Algebraic structuralism

December 29, 2016

1 Introduction

This essay is about how the notion of "structure" in ontic structuralism might be made precise. More specifically, my aim is to make precise the idea that the structure of the world is (somehow) given by the relations inhering in the world, in such a way that the relations are ontologically prior to their relata. The central claim is the following: one can do so by giving due attention to the relationships that hold *between* those relations, by making use of certain notions from algebraic logic. In the remainder of this introduction, I sketch two motivations for structuralism, and make some preliminary remarks about the relationship between structuralism and ontological dependence; in the next section, I outline my preferred way of unpacking the notion of "structure"; in §3, I compare this view to Dasgupta's *algebraic generalism*; and finally, I evaluate the view, by considering how well it can be defended against objections, and how well it discharges the motivations with which we began.

1.1 Motivations for structuralism

I don't intend to spend a lot of time discussing the motivations for structuralism: this essay is more about what the best way might be to understand structuralism, rather than about whether or not structuralism is a compelling position. Nevertheless, whether a given proposal is the best way to understand structuralism will clearly depend on the extent to which that proposal recovers motivations for structuralism and renders them compelling. So in this section, I outline two motivations for structuralism (and give thumbnail sketches of arguments for them).[1]

---

[1]There are other motivations, of course, such as metaphysical underdetermination ([Ladyman, 1998], [French, 2006], [French, 2011]), or certain features of particle physics ([Roberts, 2011],

First, consider the *argument from the history of science*.[2] The argument here is that although there is, as the pessimistic meta-induction notes, significant change in the content of our best scientific theories over time, this change does not (or does not so strongly) alter the *structural* features of those theories. Thus (to take a well-worn example), the idea is that in the transition from Fresnel's theory of light to Maxwell's, although there is indeed radical change in the ontological commitments of our theories (since the ontological commitment to the ether is dropped between the one theory and the next), there is nevertheless a structural continuity—as evidenced by the fact that Fresnel's equations governing light propagation were retained in the new theory. This is then taken to show that if scientific realists are to resist the pessimistic meta-induction, they ought to confine their belief in theories to just the structural aspects of those theories, since it is those aspects which are most likely to survive into future theoretical incarnations. Moreover, this observation is taken to show, or suggest, that it is the structural aspects of a theory to which a theory's empirical success ought properly to be attributed: as Worrall puts it,

> Roughly speaking, it seems right to say that Fresnel completely misidentified the *nature* of light; but, none the less, it is no miracle that his theory enjoyed the empirical predictive success it did; it is no miracle because Fresnel's theory, as science later saw, attributed to light the right *structure*.[3]

Second, consider the *argument from redundancy*.[4] Here, the concern is that an object-laden ontology exhibits a certain kind of redundancy, analogous to the redundancy of Newtonian spacetime with respect to Newtonian dynamics. For, consider any initial state of some system described in terms of particular individuals, and compare it with an initial state obtained by a mere permutation of those individuals (so that the distribution of qualitative properties and relations remains the same, with the only change being which individuals bear which properties and relations). The relevant observation is then that such a permutation commutes with the dynamics: the final states obtained by letting the original and permuted states evolve (under any dynamics with which we are familiar) will be related by exactly that same permutation. Therefore the permutation is, in a certain sense, "decoupled" from the dynamical evolution: differences at one time over which individuals have which properties and relations do not lead to any other

---

[McKenzie, 2014b], [McKenzie, 2014a]). I only outline two here, because they are the two which connect most directly with what I want to say; for a comprehensive overview of structuralism and the motivations for it, see [Frigg and Votsis, 2011].

[2][Worrall, 1989]

[3][Worrall, 1989, p. 157]

[4][Dasgupta, 2009], [Dasgupta, 2011]

differences at any other times. This provides a sense in which the facts about individuals do not seem to play any physically significant role, and hence seem to be redundant to our physical theorising. Moreover, insofar as we take ourselves and our processes of knowing to be physical processes, this observation seems to imply that those facts are not empirically accessible. For what it shows is that if two situations differ merely with respect to these facts (i.e., are related by a permutation of some kind), then there can be no physical process which converts that difference into any other kind of difference: a difference, for example, regarding whether I say "wow, the individuals here aren't the ones I dealt with previously!", or regarding my written response to any question about which individuals are here before me. But plausibly, a necessary condition on a certain kind of fact being empirically accessible to me is that it should (in principle) be possible to set up a reliable correlation between facts of that kind and facts about what I say, inscribe, etc. So what we see is that the physical redundancy of individuals negates the possibility of setting up such a correlation, even in principle—and hence, that the facts about individuals will not be empirically accessible.

One might be concerned that this argument does not, by itself, deliver a case for saying anything revolutionary about the ontological status of individuals. All it seems to require is denying the possibility of merely *haecceitistic* differences (differences only over which individuals play which qualitative roles);[5] so why not simply be an anti-haecceitist about objects, rather than anything more radical?[6] I'm sympathetic to this concern, but for the purposes of this essay I will bracket my sympathy and speak on behalf of the structuralist who is moved by the argument from redundancy. What such a structuralist should say, I think, is that it is not enough to just *assert* that objects lack individual haecceities; some account ought to be given of what objects are like, in a manner which illuminates or makes transparent just why they lack those haecceities.[7] The relevant analogy is that of something like electromagnetic potentials. It's one thing to assert that two potentials related by a gauge transformation are physically equivalent, i.e., to deny the possibility of merely gauge differences; it is quite another to provide an account of the invariant structure underlying those gauge-related potentials, in such a way as to make transparent just why gauge transformations do not change anything physical. In this case, electromagnetic fields provide such an account. Thus, the issue is whether the structuralist can find some account of what the world is like, which will stand to the usual account of objects as the account of electromagnetic fields stands to

---

[5]Here, I'm employing [Melia, 1999]'s definition of haecceitism as "the view that there are distinct possibilities that differ only over which objects play which roles."

[6]cf. [Pooley, 2006], [Greaves, 2011].

[7]cf. [McSweeney, ming], [Sider, 2016].

the account of electromagnetic potentials (and which will undergird anti-haecceitism, in the same way that fields undergird gauge-equivalence).

1.2 Dependence, elimination, and determination

Ontic structuralists claim that they can provide such an account, by holding that objects are ontologically dependent on structures. To unpack this thought, let's distinguish some attitudes one might have towards the relative ontological dependence of structure and objects. (At this stage, I'll leave undefined what is meant by "structure", "object", or "ontological priority": the aim is just to enumerate some slogans, the exact content of which is yet to be determined.) One could think that *structure is ontologically dependent on objects*, but not vice versa; one could think that *objects are ontologically dependent on structure*, but not vice versa; one could think that *neither objects nor structures are ontologically dependent on the other*; or one could think that *both objects and structure are ontologically dependent on one another*. Generally speaking, ontic structuralism is identified with either the second or fourth positions: that is, ontic structuralism is characterised by a belief that objects are ontologically dependent on structures. The fourth view, according to which structures are reciprocally dependent on objects, I'll refer to as *moderate ontic structuralism*;[8] the second view, according to which structures are not ontologically dependent on objects, I'll refer to as *radical ontic structuralism*.[9] Within radical ontic structuralism, there is a further subdivision. According to the *eliminativist*, objects do not—strictly speaking—exist at all. Rather, the structure is all there is; speaking of objects is a mere *façon de parler* for describing what the structure is up to. According to the non-eliminativist, objects do indeed exist—it's just that whenever they exist, they're ontologically dependent on structures.The structuralism articulated in this essay is a form of radical ontic structuralism. It can, I think, be adapted into either an eliminative or non-eliminative version; I now say a few words about what might constrain adapting it one way or the other.

