present a challenge to these structuralist interpretations that arises from considering cosmological models in general relativity. As a consequence of their high degree of spacetime symmetry, these models resist a structuralist interpretation. I then evaluate the various strategies available to the structuralist to react to this challenge.

1 Introduction

Most authors who argue for a structural realist interpretation of spacetime find their motivation to do so in their belief that the hole argument in general

relativity naturally leads to this view.[1] Often, proponents of this persuasion see their position as a *via media*, or, in Mauro Dorato's (2000) words, a *tertium quid*, between spacetime substantivalism and relationism. Structural realists about spacetime side with relationists in their conviction that the fundamental ontology of spacetime consists of relational complexes rather than individual objects such as spacetime points, but join substantivalists in their unabashed realist stance about spacetime. Some authors in this literature, such as Dean Rickles and Steven French (2006), go as far as to claim that the "sophisticated" substantivalism and the version of relationism most prominently defended today, both collapse into spacetime structuralism.[2]

An altogether different motivation can be found in Jonathan Bain (2006) who infers spacetime structuralism from the fact that general relativity can not only be expressed in the standard tensor formalism, but can also be cast in at least three alternative formalisms: twistor theory, Einstein algebras, and geometric algebra. For Bain, since these four formalisms afford nearly intertranslatable expressions of general relativity, it would be foolish to be wedded to the ontological commitments of any one of them. Instead, we should only ontologically commit to those structures necessary to support the mathematical representations of the physical systems and their dynamics, while the exact formulation of what this structure is depends on the formalism one adopts. Bain, thus, identifies what is common to all representative frameworks of the same theory, the "structure," with what we ought to be realists about.

After fixing what I take to be a sensible characterization of structural realism in Section 2, I will press on in Section 3 to propose a challenge to a structural realist interpretation of spacetime that arises from highly (spatially) symmetrical models of general relativity. In these cases, since according to the structural realist spacetime points only inherit their properties *and even their individuality* from the structure in which they partake, it turns out that the points' individuality can no longer be grounded. This result is disastrous as it suggests that the (spatial) universe consists of nothing but one lonely point. Section 4 addresses and discusses various strategies that may be pursued in resisting this unattractive implication.

---

[1]Cf. for instance Dorato (2000), Rickles and French (2006, 4), and Stachel (2006, 56f).
[2]Not everyone agrees, see e.g. Pooley (2006).

2   Defining Structure

If structural realism can be understood, roughly, as the position according to which one ought to be realist about the "structure" of a scientific theory, then this characterization remains empty unless some precisification of the term "structure" is offered. There is considerable disagreement on this point among philosophers who consider themselves structural realists, but Rickles and French (2006, 25) are certainly right in their insistence that structural realism is not a monolithic position, and neither should it be. At a quite general level, structural realism demands that an ontological commitment ought to be made not to individuals as the primary, most fundamental constituents of reality, but rather a complex of relations. John Stachel (2006, Sec. 3.1) identifies three basic types of structural realism, ranging from *timid* ("there are objects and relations, with the objects being primary and the relations being secondary"), to *traditional* ("there are objects and relations, such that relations are primary and objects are secondary"), to *radical* ("there are only relations, without underlying or accompanying relata"). Call the first two forms of structural realism *non-eliminative* and the third *eliminative* as it eliminates objects from the basic ontology. Taken at face value, eliminative structural realism as formulated by the third Stachelian option is clearly incoherent, despite the fact that the rhetoric of some British structuralists such as French and James Ladyman (2003) indeed suggests that they defend this extreme position. Given the unintelligibility of this position as taken at face value, I suspect that such an attribution is plainly false and the position really defended by these British structuralists differs from the one painted by their rhetoric. More charitably read, it perhaps means that any given relation's relata can in turn be fully individuated by relations. Like Stachel, "I would not want to be bound by the claim that this is always the case." (2006, 54)

