

Everything you always wanted to know about structural realism but were afraid to ask

Roman Frigg · Ioannis Votsis

Received: 23 July 2010 / Accepted: 7 April 2011 / Published online: 4 May 2011
© Springer Science+Business Media B.V. 2011

Abstract A structuralist perspective is one that sees the investigation of the structural features of a domain of interest as the primary goal of enquiry. This vision has shaped research programmes in fields as diverse as linguistics, literary criticism, aesthetics, sociology, anthropology, psychology, and various branches of philosophy. The focus of this paper is structuralism in the philosophy of science, and in particular those movements that have endeavoured to articulate a structural version of scientific realism, now commonly referred to as structural realism (SR). The paper provides a critical survey of the debates raging over structural realism: it provides explicit statements of the different positions as well as the arguments put forward to support them, clarifies how the different positions relate to one another, draws attention to hitherto neglected arguments, and evaluates criticisms launched against different strands of SR. Attention to the history of the field is paid in as far as this is essential to understanding the contemporary scene, but documenting the long and intricate development of SR is beyond the scope of this paper. We begin by introducing the set theoretic conception of structure on which many of the positions that we are concerned with rely (Section 2). In Section 3 we introduce the two main strands of epistemic structural realism, discuss the central objections levelled against them, most notably Newman's objection, and present the Ramsey sentence formulation. Section 4 is dedicated to a discussion of ontic structural realism. In Section 5 we offer some concluding remarks.

R. Frigg (✉)

Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK
e-mail: r.p.frigg@lse.ac.uk

I. Votsis

Philosophisches Institut, Heinrich-Heine-Universität Düsseldorf, Universitätsstraße 1,
Geb. 23.21/04.86, 40225 Düsseldorf, Germany
e-mail: votsis@phil-fak.uni-duesseldorf.de

Keywords Structural realism · Scientific realism · Philosophy of science · Theory change · Ramsey sentence · Newman problem · Identity of indiscernibles

Acronyms

SR	Structural Realism
ESR	Epistemic Structural Realism
DESR	Direct Epistemic Structural Realism
IESR	Indirect Epistemic Structural Realism
OSR	Ontic Structural Realism
EOSR	Eliminative Ontic Structural Realism
ROSR	Radical Ontic Structural Realism
RS	Ramsey Sentence
HW	Hermann-Weyl Principle
MR	Mirroring Relations Principle
NMA	No Miracles Argument
PMI	Pessimistic Meta Induction

1 Introduction

A structuralist perspective is one that sees the investigation of the structural features of a domain of interest as the primary goal of enquiry. This vision has shaped research programmes in fields as diverse as linguistics, literary criticism, aesthetics, sociology, anthropology, psychology, and various branches of philosophy. The focus of this paper is structuralism in the philosophy of science, and in particular those movements that have endeavoured to articulate a structural version of scientific realism, now commonly referred to as *structural realism* (SR).¹ The paper provides a critical survey of the debates raging over structural realism: it provides explicit statements of the different positions as well as the arguments put forward to support them, clarifies how the different positions relate to one another, draws attention to hitherto neglected arguments, and evaluates criticisms launched against different strands of SR. Attention to the history of the field is paid in as far as this is essential to understanding the contemporary scene, but documenting the long and intricate development of SR is beyond the scope of this paper.²

We begin by introducing the set theoretic conception of structure on which many of the positions that we are concerned with rely (Section 2). In Section 3 we introduce the two main strands of epistemic structural realism, discuss the central objections levelled against them, most notably Newman's objection, and present the Ramsey sentence formulation. Section 4 is dedicated to a discussion of ontic structural realism. In Section 5 we offer some concluding remarks.

¹ For a discussion of structural thinking in science see Rickart (1995), and for a detailed presentation of the structures used in fundamental physical theories see Muller (1998). Resnik (1997) and Shapiro (1997, 2000) advocate a structuralist position in the philosophy of mathematics. For a discussion of the relation between structuralism in mathematics and science see Brading and Landry (2006). Surveys of structuralist approaches in the humanities can be found in Caws (2000) and Williams (2005).

² For accounts of the history of certain strands of structuralism in the philosophy of science see Gower (2000), Votsis (2004, Ch. 2) Daston and Galison (2007, Ch. 5), and the relevant sections in Demopoulos and Friedman (1985), Solomon (1989), van Fraassen (1997, 2006), and French and Ladyman (2011).

2 Structures

Before discussing different types of structuralism, it is helpful to introduce a formal notion of structures. A structure S consists of (a) a non-empty set U of objects, which form the domain of the structure, and (b) a non-empty indexed set R (i.e. an ordered list) of relations on U , where R can also contain one-place relations (i.e. monadic properties).³ Unless stated otherwise, in what follows we do not make a distinction between monadic and polyadic relations (in common parlance often referred to as ‘properties’ and ‘relations’ respectively), and use the term ‘relation’ to cover both. It is often convenient to present a structure as an ordered tuple: $S = \langle U, R \rangle$. Two structures $S_1 = \langle U_1, R_1 \rangle$ and $S_2 = \langle U_2, R_2 \rangle$ are isomorphic iff there exists a one-to-one (i.e. bijective) mapping $f: U_1 \rightarrow U_2$ such that f preserves the system of relations of the two structures in the following sense: for all relations $r_1 \in R_1$ and $r_2 \in R_2$, the elements a_1, \dots, a_n of U_1 satisfy the relation r_1 iff the corresponding elements $b_1 = f(a_1), \dots, b_n = f(a_n)$ in U_2 satisfy r_2 , where r_1 is the relation in R_1 corresponding to r_2 in R_2 (i.e. have the same index in the indexed sets R_1 and R_2). If this is the case, f is called an ‘isomorphism’. The cardinality of a structure S is, by definition, the cardinality of its domain U .

It is important to notice that structures thus defined have no ‘material’ component. The objects in the domain of a structure are featureless dummies that have no property other than, perhaps, the metaphysical property of being an individual (more about this in Section 4). The relations in the structures are defined extensionally as sets of ordered tuples, and as such they have virtually no intensional interpretation.⁴ The extension of a relation is the set of ordered tuples to which it applies; the intension of a relation is the ‘material content’ or the meaning of a relation. For instance, the extension of the expression ‘being the father of’ is all ordered pairs $\langle a_1, a_2 \rangle$ of which it is true that a_1 is the father of a_2 ; the intension of this relation is fatherhood. A relation is specified extensionally if all we are told about the relation is the set of tuples to which it applies, while no specific intension is given. For the purpose of mathematical logic extensional specifications are sufficient because logico-mathematical properties of relations like transitivity or reflexivity only depend on their extension.