The relevant constraint is the extent to which the structural facts determine the objectual facts. That is, what kinds of objectual facts are *fully fixed* by the structural facts, and which are left open? It seems implausible that any structuralist view will enable the recovery of *haecceitistic* structural facts, facts about which objects are which—indeed, given that the avoidance of haecceitism is a motivation for structuralism, it would be a problem if this were so. Formally, we can model this by supposing that on any structuralist position, the structural facts should be expected to determine the objectual

---

[8][Esfeld and Lam, 2010]
[9][French and Ladyman, 2003], [French and Ladyman, 2010]

facts only up to qualitative isomorphism: i.e., that given two objectual representations (Tarski-models) $\mathcal{M}$ and $\mathcal{N}$ which are qualitatively isomorphic (so that there is a bijection $f : |\mathcal{M}| \to |\mathcal{N}|$ such that for any qualitative predicate P, $f(P^{\mathcal{M}}) = P^{\mathcal{N}}$), $\mathcal{M}$ and $\mathcal{N}$ will correspond to the same structural representation. The idea behind this formulation is to leave open how to best model haecceitism: whether as holding that isomorphic models may represent distinct possibilities,[10] or as holding that one ought to include non-qualitative predicates (e.g. predicates representing haecceities) in one's models.[11]

However, this only commits us to thinking that qualitative isomorphism should be a sufficient condition for sameness of structural representation, not that it should be a necessary condition. So the further question remains: does a full specification of the structural facts determine the objectual facts up to qualitative isomorphism, or might a given configuration of structural facts be consistent with two objectual representations which are not qualitatively isomorphic to one another? I will make the following assumption: the radical non-eliminativist must assume that the structural facts suffice to fully fix the objectual facts (up to qualitative isomorphism); the radical eliminativist, however, may allow that two qualitatively non-isomorphic configurations of objects are consistent with the same structural facts.

The idea behind this assumption is that the latter, but not the former, can to at least some extent be an anti-realist about the objectual facts: they can think, for example, that there may just be no (genuine) fact of the matter about how many objects there are, or about precisely what those objects are up to. After all, if object-talk is just a way of characterising structure-configurations, then it may well turn out that two different descriptions in the "objectual mode" can characterise the same structure-configuration. In such a case, the eliminativist structuralist can shrug their shoulders, and just regard these descriptions as equivalent. That's not a problem, for although the two descriptions disagree about what exactly objects are up to, they agree on the structural facts; and only the latter are (genuine) facts. Such an attitude is analogous to (say) the attitude of most physicists towards electromagnetic potentials, or of ordinal utility theorists towards utilities. Two potential-distributions related by a gauge transformation, or two utility-functions related by an increasing monotonic transformation, are regarded as equivalent. That's not a problem, for although those descriptions disagree over what exactly the potentials or the utilities are up to, they agree on the facts regarding fields or preferences—and only the latter constitute (genuine) facts.

The non-eliminativist structuralist, by contrast, does not have this same freedom. For,

---

[10]cf. [Kaplan, 1975]

[11]I read [Weatherall, 2016] as advocating that haecceitism should be modelled in this latter fashion.

I take it, what it is to be committed to the existence of objects is to be committed to regarding (at least) the qualitative facts about objects as genuine, *bona fide* facts. So on the one hand, they have to get those facts from *somewhere*; and on the other, if they are really serious that all the fundamental facts are structural facts, then it seems the only place from which they could get them is from facts about the structure. That's a problem if it turns out that the same structure-configuration is compatible with multiple object-configurations: for then, it's not clear in what sense the objectual facts are being derived from the structural facts. In other words, the following principle seems plausible: if the x-facts fully depend upon the y-facts, then the y-facts should fix or determine the x-facts. (For instance, it seems like a necessary condition on taking the thermodynamical facts to depend upon the statistical-mechanical facts is that the statistical-mechanical facts fix the thermodynamical facts.) The moderate structuralist denies that the object-facts do fully depend upon the structural facts; the eliminative structuralist denies that the qualitative distribution of objects really ought to be (always) regarded as facts at all; but the radical, non-eliminative structuralist can do neither.

2 Structure as the web of relations

2.1 No relations without relata?

Thus, our challenge is to characterise what the "structural facts" are, which the "objectual facts" could be thought of as (perhaps superfluously) representing. The good news is that there is a fairly well-known answer to this which radical ontic structuralists have offered: that the structure of the world consists in the web or family of relations instantiated in the world. Thus, for example, [French and Ladyman, 2010] write that

> On either form of [radical ontic structuralism], we are basically saying the following: although the set-theoretic representation of structure, $\langle A, R \rangle$ (where A is a set of elements and R a family of relations) must be written down from left to right, from an ontological perspective it should be read semitically, that is, from right to left.[12]

Later, the identification is made even more explicit when they remark (in passing) that "as far as the structuralist is concerned, there is nothing but the structure—that is, the set of relations" [French and Ladyman, 2010, p. 14?]. Other commentators have also suggested this as a reading of radical ontic structuralism: for example, [Wolff, 2012]

---

[12][French and Ladyman, 2010, p. ???]

characterises the view as holding that "structure was to be understood as something like a web of relations."[13]

The bad news, however, is that when radical ontic structuralism is combined with this conception of structure, the view is regarded by many as incoherent: for, the slogan goes, one cannot have relations without relata. Thus consider, for instance, Chakravartty's attack:

> Consider the idea that SR [structural realism] tells us nothing about entities that partake in relations. This is false in at least the banal sense that, as noted earlier, the existence of relations allows us to ascribe, if nothing else, existence to the entities so related, whatever they may be exactly and however we may describe them. A famous dictum of Quine's holds that 'to be is to be quantified over' in some first order language. [...] This much must seem obvious—one cannot intelligibly subscribe to the *reality* of relations unless one is also committed to the fact that *some things* are related.[14]

My preferred reading of this criticism is that it faces the structuralist with a dilemma. Suppose that we do indeed seek to identify structure[15] with a set of relations $\{R_i\}$. The dilemma is this: are the elements of this set understood as primitive entities, or as being (say) sets of tuples? Somebody who uses the latter construal cannot, it seems, be a radical structuralist: for in their ontology, relations are not fundamental entities, but are instead derivative constructions. So, the radical structuralist must have in mind the first construal. But, goes the other horn of the dilemma, this just isn't any account of *relations* at all—if she is serious about taking relations as primitive, and about taking "structure" to be the set of relations, then all she is saying is that structure is constituted by a set of primitive "we-know-not-what"s; the only structural fact, on such a reading of "structure", would be facts about how many relations there are (since the only facts that a set of unstructured primitives supports are facts about its cardinality).

2.2 Boolean structuralism

I now propose a way out of this dilemma: specifically, a way in which the second horn can be neutralised, so that the eliminative structuralist can maintain both that relations are primitive entities in the ontology (i.e., are not structured entities that have been

---

[13][Wolff, 2012, p. 607]
[14][Chakravartty, 1998, p. 399]
[15]The structure of what? For the purposes of this discussion, either the world, or some model-theoretic representation thereof.

constructed out of objects), whilst having a suitably rich conception of structural facts. The trick is to reject the claim that the structure be identified with a *mere* set of relations, and instead identify the structure with a *structured* set of relations: specifically, with an *algebra* of relations. In effect, the idea is that the structuralist's ontology should include both the relations (as primitive entities), *and* the relationships that those relations bear to one another. As I'll show below, this gets around the objection above; it also has the advantage that the appropriate kinds of algebraic structures have been introduced and studied extensively by mathematicians and logicians, and so we have "off-the-shelf" ways of making the view precise.