Apart from those varieties of structural realism listed by Stachel (2006), there exists another, in my view quite sensible, characterization of the position. Instead of insisting that either the objects or the relations must be primary—where the disjunction is understood to be exclusive—, i.e. instead of defending either timid or traditional structural realism, one can refuse to award ontological primacy to either side and consider relations and relata ex aequo necessary constituents of either the structure in the world or the structural knowledge the structuralist hopes to attain. Michael Esfeld (2004) and Esfeld and Vincent Lam (2008) have explicitly endorsed this what they

3

dub *modest* version of structural realism, and Oliver Pooley (2006, 98) has argued that the same view—he calls it the *no priority view*—is the most plausible form of structural realism. I prefer to call it *balanced* structural realism exactly because of the ontological balance it strikes between objects and relations. For present purposes, I shall presume that an attractive version of structuralism can be presented without insisting either on privileging relations over relata or on eliminating objects altogether. More precisely, let me assume the following definition: A *structure* $\mathfrak{S}$ is a pair $\langle \mathfrak{O}, \mathfrak{R} \rangle$ which consists of a non-empty set of concrete relations $\mathfrak{R}$ ("ideology") as well as a non-empty set of physical relata $\mathfrak{O}$ ("ontology"), the *domain* of $\mathfrak{S}$.

With this definition of structure under our belt, balanced structural realism as defended by Esfeld (2004) and by Esfeld and Lam (2008) then amounts to an endorsement of the following philosophical position. The "structural" aspects of a scientific theory, toward which we should entertain a realist attitude, relevantly capture the structure of the external world. The fundamental scientific theories thus reveal the structure of the external world, where structure is used in the sense of the definition in the previous paragraph. The objects $x \in \mathfrak{O}$, i.e. the things which exemplify the relations $R \in \mathfrak{R}$, do not have any intrinsic properties, but *only relational ones.* So what is really there according to the structural realist is a network of relations among objects which do not possess any intrinsic properties but are purely defined by their "place" in $\mathfrak{S}$. How these structural aspects are identified in a given theory is, of course, a highly non-trivial matter and will largely depend on one's interpretation of the theory at stake.

In the case of structural realism as the preferred interpretation of generally relativistic spacetimes, the spacetime is the structure about which the structuralist wants to be realist. A spacetime in general relativity is a triple $\langle \mathcal{M}, g_{\mu\nu}, T_{\mu\nu} \rangle$ of a four-dimensional Lorentzian manifold $\mathcal{M}$, a metric tensor $g_{\mu\nu}$, and a stress-energy tensor $T_{\mu\nu}$ that satisfies Einstein's field equations, $G_{\mu\nu}[g_{\mu\nu}] = 8\pi T_{\mu\nu}$ where $G_{\mu\nu}$ is a tensorial functional of $g_{\mu\nu}$, the so-called Einstein tensor. Since the triples $\langle \mathcal{M}, g_{\mu\nu}, T_{\mu\nu} \rangle$ are also the models of general relativity in the model-theoretic sense, interpreting spacetime structurally really means offering a structural realist interpretation of general relativity. The ontology $\mathfrak{O}$ of the structure at stake, in this case, will just consist of the points of the manifold $\mathcal{M}$. The points in $\mathcal{M}$ will then be the relata that stand in certain physically admissible relations to one another. The structural realist insists that the manifold points possess *quiddity*, i.e. some essence that it is their nature to possess, but certainly no *haecceity*, i.e. "primitive

thisness" or some particular characteristic of what it is to be that particular point, as such haecceities would surely be intrinsic—whatever else they might be. More generally, a structural realist will restrict herself to an ideology of purely qualitative properties. A property is called *qualitative* just in case its exemplification does not depend upon the existence of any particular individual. Haecceistic properties are thus excluded. It is exactly this feature which makes structural realism such an attractive response to the hole argument.