Structures thus defined are also referred to as ‘abstract structures’ to emphasise that neither the objects in their domain nor the relations have any material content. Abstract structures contrast with concrete structures, ones whose objects and relations are interpreted. For instance the structure with the domain $U = \{\text{Philip, Alexander}\}$ and a set R containing only the relation $r = \text{‘being the father of’}$ is a concrete structure while the structure with the domain $U = \{a_1, a_2\}$ and R containing

³ Two remarks regarding this definition of structures are in order. First, sometimes structures are defined such that they also involve an indexed set O of operations on U as a third ingredient. Although it is convenient in certain contexts to list operations separately, they are ultimately unnecessary because they can be reduced to relations (see Boolos and Jeffrey 1989, 98–99; Shapiro 1991, 63). Second, Logicians often regard a set of symbols denoting the elements of $\langle U, R \rangle$ as part of the structure; see for instance Hodges (1997, 2).

⁴ Russell (1919, 59–62) provides a detailed discussion of this aspect of structures; see also Newman (1928, 139) and Redhead (2001a, 74–75).

only the relation $r = \langle a_1, a_2 \rangle$ is an abstract structure. In what follows we always take ‘structure’ to refer to abstract structures unless stated otherwise.

Two points deserve mention. First, the extension of a relation does not uniquely determine its intension: there can be relations with the same extension but different intensions. Quine offers the by now canonical example of the co-extensional expressions ‘creature with a kidney’ and ‘creature with a heart’. But an extensional conception of relations does not entail an absolute absence of meaning. The mere identification of a given relation’s extension puts constraints on the sort of intensions that can be associated with it. For example, suppose that the extension of an indirectly known relation r is given by set E . Suppose further that E is the extension of the expressions ‘creature with a heart’ and ‘creature with a kidney’. The two expressions have different intensions, so merely knowing E does not tell us which one refers to r and hence which is the intended intension. Nevertheless, knowing E tells us that other expressions, e.g. ‘creature with a pharynx’, and their associated intensions do not refer to r for the simple reason that they possess different extensions (e.g. planaria have a pharynx but no circulatory system and hence neither a heart nor a kidney).

Second, structures can be instantiated in different systems. As an example consider the structure S with the domain $U = \{a_1, a_2, a_3\}$ and a set R that contains only one relation, $r = \{\langle a_1, a_2 \rangle, \langle a_2, a_3 \rangle, \langle a_1, a_3 \rangle\}$. Different physical systems can have structure S : three persons of different heights, where r is ‘taller than’; three not equally expensive books, where r is ‘more expensive than’; three suitcases of different weights, where r is ‘heavier than’; etc. So the structure S can be instantiated in many different systems, but the structure itself does not depend on these instantiations, and can be defined entirely independently of them.⁵

Two further items are frequently referred to as structures in discussions about SR: equations and Ramsey sentences. This raises the question of how they relate to the above notion of structure. Let us take equations first. The main idea is that equations specify relations on a domain. As a simple example consider Hooke’s law, which says that the restoring force F of a spring is proportional to its elongation s : $F = -ks$, where k is a constant. This equation involves two quantities that can take real values, so the domain of the structure is the real plane \mathfrak{R}^2 . The relation defined by the equation is $r = \{(x, y) : y = -kx\}$, i.e. the set of all tuples $(x, y) \in \mathfrak{R}^2$ such that y is equal to $-kx$. If we now let R be the set containing (only) the relation r , then $S = \langle \mathfrak{R}^2, R \rangle$ is the structure defined by $F = -ks$. We also see that the ‘material character’ of the original terms in the equations has dropped out: the elongation s became the real number x and the force F became the real number y . But nothing in the structure depends on x being the elongation and y being the force of a spring; in fact, we could interpret x and y differently and nothing in the structure itself would change, just as nothing depends on the relation ‘being the father of’ in the above example. If the equations become more complex (many laws of physics are differential equations), the structures specified by the equations become more complex too. But the basic idea of the connection between structures and equations remains the same. Hence, referring to equations as structures is an elliptical way of referring to the structure specified by the equation in the way just indicated.

⁵ There is, of course, a question of how to understand the ontological status of such structures. For a discussion of this point see Hellman (1989, 1996, 2001), Resnik (1997) Shapiro (1983, 1997, 2000).

Next in line are Ramsey sentences. At this point we merely present a rough sketch of the leading idea behind the Ramsey sentence approach; we return to the issue in detail in Sections 3.3 and 3.4 where we see that things are in fact far less straightforward than one would expect. The approach is best introduced with an example. Consider the simple physical law that any two bodies with opposite electrical charges attract each other. A formal rendering of this law in first-order logic is $\forall x\forall y [(Px \ \& \ Ny) \rightarrow A(x, y)]$, where ‘ P ’ stands for the property of having positive charge, ‘ N ’ for having negative charge, and ‘ A ’ for the relation of attracting each other. The sentence $\forall x\forall y [(Px \ \& \ Ny) \rightarrow A(x, y)]$ is true in a structure, namely the structure in which U consists of all objects in the world, π is the set of all positively charged objects, ν is the set of all negatively charged objects and α is the set of all tuples of objects that attract each other. If all tuples with one element from π and one from ν are in α then the sentence is true. Notice that for the purpose of expressing the relevant structure, nothing depends on π being the set of positively charged objects, etc. All that matters is that there are sets π , ν , and α on a domain such that tuples with one element from π and one from ν are in α —so, again, the ‘material’ content of P , N , and A has dropped out. Given that, why include it in the first instance? The Ramsey sentence (RS) is a tool to get rid of the material aspects of the sentences in question right from the start. We obtain the RS of a sentence by replacing predicates by variables and existentially quantifying over them. The RS in our example is $\exists X\exists Y\exists Z\forall x\forall y [(Xx \ \& \ Yy) \rightarrow Z(x, y)]$. The sentence says that there exist Relations X , Y and Z such that for all objects x and y the claim $(Xx \ \& \ Yy) \rightarrow Z(x, y)$ is true. Any mention of charge and attraction has dropped out and all that matters for the truth of RS is that there is a domain with appropriate relations defined on it, no matter how these are interpreted.

Ramsey sentences can express a variety of claims made by a theory, including equations. For this reason structures are sometimes analysed in terms of the RS with the understanding that equations are covered by that approach. We follow this convention and don’t discuss equations separately.

What we have discussed so far is the common set-theoretic notion of structure. This notion is central to our discussion because until recently most versions of structural realism have been based on this notion of structure. However, the hegemony of the set-theoretic approach has come under attack. Landry (2007) argues that the focus on the set theoretic notion of structure, or indeed that on any single formal framework, is unduly restrictive because different contexts demand different notions of structure. Muller (2010) makes the stronger assertion that both set theory and category theory are inadequate frameworks within which to formulate the claims of SR and calls for a new theory of structure that axiomatises the notion instead of defining it in terms of other known notions. Leitgeb and Ladyman (2008) introduce a graph-theoretic notion of structure,⁶ and Floridi (2008) relates structuralist ideas to information theory. Roberts (2011) identifies quantum mechanical structures with symmetry groups. Since most of the papers we discuss

⁶ A graph (in this sense) is a mathematical structure whose specification requires two types of things: edges and nodes. Intuitively we may think of nodes as objects and edges as relations. Structural realists like Ladyman are interested in so-called ‘unlabelled’ graphs because in such graphs different nodes are indistinguishable, i.e. no additional information is given about the nodes (no labelling or intension) other than potentially the edges that link them to other nodes. Leitgeb and Ladyman utilise graph theory to show that even weak versions of the principle of identity of indiscernibles can be violated.

in this review rely, either explicitly or at least implicitly, on the set-theoretic notion of structures, we stick to this notion here.