Let's begin with a simple case. Suppose that we wished to capture the "structure" of a model $\mathcal{M}$ of predicate logic, for a language containing only monadic predicates $\{F, G, H, \dots\}$. $\mathcal{M}$ consists of a set $|\mathcal{A}|$, together with extensions $F^{\mathcal{M}} \subseteq |\mathcal{M}|$, $G^{\mathcal{M}} \subseteq |\mathcal{M}|$, etc. The relevant observation is that there are, in general, various important relationships between those extensions: for instance, perhaps (it turns out) the extension of F is disjoint from the extension of G; perhaps, even, the two extensions are the complement of one another. Or perhaps G lies entirely within H. Or perhaps the intersection of F and H strictly includes G. On the object-oriented viewpoint, such claims are understood in terms of the set-theoretic makeup of the extensions: that is, we interpret "intersection" as set-theoretic intersection, complement as set-theoretic complement, etc.

The proposal to hand, however, is to instead take them as self-standing claims in their own right, asserted to hold between the extensions construed as primitive entities (rather than as sets of objects).[16] In order to do so, we had better say *something* about the interpretation of terms such as "intersection" or "complement". Here's the trick: rather than interpreting them explicitly as certain kinds of set-theoretic constructions, we instead interpret them *implicitly* by providing an appropriate axiomatisation. In this case, an appropriate set of axioms is the axioms for *Boolean algebras*. Recall that a Boolean algebra consists of a set equipped with two binary operations $\cdot$ ("meet") and $+$ ("join"), one unary operation $-$ ("complement"), and two privileged elements 0 ("bottom") and 1 ("top"), obeying the following axioms: for any elements a, b, c of the algebra,

(B1) $a + (b + c) = (a + b) + c$

(B2) $a \cdot (b \cdot c) = (a \cdot b) \cdot c$

(B3) $a + b = b + a$

---

[16]Thus, this proposal—especially if taken in an eliminativist vein—may naturally be viewed as a form of the bundle theory.

(B4) $a \cdot b = b \cdot a$

(B5) $a + 0 = a$

(B6) $a \cdot 1 = a$

(B7) $a + (b \cdot c) = (a + b) \cdot (a + c)$

(B8) $a \cdot (b + c) = (a \cdot b) + (a \cdot c)$

(B9) $a + -a = 1$

(B10) $a \cdot -a = 0$

In what sense are these axioms appropriate? Recall that we are postulating these axioms in order to implicitly define the relationships amongst the extensions, as an alternative to defining those relationships explicitly in set-theoretic terms. So a minimal desideratum is that all of these axioms should be satisfied when those relationships are given their set-theoretic interpretations (i.e. interpreting meet as intersection, join as union, and complement as complement). It's not hard to verify that this is indeed the case, subject to one important proviso: in order to get a Boolean algebra out of our model $\mathcal{M}$, it is not enough to consider the *atomic* extensions $F^{\mathcal{M}}$, $G^{\mathcal{M}}$, etc.; we need to also include the *complex* extensions (e.g. the set of objects in $|\mathcal{M}|$ satisfying $(Fx \wedge Gx)$) in the algebra. Thus, one aspect of the appropriateness of the Boolean-algebra axioms is that they are (in a certain sense) "sound" with respect to the set-theoretic interpretation: for any model $\mathcal{M}$ in a monadic language, the definable extensions over $\mathcal{M}$ (together with the set-theoretic operations of intersection, union and complement) constitute a Boolean algebra. We will denote this algebra as $B(\mathcal{M})$.

However, there's another factor which makes the Boolean-algebra axioms appropriate. Not only can any set of subsets be interpreted as a Boolean algebra; any Boolean algebra can be interpreted as a set of subsets of some domain. This result is known as *Stone's theorem*: stated more carefully, it is the result that for any Boolean algebra B, there is some set X such that B is *isomorphic* to a Boolean algebra constructed out of subsets of X.[17] Thus, the axioms of Boolean algebras are, in a certain sense, "complete" with respect to the set-theoretic interpretation: any Boolean algebra B can be thought of as arising from, or as represented by, subsets of some domain of objects. That said, it will not always be the case that this representation will take the form of a model of (monadic) predicate logic, as can be seen by a simple cardinality argument: at least in a standard

---

[17]Isomorphic, in this context, meaning that there is a bijection f between the two algebras which preserves the algebraic operations (so $f(a \cdot b) = f(a) \cdot f(b)$, $f(-a) = -f(a)$, etc.).

first-order language, there are only countably many formulae, and hence a model of such a language will contain only countably many definable extensions; thus, no such model can represent a Boolean algebra containing uncountably many elements. In a slogan, we could say that although the axioms (B1)–(B10) are complete with respect to the *set*-theoretic interpretation of Boolean algebras, they are not complete with respect to the *model*-theoretic interpretation. But that seems like a reasonable form of completeness to demand; demanding completeness with respect to the model-theoretic interpretation would amount to fetishising finitary first-order logic.

If someone were to understand structure in this sense, they would have to deny that the structural facts suffice to fully determine the objectual facts. The reason for this is the following technical observation: in general, there are non-isomorphic models of monadic predicate logic with isomorphic Boolean algebras, as the following (trivial) example illustrates. Suppose that our signature has just one monadic predicate, K, and let $\mathcal{M}_1$ and $\mathcal{M}_2$ be models for that signature defined as follows: $\mathcal{M}_1$ has one element in its domain, $\mathcal{M}_2$ has two elements, and in both cases, $K^{\mathcal{M}_i} = |\mathcal{M}_i|$. Then in both cases, the Boolean algebra of definable extensions is (or rather, is isomorphic to) the trivial two-element algebra **2**—even though the two models are not isomorphic (since they disagree over the number of objects that exist). A less trivial illustration is as follows. For any model $\mathcal{M}$ of any language of monadic predicate logic, consider the following equivalence relation $\approx$ over $|\mathcal{M}|$: $a \approx b$ whenever there is an automorphism h of $\mathcal{M}$ such that $b = h(a)$. Let $\mathcal{M}/_{\approx}$ be the quotient of $\mathcal{M}$ by $\approx$, i.e., the model whose elements are $\approx$-equivalence classes of $\mathcal{M}$, and where for any predicate P, $[a] \in P^{\mathcal{M}/_{\approx}}$ iff $a \in P^{\mathcal{M}}$ (this is well-defined, since it is guaranteed that if $a \approx b$, then $a \in P^{\mathcal{M}}$ iff $b \in P^{\mathcal{M}}$). Then $\mathcal{M}$ and $\mathcal{M}/_{\approx}$ will give rise to the isomorphic Boolean algebras, and hence have the same Boolean structure.[18] These examples show that for the "Boolean structuralist", fixing the structure of the world does not suffice to determine the objectual character of the world. As discussed above, that suggests that the radical Boolean structuralist would have to be an *eliminativist* structuralist; if structural equivalence really is to suffice for *bona fide* equivalence, then the Boolean structuralist will have to be an anti-realist about objects.