It is important to stress at this point that the balanced structural realist rejects that members of the ontology exemplify any intrinsic properties. Although there is considerable debate in metaphysics as to which properties ought to qualify as "intrinsic," it will be sufficient for our purpose to assume that *intrinsic* properties are all and only those qualitative properties whose exemplification is independent of the existence or non-existence of other contingent objects. Intrinsic properties can thus be attributed independently of accompaniment or loneliness. All non-intrinsic properties are either *extrinsic* if they are monadic, or else relational. Any structural realist who pays more than lip service to the structuralist aspect of her position should repudiate intrinsic properties as a vestige of an object-based ontology, taking to heart the structuralist dictum that the individuation of the members of $\mathfrak{O}$ occurs by virtue of their being embedded in a structure. Consequently, the structural realist will also reject the thesis of Humean supervenience, which assumes a geometry of spatio-temporally related points endowed with perfectly natural *intrinsic* properties. In fact, it is the renunciation of Humean supervenience that seems to motivate much of Esfeld's structural realism.

Let me attempt to capture the idea of structuralism more formally. What the balanced structural realist as characterized above demands is that the world is fundamentally described by a structure $\mathfrak{S}$ with objects which exemplify the intra-structural relational properties, but no other properties. Any other properties would introduce a unwanted reference to something beyond the purely structural. These intra-structural relational properties are exactly those which are invariant under automorphisms of $\mathfrak{S}$, where an *automorphism f* is defined as a map from a domain A onto itself which preserves the structure of A (i.e. is an isomorphism of a set onto itself). A property is *invariant under an automorphism* when any element $a \in A$ has the property iff its image $f(a)$ has it. Thus, the set $\mathfrak{R}$, which figures in the definition of a structure, and consequently in the definition of structural realism, contains only automorphically invariant relational properties. Furthermore, let us assume that the set of properties does not range over properties which

are in some philosophically intricate sense pathological, such as disjunctive or gruesome properties. I am painfully aware that ingenious metaphysical moves will have to be performed before we can hope to attain some robust criteria that would allow us to exclude pathological properties on a principled basis. For present purposes, I cowardly assume to have resolved these issues.

3   The Challenge from Cosmology

It can be shown that structural realism as characterized in the previous section and as applied to the present spacetime context suffers from serious difficulties in accommodating highly symmetric spacetimes. For these particular spacetime solutions with a high degree of symmetry, a devastating argument in full analogy to the one run by Jukka Keränen (2001) against structuralism in the philosophy of mathematics can be given. To this end, consider the highly symmetric cosmological standard model in classical general relativity, the so-called Friedmann-Lemaître-Robertson-Walker (FLRW) spacetimes.

Modern cosmology relies on the so-called *cosmological principle*, according to which no position in space—including ours—is privileged in any way. It is generally interpreted to mean that the universe must exhibit spatial homogeneity and isotropy, at least approximately so. A spacetime is *spatially homogeneous* just in case there exists a one-parameter family of spacelike hypersurfaces Σ_t such that for all t and for any points $p, q \in \Sigma_t$, there exists an isometry f of the metric $g_{\mu\nu}$ on $\mathcal{M}$ with $f(p) = q$. Σ_t is thus a foliation of spacetime, i.e. a partition of $\mathcal{M}$ into a union of disjoint subsets Σ_t. A spacetime is *spatially isotropic* just in case, roughly, it is impossible to find a geometrically preferred spatial direction in any of the spacelike hypersurfaces Σ_t. A theorem due to Walker (1944) establishes that there exists exactly one foliation that preserves the spatial symmetries. This preferred foliation can be labelled by a time coordinate $t \in \,]a, b[\, \subseteq \mathbb{R}$. A time t thus privileged is called a *cosmological time*. An isometry f is an isomorphism of $\mathcal{M}$ onto itself such that the metric structure is preserved. Isometries of $g_{\mu\nu}$ on $\mathcal{M}$ form a group called the *isometry group* of $\mathcal{M}$. Spatial homogeneity thus means that for every "snapshot" of the universe, there exists a map from any point in the universe to any other that preserves the structure of the universe at that time, and in particular the metric structure. In total, these symmetries—homogeneity and isotropy—are mathematically encoded in the

action of a group of isometries on $\mathcal{M}$ with spacelike hypersurfaces as the group's surfaces of transitivity.[3] Any point on any spacelike hypersurface is thus equivalent to any other point on the same hypersurface. In particular, the symmetries imply that the spatial curvature of all the spacelike hypersurfaces Σ_t of the preferred foliation is constant. Spacetimes $\langle \mathcal{M}, g_{\mu\nu}, T_{\mu\nu} \rangle$ with exact spatial homogeneity and isotropy are the FLRW spacetimes mentioned above. The challenge I am about to mount against spacetime structural realism exploits these exact spatial symmetries.