3 Epistemic structural realism

An important family of structuralist positions are answers to the fundamental epistemic question: what can we know about the world? *Epistemic structural realism* (ESR) is the view that all we can know about the unobservable world is its structure. This contrasts with fully-fledged realism ('realism' henceforth), which does not deny that we have structural knowledge, but holds that we also have non-structural knowledge. Different versions of ESR are at variance about how they qualify the claim, and about how they justify it. One of the most important qualifications concerns the issue of what counts as observable. Broadly, there are two views on this. The first regards, as is common in contemporary philosophy of science, properties or objects that are accessible to observation with the unaided senses as 'observables', and those (putative) properties or objects that are not as 'unobservables'. There are of course important questions about where and how to draw the line between observables and unobservables in that sense, but intuitively we would qualify spin, entropy, and neutrinos as unobservables, while thermometer readings and rabbits come down on the side of observables. *Direct epistemic structural realism* (DESR) holds that we can have fully-fledged knowledge (i.e. structural and non-structural knowledge) about the observable parts and aspects of the world (thus understood), but that our knowledge of the unobservable is only structural. *Indirect epistemic structural realism* (IESR) is a descendant of indirect realism, the position that we only have direct epistemic access to our sense data, our perceptions or at any rate something sensory, and only through them do we have (indirect) access to the world itself. For this reason not only neutrinos and the like, but also chairs, tables and rabbits count as unobservable, and claims about the external world *tout court* can be only structural.⁷

It has become customary to present ESR in a 'Ramsey sentence version'. We do not follow this convention; the Ramsey sentence is only one way to formulate ESR and should not be conflated with ESR itself. For this reason we introduce DESR and IESR in Subsections 3.1 and 3.2 respectively, and then turn to the Ramsey sentence in Subsections 3.3 and 3.4. Subsection 3.5 examines the 'nature versus structure' distinction, and Subsection 3.6 discusses an alternative yet related position, namely semi-realism.

3.1 The upward path to ESR

Following Psillos (2001a), we distinguish between different motivations for ESR, namely the 'upward path' and the 'downward path'; we begin with a discussion of

⁷ The term 'structural realism' was coined by Grover Maxwell (1968). Our distinction between DESR and IESR corresponds roughly to Ainsworth's (2009) distinction between 'weak ESR' and 'strong ESR'. We prefer 'DESR' and 'IESR' to 'weak ESR' and 'strong ESR' because DESR sanctions some claims that are stronger than claims sanctioned by IESR. For instance, DESR accepts that we can have knowledge of relations between entities that have no perceptual analogue, something that IESR denies.

the former and turn to the latter in Section 3.2. Those who defend ESR ‘bottom up’ take as their starting point the question of what we can assert about the world on the basis of our basic sensory experiences and argue that all we can infer from these experiences are structural features of the world. Hence, proponents of this view end up defending IESR (although it is not necessary that they must). This defence takes different forms depending on what sensory experiences are regarded as epistemically primary. One approach, which originates in the work of Bertrand Russell, focuses on perception; another approach due to Henri Poincaré departs from the notion of transmissibility.

3.1.1 *The argument from perception*

The first statement of IESR can be found in Russell’s *The Problems of Philosophy*, first published in 1912. Having recently read and been influenced by the British empiricists, Russell regarded the basic units of perception, which at the time he took to be sense-data, as the foundation of all knowledge. He put forward the view that although we have good reasons to believe that the causes of the sense-data are physical objects, all we can *know* about these objects is their structure:

‘... although the *relations* of physical objects have all sorts of knowable properties, derived from their correspondence with the relations of sense-data, the physical objects themselves remain unknown in their intrinsic nature.’ (1912, 17; original emphasis)

What are these knowable properties? Russell thinks that they are logico-mathematical properties like transitivity or reflexivity. While of *appearances* we can know both their logico-mathematical properties and their intrinsic natures—what we earlier identified as being specifiable in terms of full-blown intensions—we can know only the logico-mathematical properties of *physical objects themselves* (which, of course, is not to deny that physical objects have intrinsic natures; the claim is just that we cannot know them). So Russell concludes, as he puts it in *The Analysis of Matter*, that ‘[t]he only legitimate attitude about the physical world seems to be one of complete agnosticism as regards all but its mathematical properties’ (1927, 270). This is a structuralist position because knowing the logico-mathematical properties of relations without knowing anything else about either the intension of the relations or the relata is tantamount to saying that we only know the structure of the external world.⁸ It is worth emphasising how comparatively little this is: on this view all we can know about the world is that there are a number of objects, about which we cannot assert anything beyond the fact that they are objects which stand in relations that are reflexive, anti-symmetrical, transitive, and so on, but about whose intension we can assert next to nothing.

⁸ Maxwell (1971 18–19) summarises this position as the claim that we cannot know the first order properties of physical objects and that we can only know their second or higher order properties. This way of stating the position is misleading in two ways. First, it is important to notice that this use of ‘first order’ and ‘second order’ bears no connection with the distinction between first and second order logic, which will become important later on. Second, and more importantly, if our knowledge is limited to structural features, then even first order properties can be known, albeit of course only structurally.

Before turning to the question of how this position can be justified, it is worth drawing attention to a possible misunderstanding. ESR (both direct and indirect) is often summarised in the slogan that while we can have knowledge about unobservable *relations*, we have to remain agnostic about the intrinsic nature of unobservable *objects*. This suggests that realism and ESR agree that we can have knowledge about unobservable relations, and that the bone of contention is whether we can also have knowledge about unobservable objects. This construal of the disagreement is misleading because it suggests that realism and SR also agree on *what* we can know about such relations (as opposed to only agreeing *that* we can have knowledge about them), which need not be the case. The realist typically believes that we know the intensions of unobservable relations, i.e. that we know what the relation ‘in itself’ is. For instance, the realist believes that when we say that mercury is denser than iron, we do have direct knowledge of the unobservable relation ‘denser than’. ESR denies this by insisting that all we can know about the relation are formal properties like transitivity, which derive from the relation’s extension. The same holds true for monadic relations and objects. While the realist believes that we have knowledge of monadic relations and objects, the structuralist holds that all we can know about monadic relations is the cardinality of their extension (and the logico-mathematical relations of any relations they stand in) and all we can know about objects, other than structural knowledge of the monadic and polyadic relations they instantiate, is that they are ‘some distinct things’. The emphasis on relations in debates over SR is due to the fact that polyadic relations can have all kinds of interesting formal relations, while not much of interest can be said about monadic relations and objects at the purely formal level (in fact, the only formal property of monadic relations is, as we have just pointed out, the cardinality of their defining set). What lies at the heart of the conflict between realism and SR is the issue whether we can have substantive non-extensional knowledge with respect to unobservables, and this question equally arises in the case of polyadic relations, monadic relations, and objects.