How much of an anti-realist do they have to be? That is, how much disagreement over objects must the Boolean structuralist be willing to brook? To answer this question, first recall the notion of a *definitional expansion* of a model $\mathcal{M}$ of (not necessarily monadic) predicate logic. If the vocabulary of $\mathcal{M}$ is Σ, then a definitional expansion of $\mathcal{M}$ is a

---

[18]This corresponds to the well-known observation that the bundle theory is (at least prima facie) committed to the identity of indiscernibles.

model whose vocabulary is Σ^+, where $\Sigma \subset \Sigma^+$, such that:

- $|\mathcal{M}^+| = |\mathcal{M}|$

- $\mathcal{M}^+ \upharpoonright_\Sigma = \mathcal{M}$ (where $\mathcal{M}^+ \upharpoonright_\Sigma$ is the *reduct* of $\mathcal{M}^+$ to Σ)

- For every $\Pi \in \Sigma^+ \setminus \Sigma$, $\Pi^{\mathcal{M}^+}$ is definable in $\mathcal{M}$

We then say that two models $\mathcal{M}$ and $\mathcal{N}$ are *definitionally equivalent* if they have a common definitional expansion: that is, if there are definitional expansions $\mathcal{M}^+$ and $\mathcal{N}^+$ such that $\mathcal{M}^+ = \mathcal{N}^+$. Second, recall that two models are said to be *elementarily equivalent* if they satisfy exactly the same sentences of their associated languages. We then have the following result:

Proposition 1. Suppose that $\mathcal{M}$ and $\mathcal{N}$ are two models of models of two monadic first-order languages (without identity).[19] Then $B(\mathcal{M}) \cong B(\mathcal{N})$ iff $\mathcal{M}$ is elementarily equivalent to some model definitionally equivalent to $\mathcal{N}$ (that is, iff there are definitional expansions $\mathcal{M}^+$ and $\mathcal{N}^+$ such that $\mathcal{M}^+$ is elementarily equivalent to $\mathcal{N}^+$).

Proof. See [Monk, 2000, Theorem 3.3]. □

2.3 Cylindric structuralism

Of course, the limitation to monadic extensions is somewhat antithetical to the structuralist credo. But the algebraisation of a collection of subsets by Boolean algebra is only the simplest example of this kind of construction. The extension to binary relations yields the notion of a *relation algebra*: this encodes not only the Boolean operations, but also the operations of conversion and composition. The theory of relation algebras is sufficient to express any first-order formula containing no more than three variables—and hence, to express virtually all the mathematics one is liable to use—but if one did want to go further, then two natural options are the theory of *polyadic algebras*,[20] or that of *cylindric algebras*[21]. A polyadic algebra has, in addition to the standard Boolean operations, an operation corresponding to existential quantification; a cylindric algebra has the standard Boolean operations, an operation corresponding to existential quantification, and privileged elements corresponding to the identity.[22] Thus, which kind

---

[19] I stress that the language is without identity in order to stress that the elementary equivalence here is with respect to that language: i.e., that $\mathcal{M}^+$ and $\mathcal{N}^+$ will *not* (in general) satisfy the same identity-sentences.

[20] [Halmos, 1954], [Halmos, 1962]; for interesting reflections on the development of polyadic algebras, see [Halmos, 2000].

[21] See [Monk, 2000].

[22] As this suggests, one can think of cylindric algebras as polyadic algebras enriched with an operation of identity: more precisely, one can show that the "locally finite" polyadic algebras with identity are equivalent to the "locally finite" cylindric algebras (see [Galler, 1957] or [Plotkin, 2000, §1.2]).

of algebra is appropriate for capturing the structure will depend on the structuralist's attitude towards identity. For structuralists willing to admit identity as a bona fide autonomous relation,[23] then cylindric algebras would seem to be the more appropriate choice; for those unwilling to do so, polyadic algebras recommend themselves. For the sake of concreteness and generality, I will discuss cylindric algebras (but extending my remarks to polyadic algebras should be straightforward).

So, let us start by describing the cylindric-algebraic structure of a model $\mathcal{M}$ of first-order (polyadic) predicate logic. We will suppose that the variables of the associated first-order language come with some specific ordering, $x_1, x_2, \ldots$. Now, consider any infinite sequence $\mathbf{a} = \langle a_1, a_2, \ldots, \rangle$ of elements of $|\mathcal{M}|$, and any formula $\phi(\xi_1, \ldots, \xi_n)$ of the language of $\mathcal{M}$. We will say that $\mathbf{a}$ *satisfies ϕ in $\mathcal{M}$* just in case $\mathcal{M}$ satisfies ϕ relative to some (hence, any) variable-assignment sending each ξ_i to a_i. We will denote this as $\mathcal{M} \models \phi[\mathbf{a}]$. Let the *extension of ϕ* be $\phi^{\mathcal{M}} = \{\mathbf{a} : \mathcal{M} \models \phi[\mathbf{a}]\}$. Thus, we can form the set of all $\mathcal{M}$-definable extensions, $\mathbf{M} = \{\phi^{\mathcal{M}} : \phi \text{ is a formula of the language of } \mathcal{M}\}$. We will take this set as the domain of the algebra.[24] Since each element of $\mathbf{M}$ is a subset of the set of all infinite sequences, we can equip $\mathbf{M}$ with Boolean structure in the same fashion as the monadic case; and since the language is closed under the propositional operations, $\mathbf{M}$ is closed under the Boolean operations. Then, as mentioned above, we also equip $\mathbf{M}$ with two further pieces of structure. First, for any $i, j \in \mathbb{N}$, we define the (i, j)-*th diagonal set* D_{ij} to be the extension $(x_i = x_j)^{\mathcal{M}}$. Second, for any $i \in \mathbb{N}$, we define i-*cylindrification* to be the mapping $C_i : \mathbf{M} \to \mathbf{M}$ given by $C_i(\phi^{\mathcal{M}}) = (\exists x_i \phi)^{\mathcal{M}}$.

As in the Boolean case, the *cylindric structuralist* proposes that the elements of this algebra—the relations—not be given a set-theoretic interpretation (i.e., that they not be interpreted as sets of sequences) but rather taken as primitive, fundamental entities. As in the Boolean case, the legitimacy of this proposal depends on the possibility of replacing the explicit, set-theoretic interpretation of the algebraic structures by an implicit, axiomatic interpretation. Fortunately, we do indeed have such an interpretation. An (abstract) *cylindric algebra* consists of a set $\mathbf{A}$, equipped with

- Boolean operations $\cdot, +, -$, and privileged elements $1, 0$;

- For each $i, j \in \mathbb{N}$, a privileged element d_{ij};

- For every $i \in \mathbb{N}$, a unary operation $c_i : \mathbf{A} \to \mathbf{A}$

---

[23]cf. [Leitgeb and Ladyman, 2008]

[24]Thus, the elements of the algebra are sets of infinite sequences, not—as one might have thought—sets of n-tuples. This just avoids certain technical wrinkles that we would have to deal with if working with sequences of different lengths: see the discussion in [Németi, 1991, §4].

obeying the Boolean axioms (B1)–(B10), as well as the following axioms (for any $r, s, t \in$ **A**, and any distinct $i, j, k \in \mathbb{N}$):

(C1) $c_i 0 = 0$

(C2) $r \le c_i r$

(C3) $c_i(r \cdot c_i s) = (c_i r) \cdot (c_i s)$

(C4) $c_i c_j r = c_j c_i r$

(C5) $d_{ii} = 1$

(C6) $d_{ik} = c_j(d_{ij} \cdot d_{jk})$

(C7) $c_i(d_{ij} \cdot r) \cdot c_i(d_{ij} \cdot (-r)) = 0$

In what sense are these axioms appropriate? Well, first, we have soundness (with respect to the set-theoretic interpretation) in the same sense as for Boolean algebras: for any model $\mathcal{M}$, the axioms (B1)–(B10) and (C1)–(C7) are satisfied if we take **M** as the underlying set, the set-theoretic operations as the Boolean operations, the diagonal sets D_{ij} as the d_{ij} and the cylindrification operations C_i as the c_i. We will denote this cylindric algebra as $C(\mathcal{M})$. More generally, given any set U and ordinal α, we can construct cylindric algebras whose elements are sets of α-length sequences of elements of U (i.e., all functions $\alpha \to U$). Such an algebra is referred to as a *cylindric set algebra* (of dimension α).[25]