The symmetries of isotropy and homogeneity thus demand that the universe be the "same in every location," spatially understood, with respect to all physically relevant qualities. This idea can be formalized as the proposition, valid for all $p, q \in \Sigma_t \subset \mathcal{M}$, that

$$\forall F \in \mathfrak{P} \ (Fp \leftrightarrow Fq), \tag{1}$$

where $\mathfrak{P}$ is the set of admissible physical properties with respect to which points in Σ_t must be the "same." Proposition (1) is valid for any Σ_t of a FLRW-spacetime. For the spacetime structuralist described in the previous section, $\forall F$ must range over all and only over automorphically invariant relational properties that do not depend in their exemplification on the existence of any particular individual.

The individuality of the objects in the ontology $\mathfrak{O}$ must be ascertained by an identity criterion with whose help objects can be distinguished. The most prominent such criterion is of course Leibniz's Principle of the Identity of Indiscernibles, or a modernized version thereof. The core idea of this family of principles is to utilize distinction between objects in terms of the properties they exemplify as a criterion to individuate them. In more formal words,

$$\forall F \in \mathfrak{Q} \ (Fx \leftrightarrow Fy) \rightarrow x = y. \tag{2}$$

Varieties of this Principle of the Identity of Indiscernibles (PII) typically differ in what is taken to be in $\mathfrak{Q}$. French (2006) suggests the following three basic possibilities: (i) $\forall F$ ranges over all possible properties, (ii) $\forall F$ ranges over all possible properties except spatio-temporal ones, and (iii) $\forall F$ ranges only over intrinsic properties. The properties here at stake, even in the weakest version (i), are all qualitative properties.

---

[3]The *surface of transitivity* of a group G acting on the manifold $\mathcal{M}$ is the set $\Omega \subset \mathcal{M}$ of all points such that the group action $G \times \Omega \rightarrow \Omega$ possesses only a single group orbit, i.e. for every pair of elements $x, y \in \Omega$, there exists a group element $g \in G$ such that $gx = y$.

The strongest version of PII, viz. PII(iii), claims that no two individuals can possess the same intrinsic properties. This principle is clearly violated in classical physics, where distinct particles may be regarded as indistinguishable as far as intrinsic properties are concerned. Max Black (1952) has proposed a counterexample against such a strong version of PII by placing two indistinguishable spheres in an otherwise empty universe. Assuming that we placed *two* individual spheres into the vacuum universe, Black's example not only violates PII(iii), but also PII(ii). Does it also violate PII(i)? Not as long as the universe into which the spheres are put is interpreted as a *fixed* background of topology $\mathbb{R}^3$. In this case, the set of properties attributed to the two spheres are not entirely identical: at least properties based on their spatio-temporal location will not coincide. So at least the weakest form of PII is usually taken to be valid in classical physics, where spatio-temporal trajectories of rigid bodies do not overlap.

The structural realist about spacetime must re-interpret PII to adapt it to her purpose. The only acceptable version of PII for her, clearly, holds that $\forall F$ ranges over the set of automorphically invariant relations, excluding any intrinsic properties. Now the analogue of Keränen's argument (2001) can be derived easily: for an ontology consisting of all points of a spacelike hypersurface Σ_t of a spatially homogeneous and isotropic spacetime, we get from (1) and (2) by modus ponens that $x = y$ iff $\mathfrak{Q} \subseteq \mathfrak{P}$, where $\mathfrak{Q}$ is the set of all structurally admissible properties and $\mathfrak{P}$ the set of all physical properties as admitted by the cosmological principle. Importantly, both $\mathfrak{Q}$ and $\mathfrak{P}$ only consist of automorphically invariant relational properties, thus ascertaining $\mathfrak{Q} \subseteq \mathfrak{P}$. But since x and y have been arbitrary elements in $\Sigma_t = \mathfrak{O}$, all points of Σ_t coincide and there is only one point in Σ_t. In other words, the universe consists of one point only! Since the group of automorphisms of Σ_t are the isometries of Σ_t, and the group of isometries of Σ_t is transitive, i.e. any point in Σ_t can be moved into any other point in Σ_t, all points in Σ_t must share the same properties. But if all points in Σ_t share the same automorphically invariant properties, they can only constitute one individual according to even the weakest form of the Principle of the Identity of Indiscernibles, PII(i).