In *The Analysis of Matter* (1927) Russell presents the first detailed defence of IESR.⁹ The centrepiece of this defence is a causal theory of perception which rejects ‘the view that perception gives direct knowledge of external objects’ (*ibid.*, 197): we only have direct knowledge of the ‘intrinsic character’, ‘nature’, or ‘quality’ of *percepts*, i.e. the basic units of our perception, but not of objects in the external world.¹⁰ In other words, percepts are the only things of which we can have knowledge by acquaintance. Nevertheless, percepts lie at the end of causal chains which originate in the external world, and they—the percepts—encode information about the external world. Therefore the only way to attain knowledge of the external world, according to this view, is to draw inferences from our perceptions. The crucial

⁹ As Demopoulos and Friedman (1985, 625–627) point out, Russell’s position has close affinities with other work done at the time, in particular by Schlick and Carnap. For further discussions of these positions see Creath (1998), Psillos (1999, Ch. 3; 2000a, 2000b, 2006b) and Salmon (1994).

¹⁰ Many readers familiar with Russell’s sceptical attitude towards causation in *The Problems of Philosophy* (1912) and in *Mysticism and Logic* (1918) may find his endorsement of causation here puzzling. In spite of his scepticism, it is well known among Russell scholars that a deflated notion of causation played a central role in his philosophy.

question then becomes: what inferences can we draw from percepts and what underwrites these inferences? Russell's answer to the first part of the question is clear cut: all we can infer from percepts is structure. In response to the second part of the question, Russell suggested the adoption of two principles:

Helmholtz-Weyl Principle (HW): Different effects (i.e. percepts) imply different causes (i.e. stimuli/physical objects) (*ibid.*, 255).^{11,12}

Mirroring Relations Principle (MR): Relations between percepts correspond to relations between their non-perceptual causes in a manner which preserves their logico-mathematical properties (*ibid.*, 252).

Armed with these principles, Russell argues that from the structure of our perceptions we can 'infer a great deal as to the structure of the physical world, but not as to its intrinsic character' (*ibid.*, 400); in other words, he argues that all we can assert is that the structure of our perceptions is (at best) isomorphic to the structure of the physical world.¹³

In the late 1960s and early 1970s, Grover Maxwell published a series of articles defending an epistemic version of structural realism that owes much to Russell.¹⁴ Echoing his predecessor, he posits that

'...although we have no *direct* knowledge of the [things in themselves], the bulk of our common sense knowledge and our scientific knowledge *is* of them ... all of this knowledge is purely structural.' (1968, 155).

Closely adhering to Russell's views on observation, Maxwell holds that 'all of the external world including even our own bodies is unobserved and unobservable' (*ibid.*, 152). For both philosophers, 'observed' and 'observable' denote items of experience that are 'wholly in the mind'.¹⁵

Like Russell, Maxwell adopts a causal theory of perception, and also like Russell he emphasises that 'it is *not* essential to the position [i.e. IESR] that the sense impressions or perceptual experiences, or whatever we decide to call them "resemble" the physical

¹¹ Psillos (2001a) suggested this name for the principle on the basis of Helmholtz's and Weyl's appeal to it. It is worth noting that Russell sometimes uses the principle in its contrapositive (but equivalent) form, namely as the claim that same causes imply same effects. Quine independently endorses a modified version of the HW principle focussing on similarity rather than sameness (1998, 19). The principle in one form or another has also been independently endorsed by Locke in *An Essay Concerning Human Understanding* ([1690] 1975, Book II, Ch. XXXII, §15), Hume in the *Treatise* ([1739] 1975, Book II, Part III, §1), Descartes in the 6th Meditation ([1641] 1993) and Mill in *A System of Logic* ([1874] 2008, p.423).

¹² Stimuli, according to Russell, are 'the events just outside the sense-organ' (1927, 227). They are thus classified as physical events.

¹³ See also (*ibid.*, 251, 253, 254, 263–4, 270–1; 1919, 59–62; 1912, 32, 34). Although Russell uses different terminology, his definition of structure (see, for example, (1927, 250)) is equivalent to the definition of structure given earlier. For more on this issue see Solomon (1990).

¹⁴ Maxwell also credits Poincaré, Schlick, and Wittgenstein, as well as Beloff, Mandelbaum, Aune and Pepper with having developed versions of ESR (see his 1968 for references).

¹⁵ This understanding of these terms, of course, conflicts with the prevalent understanding in the scientific realism debate. Seemingly paradoxically, Maxwell is best known among philosophers of science for his critique of the observable/unobservable distinction; see his (1962). The apparent tension is dissolved once we realise that in the context of his discussion of ESR the entire external world is unobservable, and that therefore the distinction he criticised in his (1962) is of an altogether different kind.

objects which may be among their causal antecedents' (*ibid.*, 155). All we need is that 'at least a certain subset of the features of the [sense] impression are isomorphic with a subset of the features of the physical object' (*ibid.*, 156). Without this type of correspondence, Maxwell insists, there can be no knowledge of the external world. His justification for this requirement proceeds along familiar Russellian lines and invokes principles like HW (*ibid.*, 156) and the claim that causal chains leading up to our perceptions are structure-preserving, i.e. MR (1971, 25).

The *argument from perception* can then be summarised as follows:

- (1a) All knowledge is ultimately based on perceptions.
 - (1b) We can have both structural and non-structural knowledge about perceptions.
 - (1c) We have no good reason to believe that the non-structural aspects of perceptions can tell us anything about the non-structural aspects of their external world causes.
 - (1d) We have good reason to believe (given MR and HW) that the structure of our perceptions is isomorphic to the structure of their external world causes.
- ∴ We have good reason to believe that we can have knowledge about the external world and that this knowledge is only structural.

Premise (1a) is the basic posit of traditional empiricism, a discussion of which is beyond the scope of this paper. Premise (1b) is uncontroversial. Premise (1c) seems to have some scientific basis. Science tells us that two people need not have *qualitatively* identical perceptions of the same object to correctly identify that object.¹⁶ As such the qualities of perceptions do not seem to tell us anything about the qualities of their external world stimuli. One concern with this whole approach is that it presupposes that external world stimuli have qualities, i.e. non-structural aspects. We do not find this concern critical, for the IESRist may easily take a softer stance by advocating agnosticism with respect to the existence of such non-structural aspects. A more serious concern is that supposing that such aspects do exist, it is sufficient for only some of these (even a tiny minority) to be reflected in the non-structural aspects of the perceptual world to yield non-structural knowledge. However, the structuralist need not commit herself to the (strong) claim that we have no non-structural knowledge; all she needs to claim is that we have no justification for believing that we indeed have non-structural knowledge of the external world (or indeed to identify certain non-structural bits as knowledge about the external world rather than just our perceptions).