The issue of completeness is slightly more delicate. There is no straightforward analogue of the Stone theorem: in general, there are cylindric algebras (i.e. structures satisfying the axioms (B1)–(B10) and (C1)–(C7)) which are not isomorphic to any cylindric set algebra. Indeed, it is known that the class of those cylindric algebras which are isomorphic to some cylindric set algebra—the so-called *representable* cylindric algebras[26]—cannot be finitely axiomatised.[27] However, it can be recursively axiomatised; moreover, there exist various other (non-trivial) ways of picking it out.[28] So we do have some means of characterising the representable cylindric algebras, and one could regard these means as providing the implicit definition of the algebraic operations that

---

[25]The cylindric set algebras of the form $C(\mathcal{M})$, for some model $\mathcal{M}$ of predicate logic, are exactly the ω-dimensional cylindric set algebras which are "regular" and "locally finite": see [Monk, 2000, §11].

[26]Strictly speaking, the representable cylindric algebras are those which are isomorphic to some subdirect product of cylindric set algebras. But one can show that any such algebra is directly representable as a cylindric set algebra (see [Németi, 1991, p. 503]).

[27][Monk, 1969]; see [Németi, 1991, Theorem 4].

[28]See [Henkin et al., 1986] for a survey.

we sought. Alternatively, if one thought that the implicit definition ought to be appropriately "surveyable" in some sense (and denied that recursively axiomatisability was sufficient for surveyability), then one could take (C1)–(C7) to provide at least *some* grasp of the meanings of the algebraic operations with which we have spelled out our notion of structure—albeit not one that permits us to infer that the structural facts will always be representable by some distribution of objects equipped with properties and relations. (A really radical structuralist might deny that this is even a problem: perhaps, they think, there are structural possibilities which correspond to no objectual possibility.)

Again, the (cylindric) structural facts underdetermine the objectual facts, in the following sense: there are non-isomorphic models $\mathcal{M}$ and $\mathcal{N}$ such that $C(\mathcal{M}) \cong C(\mathcal{N})$. In fact, we have the following analogue of Proposition 1.

Proposition 2. For any models $\mathcal{M}$ and $\mathcal{N}$ of first-order predicate logic with identity, $C(\mathcal{M}) \cong C(\mathcal{N})$ iff $\mathcal{M}$ is elementarily equivalent to a picture definitionally equivalent to $\mathcal{N}$ (that is, iff there are definitional expansions $\mathcal{M}^+$ and $\mathcal{N}^+$ such that $\mathcal{M}^+$ is elementarily equivalent to $\mathcal{N}^+$).

Proof. See [Monk, 2000, Theorem 3.3]. □

Note that the elementary equivalence here is with respect to a language containing identity, unlike the elementary equivalence referred to in Proposition 1. Hence, the cylindric structuralist does not face as much underdetermination as the Boolean structuralist: indeed, if $\mathcal{M}$ and $\mathcal{N}$ are finite, then $C(\mathcal{M}) \cong C(\mathcal{N})$ iff $\mathcal{M}$ is definitionally equivalent to $\mathcal{N}$ (since for finite structures, elementary equivalence implies isomorphism).

3 Algebraic generalism

I'm not the first to defend the idea that ideas and resources from algebraic logic could be used to articulate and defend some form of structuralism. [Dasgupta, 2009] uses (a modified form of) Quine's predicate-functor logic to express a structuralist view, which he refers to as "generalism". For our purposes, the relevant observation is that Dasgupta spells out his "generalist language" G by using *algebraic* resources, as follows:

> The generalist's ontology consists of a domain of properties, each with a given adicity. Properties of adicity greater than one might be better called relations, but I will use the term 'property' to cover them all. To talk about these properties the generalist introduces to language G terms of the form $\mathbf{P}^n$, where n indicates the adicity of the property denoted. [...] According

14

> to the generalist, the domain of properties instantiates a certain algebraic structure. To talk about this structure, she introduces six expressions to G: $\&$, $\sim$, $\mathbf{c}$, $\mathbf{p}$, ι, and σ. Syntactically, they apply to terms to give a complex term—much like 'the father of x' in English applies to terms to give a complex term—so we can call them 'term-functors'. [. . .]
>
> In addition to these term-functors, the generalist also helps herself to a 2-place term $\mathbf{I}^2$ which denotes what we ordinarily think of as the relation of identity.[29]

These term-functors are drawn from Quine's "predicate-functor" logic. The idea is that they may be applied to terms for properties in order to form other, complex terms for properties. Using λ-calculus notation, and letting $\mathbf{\Phi}^n = \lambda x_1 \ldots x_n.\phi(x_1, \ldots, x_n)$ and $\mathbf{\Psi}^m = \lambda x_1 \ldots x_m.\psi(x_1, \ldots, x_m)$, the relationship between the expressions of G and those of ordinary first-order predicate logic are as follows:

- $\sim \mathbf{\Phi}^n = \lambda x_1 \ldots x_n.\neg\phi(x_1 \ldots x_n)$

- $(\mathbf{\Phi}^n \& \mathbf{\Psi}^m)^k = \lambda x_1 \ldots x_k.(\phi(x_1, \ldots, x_n) \wedge \psi(x_1, \ldots, x_m))$, where $k = \max(n, m)$

- $\sigma \mathbf{\Phi}^n = \lambda x_1 \ldots x_n.\phi(x_n, x_1, x_2, \ldots, x_{n-1})$

- $\iota \mathbf{\Phi}^n = \lambda x_1 \ldots x_n.\phi(x_2, x_1, x_3, \ldots, x_n)$

- $(\mathbf{c}\mathbf{\Phi}^n)^{n-1} = \lambda x_1 \ldots x_n.\exists x_1 \phi(x_1, \ldots, x_n)$

- $(\mathbf{p}\mathbf{\Phi}^n)^{n+1} = \lambda x_1 \ldots x_{n+1}.\phi(x_1, \ldots, x_n)$

G is then equipped with a single predicate for 0-place terms, **obtains**: asserting that Φ^0 **obtains** is what the generalist does in lieu of asserting that ϕ.

The resemblance to the cylindric-algebra apparatus is hopefully clear: clearly, the operations being used here are designed to do the same sort of thing as the cylindric-algebraic operations. I don't want to go into the details of the relationship between the two; the more important thing is that the only explanation Dasgupta gives of the language G is in terms of ordinary predicate logic, PL. In particular, he gives a semantics for G which is based on ordinary, predicate-logic models—essentially, the semantics obtained by "pulling back" the standard Tarskian semantics for predicate logic by the correspondence given above. (To be more accurate, he gives a semantics for a closely related language Q, in terms of which he defines a consequence relation on G. I'll gloss

---

[29][Dasgupta, 2009, pp. 52–53]

over this, since I don't believe anything significant hangs on it.) This then leads to a concern that what we have here is not a new metaphysics, but just a curious notation for expressing the old, object-oriented metaphysics. Dasgupta's main result, that one can translate at will between G and PL, is hardly likely to reassure us in this regard: that just adds to the suspicion that G is no more generalist (structuralist) than PL.