In other words, the structural realist cannot distinguish between the elements of $\mathfrak{O}$. But if the objects in $\mathfrak{O}$ cannot be distinguished, they must be identified according to (2). However, when every other object that there might have been must be identified with one particular object, then we say

that there *is* only one object. Formally speaking,

$$\exists!x \left[x \in \mathfrak{O} \right] \Leftrightarrow_{\text{def}} \exists x \left[x \in \mathfrak{O} \wedge \forall y \left(y \in \mathfrak{O} \to x = y \right) \right]. \qquad (3)$$

Thus, we can see that the numerical plurality of the elements in the ontology cannot be grounded by the relata's being situated in a relational structure if this structure exhibits a high degree of symmetry, as is the case in the FLRW cosmological models of general relativity. The point(s?) of Σ_t cannot gain their individuality from the relations which they exemplify. Thus, they don't fulfill the balanced programme of structural spacetime realism. Call this result the *abyssal embarrassment* for the spacetime structuralist.

4 Exit Strategies for the Structuralist

Can the spacetime structuralist recover from the abyssal embarrassment? If so, what are the strategies at her disposal to cope with this result? The most immediate possibility would be to simply redefine structural realism and to hope that the recharacterization does not suffer from the abyssal embarrassment. A comprehensive analysis of this move is beyond the confines of the present essay. Suffice it to say here that I suspect that most interesting versions of structural realist attempt to interpret generally relativistic spacetime will be subject to this result. Apart from changing the discourse in this manner, I see at least four defensible strategies:

1. deny the relevance of homogeneous cosmological models (and similarly symmetric spacetimes);

2. reject PII(i) *even for classical physics* and urge its replacement with another criterion of individuation;

3. claim that PII(i) is inapplicable to the case at hand because no criterion of individuation is needed at all; or

4. argue that we have misidentified the spacetime structure, i.e. a spacetime should not be identified with a manifold of points endowed with certain physical fields at all.

The first path is motivated by the realization that these cosmological models are highly idealized. According to this view, structuralism is still

true of our complicated world despite its failure in these stylized, simple models. Consequently, it denies the relevance of these highly symmetric spacetimes on which the challenge relies. Such denial could occur on the basis of a measure-theoretic consideration. The symmetric spacetimes of the type FLRW, this defence will run, are of measure zero in the space of solutions of the field equations of classical general relativity and thus of not much concern. Since the Einstein equations have not been solved in their full generality and since this space is therefore not explicitly known, this line of argument incurs a promissory note recording the debt of producing this space including a natural measure defined on it. Even if we accept this promissory note for the time being, however, the contention that the abyssal embarrassment will only arise in almost no possible worlds is not implied. The structural realist will still have to establish that the displeasingly symmetric spacetimes are in fact of measure zero. Even if we grant this much this line of defence would be rather inelegant, as it is simply not the case that e.g. the FLRW models are irrelevant *regardless of their measure in the generally relativistic solution space.* An advocate of this response will then strictly speaking be a spacetime structuralist with respect to an empirically adequate proper subset of the set of models of the full theory. It would fall short of establishing that the full theory should be interpreted structurally. Even worse, this line of defence does not signficantly alleviate the pain that the abyssal embarrassment inflicts on the spacetime structuralist; after all, it leaves exposed the complete failure of the structural approach to accommodate what is arguably the most important family of spacetimes in general relativity.