There are subtleties concerning premise (1d). MR is strong enough to establish an isomorphism between the unseen world and our percepts, and hence the argument is valid without any appeal to HW, which then becomes an idle wheel. By contrast, HW is not strong enough on its own to establish an isomorphism between the structure of the world and the structure of our percepts because it does not preclude different stimuli having the same perceptual effect.¹⁷ For those unwilling to assume MR this leaves two options. The first option involves the addition of a further

¹⁶ It is important to emphasise that this does not mean that *structural* differences in perceptions have no corresponding *structural* differences between external world causes. This kind of correspondence is in fact required by HW.

¹⁷ Indeed, MR entails HW but not vice versa.

assumption to make HW strong enough to establish a bijective relation. The missing assumption is in fact the converse principle of HW, which we hereafter call ‘WH’, namely that different stimuli/physical objects imply different percepts. The second option involves coming to terms with a weaker form of structural knowledge, one that rests content with at best some sort of embedding.¹⁸

What justifies HW and WH? It is hard to imagine how we can interact with the world without accepting some version of these principles. Consider what would happen if HW did not hold at least most of the time; that is, suppose that the same (or sufficiently similar) stimulus repeatedly gave rise to different (or sufficiently dissimilar) perceptions in the same person. Recognising and evading a predator would then be a miracle since the same (or sufficiently similar) stimulus, e.g. a tiger, would hardly ever give rise to the same (or sufficiently similar) perceptions, e.g. a tiger perception. A similar justification can be given for WH. Suppose that different (or sufficiently dissimilar) stimuli repeatedly gave rise to the same (or sufficiently similar) perception in the same person. Correctly gauging changes in people’s behaviour would then be a miracle since different (or sufficiently dissimilar) behaviour, e.g. smiling vs. crying, would hardly ever give rise to different (or sufficiently dissimilar) perceptions. Hence, having a neurophysiology that functions in accordance with both HW and WH confers significant evolutionary and learning advantages.

Psillos (2001a, S13–S16) criticises Russell’s justification for premise (1d). More specifically, he holds that Russellian ESR faces a dilemma: On the one hand, the HW principle is too weak to establish isomorphic relations, as required by ESR, for it can only establish embeddability relations. Without isomorphic relations, he argues, the epistemic structural realists cannot maintain inferential knowledge about the structure of the external world. On the other hand, HW and WH allow for the establishment of isomorphic relations but in so doing they forbid any variance between the structure of the external world and the structure of our perceptions. This, Psillos argues, is too stringent a constraint on knowledge.

In reply, Votsis (2005) argues that Psillos’ ‘dilemma’ for the Russellian version of ESR is ill conceived. With respect to the first horn of the dilemma, Votsis argues that Psillos neglects the MR principle, which can be invoked to establish isomorphic relations. With respect to the second horn, Votsis argues that ESR is not, contra Psillos, a position that *requires* isomorphic mappings. The claim is rather that an isomorphism between the structure of percepts and the structure of reality is the ideal case, but we can fall short of this ideal and still have knowledge. Some variance between the two structures is therefore allowed. Having said this, for perception to be a reliable guide to the external world this variance must not be too prevalent.¹⁹

¹⁸ A structure $S_1 = \langle U_1, R_1 \rangle$ is *embedded* into a structure $S_2 = \langle U_2, R_2 \rangle$ iff there exists an injective mapping $f: U_1 \rightarrow U_2$ such that f preserves the system of relations of S_1 in the following sense: for all relations $r_1 \in R_1$ and $r_2 \in R_2$, if the elements a_1, \dots, a_n of U_1 satisfy the relation r_1 then the corresponding elements $b_1 = f(a_1), \dots, b_n = f(a_n)$ in U_2 satisfy r_2 , where r_1 is the relation in R_1 corresponding to r_2 in R_2 (i.e. have the same index in the indexed sets R_1 and R_2). We typically speak of embeddings when the cardinality of U_2 is greater than the cardinality of U_1 . In those cases, an embedding is just an isomorphism between S_1 and a part—a ‘substructure’ as it is sometimes called—of S_2 .

¹⁹ How much variance can be afforded before the reliability of perception breaks down is not an easy question to answer.

Otherwise, or so the argument goes, we would not be able to rely on perceptions, as often as we do, to successfully navigate the external world.²⁰

Another issue worth discussing concerns the nature of the relation of the isomorphism or embedding relation. These are relations that, by definition, hold between set-theoretical structures. Yet up to now we have been attributing such relations to non-mathematical things like parts of the perceptual and parts of the physical world. The position has thus tacitly assumed that these parts are structured in a way that can be represented set-theoretically. More controversially, it has tacitly assumed that the structures assigned to these parts are unique and objective, i.e. independent of our interests. Establishing the truth of this assumption is not a straightforward matter. We return to this issue at the end of section 3.4.²¹

3.1.2 The transmission argument

An alternative defence of IESR departs from Poincaré's observation in *The Value of Science* that 'nothing is objective which is not transmissible, and consequently that the relations between the sensations can alone have an objective value' (1913, 348).²² This argument needs some disentangling. Like Russell, Poincaré assumes that we only have direct epistemic access to our perceptions. But our perceptions are by their very nature private and cannot be transmitted. We cannot, for instance, communicate our sensory experience of two particular shades of green to someone else.²³ What we can communicate, however, are relations between these sensory experiences, for instance that one shade of green is darker than the other. Or almost. In fact, what we can transmit is not the 'full' sensation of *darker than*, which is as private as the two shades of green themselves. What we can transmit is the structure of *darker than*. So what we really transmit, on this view, is that the two sensations stand in a transitive, non-reflexive, and asymmetrical relation.

While Poincaré does not further specify what kind of transmission he has in mind, Quine makes the argument more specific by couching it in terms of the transmission of knowledge by dint of language. He invites us to consider the following scenario (1968, 161). Send a man into a room whose content we don't know, and then let him report to us what he has seen. He uses all kind of verbal expressions to describe to us what is in the room. But since we have seen nothing of what he has seen, all he manages to convey to us is structural features of its interior because he cannot

²⁰ Psillos raises another objection in that paper. He claims that the structural realist cannot account for the possibility that the unobservable world may have extra structure not manifested in the perceptual world. This claim is incorrect. Russellian ESR just requires that all, or at least most, perceptual structures have corresponding external world structures, not vice-versa.

²¹ This problem is most acute in the case of percepts since it is possible that in light of the same set of stimuli different perceivers attribute different structures to their perceptions. The epistemic structural realist may be able to bite the bullet here so long as divergent attributions of structure are the exception rather than the rule.

²² Russell (1948, 485–6) and Carnap (1928, §16) make similar remarks about the intransmissibility of everything but structure.