Dasgupta does address this concern, however: he writes that

> Earlier [. . .] I explained what the term-functors of G mean by appealing to our ordinary understanding of properties being instantiated by a *domain of individuals*. But the generalist claims that individuals are not constituents of the fundamental facts of the world. So was that method of explanation legitimate?
>
> Similarly, I just used the semantics for Q to define a relation of implication on G and a relation of equivalence between G and PL. But the semantics for Q used ordinary models of PL of the form $\mathcal{M} = (D, v)$, where D is a *domain of individuals*. Was that legitimate?
>
> I think it was. The key is to distinguish between metaphysical and conceptual priority. Metaphysically, the generalist claims that the fundamental facts of the world are those expressed by G. But conceptually, I have explained the view to you in terms that you are familiar with, namely in terms of individuals. Once you are competent with G you can then, if you so wish, become a generalist by giving up your understanding of G in terms of a domain of individuals and simply taking the term-functors and predicate of G as undefined primitives.
>
> There is no space to defend the cogency of the strategy here in full. I shall only plead innocence by association, for there are many other cases in which we adopt it. The case of velocity is [. . .] a good example. Our ordinary concept of the relative velocity between two material things is defined as the difference between their absolute velocities. Once you are competent with the concept, you can then (if you so wish) dispense with absolute velocity and either think of relative velocity as a primitive concept or define it in other terms. And this is precisely what we do when we adopt a Galilean conception of space-time.[30]

My response to this will probably be unsurprising: I don't think that this is a legitimate strategy. The mere adoption of a variable-free notation like G is *surely* not enough to

---

[30][Dasgupta, 2009, p. 66]

renounce commitment to objects—not if the only explanation that can be offered of the meaning of that notation is exactly the same as that offered to explain the meaning of its variable-laden cousin PL. Dasgupta's proposal seems to be that having used this semantics to induce a consequence-relation on G, we kick away the ladder by retaining the consequence-relation but forgetting the semantics from which it arose. On the one hand, it's then not clear what the purpose of introducing G is after all: why not perform the same manoeuvre with respect to PL, and save the time required to learn a new notation?[31] On the other, it's not clear that this retains any meaning for the language G at all. Absent a semantics, what makes G a *language* equipped with a relation of *implication* at all, rather than (merely) a set of strings equipped with a certain partial order?

What is needed is a semantics for G which explains how the world could make sentences of G true, even if the world is (as the radical eliminative structuralist claims) not fundamentally constituted by individuals standing in properties and relations, and therefore not apt to be represented by a Tarski-style model. I agree with Dasgupta that the analogy with Galilean spacetime is instructive here, although I don't think its implications are what he suggests. It's one thing to abjure talk of absolute velocities, and to promise that one will only use sentences that speak in terms of velocity-differences: i.e. to promise that terms such as v_1 or v_2 will not be used to assert anything about the world, unless they have been made safe by embedding into terms such as $(v_2 - v_1)$. Nevertheless, if the only account that one can give of the *meaning* of a term such as $(v_2 - v_1)$ is one which defers to the meanings of v_1 and v_2 (as being absolute velocities in something like Newtonian spacetime), then one has not fully succeeded in articulating a metaphysics which is free of absolute velocities.[32] The power of Galilean spacetime is precisely that it provides a structure in which a term like $(v_2 - v_1)$ is interpreted directly, not via an interpretation of its component parts. Cylindric algebras, I claim, provide just such a semantics for G—or, equivalently, for PL.

---

[31]Here's a possible answer: elsewhere, Dasgupta writes that "it is arguably analytic of the existential quantifier that existentially quantified facts are grounded in their instances." [Dasgupta, 2011, p. 131] The problem is that what makes this an analytic feature of the existential quantifier is surely the form of the semantics for it—as Dasgupta then goes on to say, "this understanding of the quantifier is arguably implicit in the standard Tarskian semantics for PL." [Dasgupta, 2011, pp. 131–132]. If so, then it will equally well be analytic of the operator **c** that facts of the form $(\mathbf{c}\mathbf{\Phi}^1)^0$ **obtains** are grounded in instantiations of $\mathbf{\Phi}^1$.

[32]Indeed, this is arguably the bind in which Newton's opponents found themselves: determined to renounce talk of absolute velocities, but unable to provide an appropriate structure for interpreting that talk (without also having to renounce talk of absolute accelerations).

4 Evaluation

In this final section, I evaluate the proposed form of structuralism: first, by considering two objections to it, and then by looking at how well it satisfies the motivations for structuralism that I introduced earlier.

4.1 Objections and replies

The first objection is that this is all (mere) metaphysics, of no relevance to any actual science. Our best scientific theories are most certainly *not* phrased in terms of first-order predicate logic: so why should we worry about such esoterica as whether representable algebras can be finitely axiomatised, the gap between elementary equivalence and isomorphism, etc? My first response to this is to plead *così fan tutte*: confining attention to the case of first-order theories and models is, for better or worse, a standard feature of the literature on structuralism (and, indeed, the metaphysics of science more generally). So insofar as my aim here is just to give a version of structuralism that is superior to its rivals on the market, the fact that I also engage in this practice is not dialectically relevant. However, I also think that a more substantial defence can be given.

For, there are good grounds for thinking that in actual, realistic theories, we can indeed find algebraic representations of the models of those theories which are sympathetic to the structuralist's aims. For instance, in spacetime theories there has been substantial interest in representing a manifold via its *smooth algebra*: the commutative, associative, unital algebra (over $\mathbb{R}$) of smooth scalar fields on the manifold. Such an algebra encodes a great deal of information about the structure of the manifold; indeed, one can show that two manifolds have isomorphic smooth algebras iff they are diffeomorphic to one another. Moreover, geometric object fields on the manifold (e.g. vector fields, a metric, etc.) can be represented by certain structures on the smooth algebra (e.g. a vector field corresponds to a derivation on the smooth algebra).[33] Thus, a model of a spacetime theory $\langle M, O_i \rangle$ (with M a manifold and the O_i being geometric object fields) can be represented by an appropriate spacetime algebra $\langle \overline{M}, \overline{O}_i \rangle$, where $\overline{M}$ is the smooth algebra and the $\overline{O}_i$ are the algebraic representations of the O_i: for example, if $\langle M, g \rangle$ is a Lorentzian manifold, then $\langle \overline{M}, \overline{g} \rangle$ is known as an *Einstein algebra*. Similarly, in analysing quantum theories it often proves insightful to consider a collection of operators on a given Hilbert space to constitute a certain kind of algebra, known as a C^*-*algebra*; in classical mechanics, the smooth scalar functions on phase space may usefully

---

[33]Indeed, in modern differential geometry one standardly *defines* a vector field as a derivation on the smooth algebra (e.g. [Baez and Muniain, 1994, Chapter 3]).