The second strategy rejects PII(i) as a valid criterion of individuation and seeks to supplant it. The first answer in this vein builds on an inversion of the objection aired in the previous paragraph and denies that PII(i) is necessarily true. It insists that its truth is contingent and if its application in the context of fundamental physics leads to absurd results, it should be given up as a criterion of individuation and be replaced by a more sensible alternative. In this sense, the above argument shows that PII(i) leads to an abyssally embarrassing conclusion for some highly symmetric spacetimes. This fact can then be taken as *evidence* that the principle must be rejected as a criterion of individuation in those possible worlds in which it thus fails. This is exactly what it means for the principle to be only contingent: while it is true in almost all possible worlds, it is false in some possible worlds, which (may) happen to be of measure zero. To find out whether the principle is true

in the actual world is thus an empirical matter. As we arguably live in a spatially inhomogeneous universe, it will likely turn out to be true in the actual world. This escape, however, must explain why a metaphysical principle of individuation should not be universally valid no matter what. In particular, if structuralism is offered as a general interpretation of generally relativistic spacetimes, then it seems unacceptable to entertain different criteria of individuation with the range of possibilities that are allegedly captured by the same interpretation.

Simon Saunders (2003) offers another, and more viable, variant of the second strategy. He proposes to find, or construct, a symmetric, but irreflexive relation $S \in \mathfrak{R}$ such as to render the objects "weakly discernible" and hence save them from being identified. He calls the principle based on weak discernability the *Principle of the Identity of Indiscriminables*. The attraction of this approach is that this irreflexivity can ground the object's individuality without recourse to some sort of primitive thisness. However, it faces at least two other difficulties. First, it will not be trivial to find such a relation defined on a homogeneous spacelike hypersurface Σ_t without disturbing its homogeneity, if homogeneity is understood as defined in (1). Second, in order to appeal to such relations, an individuation of objects must already be presupposed: how can I know that there are at least two objects such that an irreflexive relation can be exemplified on the elements of Σ_t? I do not see how this suspicion of circularity can be dispelled. Thus, the resolution of the analogous argument by Keränen (2001) as proposed by Ladyman (2005), which is based on the Principle of the Identity of Indiscriminables, is not available here.

Of course there remains at least a third possibility for enacting an exit strategy of type two. Esfeld (2004, 603) explicitly leaves open the possibility that the objects have non-qualitative properties such as primitive thisness. These haecceities seem to deliver a last resort if everything else failed. If the individuation of objects is based on them, they become ineliminable members of the set of properties that the structural realist must admit. But haecceities are non-relational properties, and certainly not automorphically invariant. Therefore, haecceities cannot be an attractive option for the structural realist. Furthermore, unobservable haecceities re-open the "gap between metaphysics and epistemology" that many structural realists such as Esfeld are so anxious to close.

Esfeld (2004) himself endorses an escape along the fourth line of defence, i.e. denying that *any* criterion of individuation is needed at all, when he

states that "[t]here is no need for the [objects] of which [relational] properties can be predicated to be distinct individuals." (p. 611) He reaffirms this attitude two pages later when he writes that "[t]he argument of this paper accepts that relations require things that stand in the relations (although these things do not have to be individuals, and they need not have intrinsic properties)." (p. 613) Insisting that objects which stand in relations need not be individuals seems highly counterintuitive: should there not be a fact of the matter whether two "places" x and y in a structure $\mathfrak{S}$ are identical or not? In other words, should the statement $x = y$ not have a definite truth value? It should, for otherwise the structuralism defended collapses into the eliminativist version that Esfeld explicitly rejects. A balanced structuralist cannot deny that there are facts of the matter as to whether x and y are identical. At the very least, the advocate of the no-criterion-of-identity-is-needed-for-balanced-structuralism position owes an account of why such a criterion is not needed *even for the balanced version.*

In a brief reaction to my challenge, Esfeld and Lam (2008, 32) criticized what they take to be my metaphysical position as a back-handed way of demanding that the individuality of the ontologically prior relata must be grounded in intrinsic properties. I respond with an emphatic "no." The balanced spacetime structuralist needs identity criteria to individuate both objects and relations. In general, no intrinsic properties are required at all in order to individuate objects. I acknowledge that individuating objects and relations is in general possible within the structural realist programme, but runs into trouble in highly symmetric cases.[4] The solution proposed by Esfeld and Lam (2008, 33) is to accept numerical plurality as a primitive, and thus to acknowledge that this is the only way of individuating fundamental physical objects. Although I agree with them that primitive numerical plurality is superior to primitive thisness, I find this answer highly unsatisfactory. Particularly since I believe there to be a more promising resolution available to the spacetime structuralist.