²³ Against this view it has been objected that one can communicate, for example, a feeling of sadness by reporting it or by using specific facial expressions. Although we agree that someone can communicate in this way that they have a sad feeling, this does not imply that the person can communicate their particular sensory experience of sadness.

communicate perceptions. Indeed had we been in the room with him, the argument goes, he would still not be able to communicate his perceptions though, of course, we would be able to connect his structural characterisations with our own perceptions.²⁴

The *argument from transmission* can then be summarised as follows:

- (2a) All knowledge (i.e. public and private) is ultimately based on perceptions.
 - (2b) Perceptions consist of individual sensory experiences and their relations.
 - (2c) If something is public knowledge about the external world then it is transmissible via language.
 - (2d) The content of individual sensory experiences is not transmissible via language.
 - (2e) The logico-mathematical properties of relations between sensory experiences are transmissible via language.
- ∴ Only the logico-mathematical properties of relations between sensory experiences but not the individual sensory experiences themselves can be publicly knowable.

Premise (2a) is identical to the first premise of the argument from perception. Premise (2b) is uncontroversial since it merely decomposes perceptions into two components. Premise (2c) is perhaps less obvious. It gets its plausibility from the idea that knowledge and in particular scientific knowledge is and ought to be publicly accessible, e.g. presentable in textbooks. Premise (2d) also has some plausibility. Language may help evoke similar feelings or sensory experiences in other persons but there is no good reason to think that the content of those experiences will be identical. Even so, this leaves untouched the question of whether such content can be transmitted in non-linguistic ways. Premise (2e) on the other hand is less troublesome. Being able to linguistically communicate the logico-mathematical properties of relations between sensory experiences is a commonplace affair. Think of the earlier example of the relation ‘darker than’: we have no problems communicating this relation (as well as other relations like it) and the logico-mathematical properties it satisfies via language.

We would like to emphasise that the argument, at least as presented by Poincaré, does not conclude that we in fact have publicly accessible structural knowledge; the argument only establishes the conditional claim that *if* we have any knowledge about the external world at all, *then* that knowledge is only structural, thus leaving open the possibility that we have no such knowledge at all. In other words, the argument establishes that relations between sensory experiences are the only candidates for public knowledge. To attain the desired ESR conclusion that we have structural knowledge of the external world, the argument needs to be augmented with an additional premise: (2f) We do have knowledge of the external world. Adding this premise would not render the argument trivial (or question begging) since the premise does not presuppose what kind of knowledge we have, which is precisely the issue at stake.

²⁴ In this context, the notion of transmission is broader. For more on Quine’s structuralism see his (1969) and (1992); for a discussion of his position see Rosner (1996).

3.1.3 The argument from predictive power

A different argument in defence of IESR has been suggested by Votsis (2004, Ch. 6).²⁵ This argument focuses on scientific theories and rests on the claim that of all the features of theories, only the mathematical structures plus the observable content of the theory possess predictive power: it seems that we do not need to presuppose anything non-structural about the entities postulated by the theory in order to make the relevant testable predictions. This is significant because all we can test in experiments is how accurately our predictions match empirical findings. If we then also grant that epistemic warrant is grounded in empirical testing, we should only take a theory's mathematical structures (and its observable content) seriously. This is the *argument from predictive power*:

- (3a) Epistemic warrant is ultimately conferred onto a claim solely through successful empirical tests, i.e. through that claim's ability to contribute to successful predictions.
 - (3b) The only parts of science that are indispensable for the production of predictions are empirically interpreted mathematical structures.
 - (3c) Empirically interpreted mathematical structures can reveal no more than the unobservable world's structure.
- ∴ Of the claims about the unobservable world only structural ones can attain epistemic warrant.

Premise (3a) is a posit closely related to traditional empiricism so, as before and in the interests of expediency, we sidestep its discussion. Possible challenges to Premise (3b) are very similar to the ones levelled against the first premise of the argument from the history of science (to which we turn below): one could argue that more than just empirically interpreted structure is involved in deriving predictions; in other words, one could argue that although necessary, empirically interpreted structures are not sufficient to derive predictions. This would undermine the conclusion. However, as we will see below when discussing Psillos' objections, what seems to be non-structural may, upon closer examination, well turn out to be structural and so the critic would at least have to produce a convincing example of a non-structural element which is crucial to a theory's predictive power.

Premise (3c) is perhaps the most contentious one. It allows empirically interpreted mathematical structures to reveal structural features of the unobservable world. That is a presupposition that some are unwilling to grant. For example, van Fraassen (2006) argues that empirically interpreted mathematical structures reveal something about the world of phenomena but nothing about the unobservable world. A pertinent question in this dispute is whether successful interventions in the unobservable world that depend on empirically interpreted mathematical structures can be accounted for without presupposing at least some structural knowledge about the unobservable world.

3.2 The downward path to ESR

Those who defend ESR 'top down' take as their starting point actual scientific theories and argue that once we strip away the non-structural elements we reach true

²⁵ This argument could *mutatis mutandis* also be put forward in support of DESR.

scientific knowledge at the bottom. The advocates of this approach endorse DESR. The downward path to ESR is articulated and defended in two different ways, one based on the history of science, the other on the nature of mathematical representation.

3.2.1 *The argument from the history of science*

The argument from the history of science originates in Poincaré's *Science and Hypothesis* (1905).²⁶ He starts by introducing what is nowadays referred to as the argument from the pessimistic meta-induction (PMI) (*ibid.*, 160): departing from the observation that the history of science is a graveyard of once successful but now discarded theories, we conclude that currently successful theories will very probably turn out to be false as well, and that we should therefore not believe that the empirical and explanatory success of scientific theories warrants the claim that they are (at least approximately) true.²⁷ Poincaré responds to this antirealist challenge in two steps. He first points out that important elements of theories, namely equations, survive theory change:

‘No theory seemed established on firmer ground than Fresnel’s, which attributed light to the movements of the ether. Then if Maxwell’s theory is to-day preferred, does that mean that Fresnel’s work was in vain? No; for Fresnel’s object was not to know whether there really is an ether, if it is or is not formed of atoms, if these atoms really move in this way or that; his object was to predict optical phenomena. This Fresnel’s theory enables us to do to-day as well as it did before Maxwell’s time. The differential equations are always true, they may be always integrated by the same methods, and the result of this integration still preserves their value.’ (*ibid.*, 160–1)

Poincaré then claims, and this is the second step, that equations describe certain aspects of reality which we should be realists about:

‘It cannot be said that this is reducing physical theories to simple practical recipes; these equations express relations, and if the equations remain true, it is because the relations they express preserve their reality. They teach us now, as they did then, that there is such and such a relation between this and that; only that something which we called *motion*, we now call *electric current*. But these are merely names of the images we substitute for the real objects which Nature will hide from our eyes. The true relations between these real objects are the only reality we can attain [...]’ (*ibid.*, 161)²⁸

²⁶ Poincaré is often thought of as a conventionalist anti-realist, not only with regard to geometry but also physics. However, Maxwell (1968), Giedymin (1982), Worrall (1982; 1989; 1994), Zahar (1996; 2001), Stump (1989), Psillos (1995; 1999), Gower (2000), and Redhead (2001a) argued, in our view convincingly, that Poincaré is an ESRist. Some have also argued that Duhem ([1914]1991), another alleged conventionalist, actually held an ESR position very similar to Poincaré’s (Worrall 1989; Chakravartty 1998; Gower 2000; and Zahar 2001).