be characterised as a *Poisson algebra*; and so on and so forth. For our purposes, the relevant observation is that in such cases, we form an algebra whose elements represent qualitative properties (smooth scalar fields in the case of smooth algebras, physical observables in the case of C^*- or Poisson algebras), and so abstract away from the individuals which would, in a more traditional presentation, be taken to be the bearers of those properties. Thus, the algebraic approach to structuralism looks as though it will generalise to more realistic scientific contexts. Moreover, by studying the (relatively!) simple case of algebras for first-order models, we can get insight into what questions to ask, and what kinds of answers it is reasonable to expect. For example, the standard definition of a smooth algebra is (roughly speaking) that it is a commutative, associative, unital algebra (over $\mathbb{R}$) which is isomorphic to the smooth algebra of scalar functions over some manifold; given the discussion of representability for cylindric algebras, we are motivated to ask whether any independent, complete characterisation of the class of smooth algebras is possible. (So far as I know, this remains an open question.) Similarly, the fact that isomorphism at the level of algebras does not necessarily lead to isomorphisms at the level of objectual representations suggests that we should look for analogous "gaps" between the algebraic and non-algebraic representations in other theories: which, in the context of C^*-algebras, is precisely the question of what to say about unitarily inequivalent Hilbert-space representations of the same C^*-algebra.[34]

The second objection is that algebraic structuralism does not succeed in truly expunging objects from the ontology. After all, a cylindric algebra consists of a bunch of things—the elements of the algebra—standing in certain algebraic relations to one another. So (the objection goes) rather than eliminating objects from the ontology, the algebraic structuralist has merely *replaced* them: they have shown how one kind of object ("ordinary" objects) can be eliminated in favour of another (elements of a cylindric algebra), but not how one might do without objects altogether. Lam and Wüthrich raise this concern in the context of Einstein algebras:

> [. . .] let us emphasize that the mere non-trivial elimination of manifold points (or their correlates) does not suffice to give [radical ontic structuralism] the upper hand against [moderate ontic structuralism]. What must be shown to establish this, instead, is that any relata whatsoever are non-trivially eliminated from the formulation of the theory; [moderate ontic structuralism] does not include or entail a commitment to any specific set of relata such as manifold points.[35]

---

[34][Ruetsche, 2011]
[35][Lam and Wüthrich, 2015, p. 624]

Similarly, Psillos remarks that "Even in group theory—Ladyman and French's paradigm case for their thesis—the group-structure is detached from any particular domain of individuals, but not from the very notion of a domain of individuals."[36]

This objection, however, rests upon a conflation of representational and metaphysical criteria. The objection concedes that the algebraic account does indeed present relations as atomic, unanalysed components of the metaphysics, lacking any internal structure; it then complains that, *in virtue of that very fact*, we are no longer talking about relations—because (says the objection) to take something as an atomic metaphysical component, devoid of internal structure, *just is* to treat it as an object rather than a relation. But this conflates *being a subject of predication* with *being an object*.[37] The fact that our representational apparatus is that in which some characteristics are ascribed of some entities does not mean that the entities must be identified as objects. Indeed, even if we had some representational apparatus that (somehow) avoided taking the form of such ascriptions,[38] merely recasting the standard object-oriented metaphysics into that apparatus would not thereby make us into structuralists; it would do so only if, by the lights of that apparatus, relations were treated as not being constituted by objects. By the same token, a view on which relations are not so constituted is a structuralist view— independently of whether that view is presented using a (broadly) subject-predicate apparatus like that employed here, or some other apparatus.

4.2 Meeting desiderata

Let us now turn to whether algebraic structuralism is an apt vehicle for the motivations considered in section 1.1. First, the argument from the history of science claimed that the "structure" of a theory will be preserved across theory change, at least to a greater extent than other aspects of a theory. Evaluating this claim is not straightforward, for two reasons: first, that (as already discussed) the only case we've considered in detail is the rather unrealistic case of first-order theories; second, that in this essay I've only talked about using algebraic means to characterise the structure of *models* of a theory, rather than the structure of a theory as a whole.[39] Nevertheless, even within these

---

[36][Psillos, 2001, p. S22]

[37]Of course, there is a long-standing philosophical tradition of identifying *substance* with the ultimate subject of predication. But the question of whether it is structures or objects which constitute the fundamental substance of the world is exactly what is at issue.

[38]It remains an open question, to my mind, whether there is or could be such an apparatus: the best hope is likely something category-theoretic, but the details remain to be made out.

[39]One might have thought that characterising the structure of a theory would just be a matter of characterising the structure of each of its models. However, that risks neglecting those aspects of a theory's structure concerning the relationships *between* models, which may be important for determining the

constraints, I think we can see some ways in which algebraic structuralism may be able to support such a claim. Since the characteristic feature of the relevant algebras is that they comprise physical properties and relations, and it is those properties and relations which most immediately connect to empirical results, then—given the assumption that the empirical success of the old theory will be incorporated into its successor—we have some grounds for thinking that the algebraic structure of the two theories should exhibit the kind of continuity the argument demands. Of course, we don't expect that the algebraic structure will be *fully* preserved, i.e., that the algebraic representations offered by the successor theory will be precisely isomorphic to those offered by its predecessor: if that were so, then it's hard to see how the two theories could (for the structuralist, at least) really be regarded as distinct. But we can expect that the structure of the predecessor theory will be somehow embedded in that of the successor theory, i.e., that the algebras associated to the latter will, in some appropriate limit, reduce to or approximate towards the algebras associated to the former. And if we attend to the examples mentioned above, then we find that this picture is (somewhat) borne out: for example, in the use of deformation quantisation to analyse the relationship between algebras of classical and quantum observables.

Second, consider the argument from redundancy. The argument here was that we need an account of what individuals are which will deliver anti-haecceitism. I claim that algebraic structuralism does, indeed, succeed in doing that. Two objectual representations which are related by a permutation are, for the algebraic structuralist, merely two ways of representing the same structural facts (by Proposition 2, given that isomorphism entails elementary equivalence). We therefore have a characterisation of the invariant reality underlying such representations, which helps to explain why anti-haecceitism holds. One concern (related to the second objection above) is that although algebraic structuralism avoids the permutation argument for individuals, it is subject to a higher-order version of the same problem: by permuting the relations in an algebra, one can obtain an algebra that differs only over which relations play which algebraic role. That appears to mean that we will face the same problems of redundancy and inaccessibility; and hence, one might question what has really been gained.[40] I don't have space for a full response here, but let me make a couple of remarks. First, the permissibility of such permutations will depend upon the extent to which one is a *nomic essentialist*: i.e., on whether one thinks that it is possible for a property or relation to play a different role

---

theory's content (see [Halvorson, 2012, §7] for an elaboration of this claim).

[40]This is, essentially, the point that [Rynasiewicz, 1992] makes in the context of spacetime algebras.

in the laws to that which it actually plays.[41] If nomic essentialism is true, and if nomological roles are unique, then no such permutation will deliver a possible world—so the arguments for redundancy and inaccessibility cannot get off the ground. Second, even if one does not find nomic essentialism plausible with respect to ordinary properties and relations, it seems to me that it is highly compelling with respect to the higher-order properties and relations constituting the algebraic relationships (e.g. $+$, c_i, etc.): given that (on the structuralist's account) the nature of these relations is constituted purely by their implicit definitions via the axioms (B1)–(B10) and (C1)–(C7), it is very hard to make sense of the idea of (say) the operation c_i coming apart from the "c_i-role" in the axioms. If this is granted, and given that permutations of (first-order) properties and relations are just present in the non-structuralist picture as the structuralist picture, then the algebraic structuralist is at least dealing with one fewer permutation problem than the non-structuralist: zero rather than one (if nomic essentialism about first-order properties is accepted) or one rather than two (if nomic essentialism about first-order properties is rejected).