---

[4]I have not discussed the individuation of properties. It seems obvious to me, however, that this will be equally necessary for a structure $\mathfrak{S}$ to be intelligible. But it might seem less problematic if relata are only related by *one* relation. For a structural realist, however, this would resurrect a similar worry as the one exposed by the challenge: will it still be possible in general to individuate objects in the ontology with the help of only one relation? The answer is yes, although in order to be able to individuate objects, symmetric structures must be excluded.

The final reply claims that PII(i) does not apply to the points of Σ_t because these are perhaps not the individuals that take the place of the relata in the fundamental spacetime structure. As it maintains that the objects in $\mathfrak{O}$ of a structure $\mathfrak{S}$ must be individuated, and since the manifold points in Σ_t cannot be individuals as shown by the abyssal embarrassment, it must consequently deny that the manifold points in Σ_t can play the role of the relata in $\mathfrak{S}$. Therefore, manifold points should not be interpreted as places in the fundamental spacetime structure. In other words, $\Sigma_t \neq \mathfrak{O}$. Thus, the structure $\mathfrak{S}$ has been misidentified in the first place. In order to correct this error, one might take a hint from Bain (2006) and seek to reformulate general relativity in terms of Einstein algebras, or of sheaves, or of something else entirely, and then reidentify the structure, i.e. the ontology and the ideology, of general relativity, and hope that this time the individuality of the objects in the fundamental ontology does indeed derive from their embedding in the relational structure, even in the case of highly symmetric spacetimes. Of course, this sketch of a proposed resolution is entirely programmatic. Much more work needs to be done to flesh out how one can identify the individuals which take the place of the relata in the structure at hand. But this is the topic for another day.

References

[1] Bain, Jonathan (2006), "Spacetime Structuralism", in Dennis Dieks (ed.), *The Ontology of Spacetime*. Amsterdam: Elsevier, 37-66.

[2] Black, Max (1952), "The Identity of Indiscernibles", *Mind* 61: 153-164.

[3] Dorato, Mauro (2000), "Substantivalism, Relationism, and Structural Space-time Realism", *Foundations of Physics*, 30: 1605-1628.

[4] Esfeld, Michael (2004), "Quantum Entanglement and a Metaphysics of Relations", *Studies in History and Philosophy of Modern Physics* 35: 601-617.

[5] Esfeld, Michael, and Vincent Lam (2008), "Moderate Structural Realism about Space-time", *Synthese* 160: 27-46.

[6] French, Steven (2006), "Identity and Individuality in Quantum Theory", in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*, Spring 2006.

[7] French, Steven and James Ladyman (2003), "Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure", *Synthese* 136: 31-56.

[8] Keränen, Jukka (2001), "The Identity Problem for Realist Structuralism", *Philosophia Mathematica* 9: 308-330.

[9] Ladyman, James (2005), "Mathematical Structuralism and the Identity of Indiscernibles", *Analysis* 65: 218-221.

[10] Pooley, Oliver (2006), "Points, Particles, and Structural Realism", in Rickles et al. 2006, 83-120.

[11] Rickles, Dean, and Steven French (2006), "Quantum Gravity Meets Structuralism: Interweaving Relations in the Foundations of Physics", in Rickles et al. 2006, 1-39.

[12] Rickles, Dean, Steven French, and Juha Saatsi (eds.) (2006), *The Structural Foundation of Quantum Gravity*. Oxford: Clarendon.

[13] Stachel, John (2006), "Structure, Individuality, and Quantum Gravity", in Rickles et al. 2006, 53-82.

[14] Walker, A G (1944), "Completely symmetric spaces", *Journal of the London Mathematical Society* 19: 219-226.