²⁷ For a discussion see Laudan (1981).

²⁸ Poincaré’s second historical example is the fact that some of the equations describing Carnot’s heat engines survived when the conception of heat as a material fluid (called ‘caloric’, on which Carnot’s theory was based, was abandoned (1905, 165).

This argument was taken up by Worrall (1982; 1989; 1994; 2007), who put it into the context of the modern debate between scientific realists and anti-realists, and made the structuralist nature of Poincaré's realism explicit.²⁹ Worrall observes that there are two powerful arguments in this debate, which, however, pull in opposite directions. The first one is the pessimistic meta-induction, which supports anti-realism.³⁰ The other supports realism and is the so-called *no miracles argument* (NMA) according to which the success of a scientific theory would be a miracle if the theory was not at least approximately true, and that therefore scientific realism is the *only* philosophical position that does not make the predictive success of science a miracle.³¹ Worrall concedes that both arguments have intuitive appeal despite pulling in opposite directions, and he submits that a tenable position in the realism versus antirealism debate has to do justice to both. Structural realism, he claims, is the position that underwrites both arguments because it occupies the appropriate middle ground between realism and antirealism, or, in Worrall's words, it is the position that allows us to have 'the best of both worlds':

'This largely forgotten thesis of Poincaré's seems to me to offer the only hopeful way of *both* underwriting the "no miracles argument" *and* accepting an accurate account of the extent of theory change in science. 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*.' (Worrall 1989, 157)

There are two important points here. The first is the association of equations with structure. Like Poincaré, Worrall believes that what survives scientific revolutions is equations, which are taken to 'encode' the structure of the theory's target domain (in his example the behaviour of light) and that the preservation of equations through theory change therefore amounts to the preservation of structure.³² The second is the association of the non-structural part of a theory with the 'nature' of things. Fresnel thought that light consisted of vibrations transmitted through an all-pervading medium, the ether. So Fresnel was wrong about the nature of light, because we now know that there is no ether. But nevertheless he got the structure of light, which is captured in his equations, right. We now see how structural realism is the best of both worlds: it underwrites the no miracles argument because it takes the success of

²⁹ A line of argument very similar to Worrall's is developed in Zahar (1994; 2001, Ch. 2; 2004). Zahar adopts a notion of observability that is very similar to that of Russell, a move which likens his position to IESR. Worrall has recently also flirted with a Russellian notion of observability, although he has not sanctioned this notion in print.

³⁰ Worrall (personal communication) now endorses a different, and rather deflationary, understanding of PMI. In its original formulation the argument assumes that scientific revolutions bring about substantive changes. Worrall now denies this and thinks that what was thrown overboard in scientific revolutions such as the shift from Fresnel's to Maxwell's theory was mere 'metaphorical puff', which may have been of some heuristic value but had no cognitive import.

³¹ See Boyd (1985), Musgrave (1988) and Psillos (1999) for a statement and defence of NMA. Critical discussions can be found in Howson (2000) and Magnus and Callender (2004).

³² The central equations of a theory need not encode the *entire* structure of the theory; all mathematically definable relations in the theory, for instance measurement scales, contribute to the theory's structure. For a discussion of this point see Redhead (2001a).

science to reflect the fact that we have got the structure of the world right; it underwrites the pessimistic meta-induction argument because it concedes that what theories prescribe as the nature of objects gets dismissed in the wake of a scientific revolution.

The *argument from the history of science* can then be stated as follows:

- (4a) Only two elements of a theory get preserved through theory change: (a) the theory's mathematical formulation,³³ and (b) the empirical interpretation of the theory's terms.
 - (4b) A theory's mathematical formulation 'encodes' the structure of that theory's target domain.
 - (4c) Preservation of an element is a reliable guide to its (approximate) truth.
 - (4d) Non-preservation of an element is a reliable guide to its (approximate) falsity.
- ∴ The preservation of structural elements through theory change is a reliable guide of their (approximate) truth. The non-preservation of non-structural elements is a reliable guide of their (approximate) falsity.

Two qualifications are in order. First, by 'elements' we mean statements about the world that have a truth value. The elements can be true, approximately true, false or approximately false.³⁴ The brackets around 'approximate' are shorthand for 'truth or approximate truth' and 'falsity or approximate falsity'. Structural elements are truth-valued statements whose content is purely structural.³⁵ Second, for most structural realists this 'purely structural' clause concerns only the unobservables. What counts as an unobservable differs in some of these accounts.

All four premises raise serious questions. Premises (4c) and (4d) incorporate an instance of NMA: history tells us which parts of the theory are genuinely successful/unsuccessful (namely the ones that survive/perish), and by NMA we regard those that are successful/unsuccessful as (approximately) true/false. As such it is subject to the kind of objections anti-realists raise against the use of NMA (see, for example, Van Fraassen 1980). The truth of Premise (4b) depends on how exactly one construes the notion of structure and, as the scare quotes indicate, the notion of encoding structure is far from clear. We turn to this issue in Section 3.3 where we discuss Newman's problem. The tenability of Premise (4a) depends on one's understanding of the history of science as well as its future. There are two objections to it. The first, explicitly considered by Worrall (1989), raises doubts about how typical the Fresnel-Maxwell case is in the history of science and hence raises doubts about the validity of the generalisation; the second takes issue with the assumption

³³ This premise should not be taken to refer only to those parts of a theory which are explicitly mathematised but also to those that can be given a mathematical formulation.

³⁴ The introduction of the class of approximately false statements is motivated by the fact that some statements assert things about the world that are neither utterly false nor approximately true. It goes without saying that any successful defence of realism will need to provide an adequate account of both the notion of approximate truth and the notion of approximate falsity.

³⁵ Although this may sound like an implicit endorsement of the syntactic view of theories, it is not. At least some structural realists, most notably Worrall (1984), believe that syntactic and semantic formulations of scientific theories are intertranslatable without loss.

that *only* structure is preserved through theory change. Let us address these objections in turn.³⁶

A number of authors have argued that the kind of preservation of equations we find in the Fresnel-Maxwell episode is in fact atypical and hence does not by itself warrant the claim that generally equations are preserved (Howson 2000, 39–40; Redhead 2001a, 86–88). In fact, Worrall grants this point and suggests ESR does not need *exact* preservation of structure to get off the ground.³⁷ It is often the case that equations of an older theory reappear as limiting cases of equations in a newer theory; for instance, Newton's equation of motion turns out to be a limiting case of the field equations of the general theory of relativity. This, he argues, is all that ESR needs if we allow for the invocation of a correspondence principle. The best-known version of such a principle, due to Heinz Post, has it that 'any acceptable new theory *L* should account for its predecessor *S* by "degenerating" into that theory under those conditions under which *S* has been well confirmed by tests' (1971, 228). In this sense the 'old' structure can be viewed as a limiting case of the 'new' structure and hence one can say that the old theory got the structure of its target domain at least approximately right, where the precise meaning of 'approximately' is given by the particular conditions alluded to in the correspondence principle.