5 Conclusion

So, to conclude: I've argued that the most charitable way to unpack the notion of "structure as the web of relations" is by characterising structure as an algebra of relations—specifically, at least in the context of first-order theories, as a cylindric algebra. There remain many open questions. Some of these are of a technical nature: what sort of algebra is appropriate in a context where we permit a stronger background logic, e.g. infinitary or higher-order logic? How would such algebras bear on the "representation" theorems, like Proposition 1 and 2? And what does the analysis look like in the context of Einstein algebras or C^*-algebras? Others are more philosophical: how exactly do technical results like Propositions 1 and 2 bear on theses about ground or ontological dependence? Does this apparatus apply to epistemic forms of structuralism, or only to ontic structuralism? Are the motivations for structuralism really compelling, after all? Nevertheless, I hope that this work has laid the foundations for good progress on these questions; and that I have convinced you that whether or not "relations without relata" is the right view to hold about what the world is like, it is at least a coherent and intelligible view.

---

[41]See e.g. [Hawthorne, 2001], [Schaffer, 2005].

References

[Baez and Muniain, 1994] Baez, J. C. and Muniain, J. P. (1994). *Gauge Fields, Knots, and Gravity*. World Scientific, Singapore.

[Chakravartty, 1998] Chakravartty, A. (1998). Semirealism. *Studies in History and Philosophy of Science Part A*, 29(3):391–408.

[Dasgupta, 2009] Dasgupta, S. (2009). Individuals: An essay in revisionary metaphysics. *Philosophical Studies*, 145(1):35–67.

[Dasgupta, 2011] Dasgupta, S. (2011). The Bare Necessities. *Philosophical Perspectives*, 25(1):115–160.

[Esfeld and Lam, 2010] Esfeld, M. and Lam, V. (2010). Holism and structural realism. In Vanderbeeken, R. and D'Hooghe, B., editors, *Worldviews, Science and Us. Studies of Analytical Metaphysics. A Selection of Topics from a Methodological Perspective*, pages 10–31. World Scientific, Singapore.

[French, 2006] French, S. (2006). Structure as a Weapon of the Realist. *Proceedings of the Aristotelian Society*, 106(1):169–187.

[French, 2011] French, S. (2011). Metaphysical underdetermination: Why worry? *Synthese*, 180(2):205–221.

[French and Ladyman, 2003] French, S. and Ladyman, J. (2003). Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure. *Synthese*, 136(1):31–56.

[French and Ladyman, 2010] French, S. and Ladyman, J. (2010). In Defence of Ontic Structural Realism. In Bokulich, A. and Bokulich, P., editors, *Scientific Structuralism*, pages 25–42. Springer Netherlands, Dordrecht.

[Frigg and Votsis, 2011] Frigg, R. and Votsis, I. (2011). Everything you always wanted to know about structural realism but were afraid to ask. *European Journal for Philosophy of Science*, 1(2):227–276.

[Galler, 1957] Galler, B. A. (1957). Cylindric and polyadic algebras. *Proceedings of the American mathematical Society*, 8(1):176–183.

[Greaves, 2011] Greaves, H. (2011). In Search of (Spacetime) Structuralism. *Philosophical Perspectives*, 25(1):189–204.

[Halmos, 1954] Halmos, P. R. (1954). Polyadic Boolean Algebras. *Proceedings of the National Academy of Sciences of the United States of America*, 40(5):296–301.

[Halmos, 1962] Halmos, P. R. (1962). *Algebraic Logic*. Chelsea Publishing Company. Google-Books-ID: BzNXYl6UZzUC.

[Halmos, 2000] Halmos, P. R. (2000). An autobiography of polyadic algebras. *Logic Journal of IGPL*, 8(4):383–392.

[Halvorson, 2012] Halvorson, H. (2012). What Scientific Theories Could Not Be. *Philosophy of Science*, 79(2):183–206.

[Hawthorne, 2001] Hawthorne, J. (2001). Causal Structuralism. *Noûs*, 35:361–378.

[Henkin et al., 1986] Henkin, L., Monk, J. D., and Tarski, A. (1986). Representable cylindric algebras. *Annals of Pure and Applied Logic*, 31:23–60.

[Kaplan, 1975] Kaplan, D. (1975). How to Russell a Frege-Church. *The Journal of Philosophy*, 72(19):716–729.

[Ladyman, 1998] Ladyman, J. (1998). What is Structural Realism? *Studies in History and Philosophy of Science Part A*, 29(3):409–424.

[Lam and Wüthrich, 2015] Lam, V. and Wüthrich, C. (2015). No Categorial Support for Radical Ontic Structural Realism. *The British Journal for the Philosophy of Science*, 66(3):605–634.

[Leitgeb and Ladyman, 2008] Leitgeb, H. and Ladyman, J. (2008). Criteria of Identity and Structuralist Ontology. *Philosophia Mathematica*, 16(3):388–396.

[McKenzie, 2014a] McKenzie, K. (2014a). On the Fundamentality of Symmetries. *Philosophy of Science*, 81(5):1090–1102.

[McKenzie, 2014b] McKenzie, K. (2014b). Priority and Particle Physics: Ontic Structural Realism as a Fundamentality Thesis. *The British Journal for the Philosophy of Science*, 65(2):353–380.

[McSweeney, ming] McSweeney, M. (forthcoming). An epistemic account of metaphysical equivalence. *Philosophical Perspectives*.

[Melia, 1999] Melia, J. (1999). Holes, Haecceitism and two conceptions of determinism. *The British Journal for the Philosophy of Science*, 50(4):639–664.

[Monk, 1969] Monk, J. D. (1969). Nonfinitizability of classes of representable cylindric algebras. *The Journal of Symbolic Logic*, 34(03):331–343.

[Monk, 2000] Monk, J. D. (2000). An introduction to cylindric set algebras. *Logic Journal of IGPL*, 8(4):451–496.

[Németi, 1991] Németi, I. (1991). Algebraization of quantifier logics, an introductory overview. *Studia Logica*, 50(3-4):485–569.

[Plotkin, 2000] Plotkin, B. (2000). Algebra, Categories and Databases. In Hazewinkel, M., editor, *Handbook of Algebra*, volume 2, pages 79–148. Elsevier, Amsterdam, Netherlands.

[Pooley, 2006] Pooley, O. (2006). Points, particles, and structural realism. In Rickles, D., French, S., and Saatsi, J., editors, *The Structural Foundations of Quantum Gravity*, pages 83–120. Oxford University Press, Oxford, UK.

[Psillos, 2001] Psillos, S. (2001). Is Structural Realism Possible? *Philosophy of Science*, 68(3):S13–S24.

[Roberts, 2011] Roberts, B. W. (2011). Group Structural Realism. *The British Journal for the Philosophy of Science*, 62(1):47–69.

[Ruetsche, 2011] Ruetsche, L. (2011). *Interpreting Quantum Theories: The Art of the Possible*. Oxford University Press, Oxford; New York.

[Rynasiewicz, 1992] Rynasiewicz, R. (1992). Rings, Holes and Substantivalism: On the Program of Leibniz Algebras. *Philosophy of Science*, 59(4):572–589.

[Schaffer, 2005] Schaffer, J. (2005). Quiddistic Knowledge. *Philosophical Studies*, 123(1-2):1–32.

[Sider, 2016] Sider, T. (2016). Theoretical equivalence. Locke Lectures.

[Weatherall, 2016] Weatherall, J. O. (2016). Regarding the 'Hole Argument'. *The British Journal for the Philosophy of Science*, page axw012.

[Wolff, 2012] Wolff, J. (2012). Do Objects Depend on Structures? *The British Journal for the Philosophy of Science*, 63(3):607–625.

[Worrall, 1989] Worrall, J. (1989). Structural Realism: The Best of Both Worlds? *Dialectica*, 43(1-2):99–124.