Redhead (2001a, 83–88), although himself a supporter of ESR, raises some concerns with respect to this line of reasoning. There are prominent cases in which the correspondence principle is satisfied, and yet the relevant structures seem to be transformed to such an extent that structural continuity becomes questionable. Redhead's prime examples are precisely the cases that are supposed to underwrite the correspondence principle, namely the transition from Newtonian to relativistic mechanics on the one hand, and to quantum mechanics on the other. In the former case the metric becomes singular in the classical limit, in the latter case non-commutative observables become commutative. This, Redhead argues, has profound implications for the respective theories and one can therefore not speak of structural continuity in any straightforward way. The question then is whether there is a not so straightforward way to understand these transitions that is supportive of ESR.³⁸

The core of the second objection to Premise (4a) is the belief that more than *just* structure is preserved through theory change. Psillos offers two arguments for this conclusion. The 'uninterpreted equations objection' has it that DESR, as formulated above, only commits us to non-interpreted equations. These, however, are not

³⁶ For an in-depth discussion and reformulation of this argument that aims to address these and other objections see Votsis (2011a). Among other things, Votsis points out that it is normal to expect that not all structures get preserved through theory change since some of them enjoy no genuine predictive success. Votsis (2011c) points out that theory parts should not be considered either (approximately) true or empirically successful because they survive; rather they should be regarded as both (approximately) true and (hopefully) surviving because they are empirically successful.

³⁷ The point is mentioned already in Worrall (1989, 120–123); it is more fully developed in his (2007, 135–6, 142–4).

³⁸ Although not a structural realist himself, Schurz (2009) provides a structural correspondence theorem that may prove useful for the structural realists. As one referee pointed out there may not be a need for a general account of structural continuity through theory change. Case-by-case demonstrations may be sufficient so long as one can show that what is preserved tells us something about the structure of the unobservables.

enough to derive any predictions at all and do not therefore deserve all the epistemic credit (1999, 69, 153–4; 2001a, 21). Closely related is the ‘non-structure preservation objection’, which submits that a closer look at the historical cases at issue reveals that at least some non-structural theoretical content is retained in theory change (1999, 147–8).

The first argument provides an opportunity to clarify the position. Although structuralists often talk as if only equations survived scientific revolutions, a careful look at the ‘fine print’ reveals that the position is rather more nuanced. The crucial point is that DESR only concerns the *unobservable* part of the theory (electrons and the like): it shares with realism a commitment to the fully-fledged *observable* content of the theory, but then insists, contra realism, that when it comes to unobservables we should only take at face value the theory’s structural claims. But what is relevant to deriving predictions, according to this view, is exactly the observable part of the theory. So the structuralists’ proclamation that all that remains constant through theory change is equations should be taken as implicitly including the qualification ‘and observable empirical content’; and when Worrall asserts that Fresnel’s equations reappear in Maxwell’s theory ‘*newly interpreted* but, as mathematical equations, entirely unchanged’ (1989, 160, emphasis added), then this should be understood as the claim that the fundamental ontology of the unobservable part of the theory has changed, without, however, there being a reinterpretation of the empirical content (Votsis 2007, 63–65).

The second argument departs from a re-examination of the Fresnel-Maxwell case.³⁹ Psillos lists three assumptions that Fresnel used in the derivation of the equations, and which he regards both as non-structural and preserved in the shift to Maxwell’s theory: a minimal mechanical assumption, the principle of conservation of energy (*vis viva*), and a geometrical analysis of the configuration of the light-rays in the interface of two media (1999, 158). A detailed discussion of these assumptions would take us too far into the particulars of the said theories, but it has to be mentioned at least that some structural realists have contested that there is anything ‘un-structural’ about these principles (Votsis 2004, 76–78).⁴⁰

3.2.2 *The argument from mathematical representation*

A different argument for structural realism departs from the nature of scientific representation. Van Fraassen makes the point as follows:

‘According to the semantic approach, to present a scientific theory is, in the first instance, to present a family of models—that is, mathematical structures offered for the representation of the theory’s subject matter. Within mathematics, isomorphic objects are not relevantly different; so it is especially appropriate to refer to mathematical objects as “structures”. Given that the

³⁹ For more details on this case, see Saatsi (2005).

⁴⁰ Furthermore, as a referee has pointed out, the preservation of non-structural elements through *one* revolution is not enough to undermine ESR; only persistent preservation through many revolutions is a reliable guide to truth. We are in general agreement with this point and would like to add that this holds also for structural elements, i.e. their survival through *one* revolution is no guarantee of their (approximate) truth.

models used in science are mathematical objects, therefore, scientific theoretical descriptions are structural; they do not “cut through” isomorphism. So the semantic approach implies a structuralist position: science’s description of its subject matter is solely of structure’ (1997, 522).

Before presenting a concise summary of the argument, three qualifications are in order. First, reference to the semantic view of theories is unnecessary since one need not be a proponent of that view to uphold the two premises, i.e. that mathematical objects are describable up to isomorphism and that mathematical objects are adequate tools for scientific representation of the world. Second, although van Fraassen uses this argument to motivate a structuralist version of constructive empiricism, the argument can just as well be used to support a structuralist version of realism.⁴¹ Third, there is a question about the scope of structuralism. On a radical reading, all scientific knowledge derives from, or can be completely reformulated in terms of, scientific representation, and all scientific representation is structural.⁴² A more cautious reading leaves it open whether all knowledge is derived from structural representations but acknowledges that in certain fields—for instance in elementary particle physics—our knowledge indeed derives from (or can be completely accounted for in terms of) mathematical models/theories. We formulate the argument in a way that remains neutral with regard to the issue of scope.

The argument can then be summarised as follows:

- (5a) Mathematical objects can only be specified up to isomorphism.
 - (5b) In certain fields, adequate scientific representation of a target domain can be achieved with empirically interpreted mathematical objects and nothing else.
 - (5c) In those fields, mathematical objects are the sole carrier of scientific knowledge.
- ∴ All we can ever know about the subject matter of such a field is its structure.

Premise (5c) is the least controversial of all premises. We learn about, say, electrons from our best theories or models about electrons, and there is no way to bypass these theories or models and gain ‘extra-theory access’ to the world. Premise (5a) is the subject of an ongoing controversy in the philosophy of mathematics; the panorama of positions is surveyed in Shapiro (2000). Premise (5b) has been disputed by Frigg (2006) and Suárez (2003) who argue that scientific representation requires resources beyond mathematical structures and isomorphism.

3.3 The Ramsey sentence and Newman’s problem

We now return to Premise (4b) of the argument from the history of science, the claim that a theory’s mathematical formulation, e.g. its equations, encode the structure of that theory’s target domain. Neither Poincaré nor Worrall explicitly define the notion of structure but instead offer equations as examples of a theory’s structure. In recent

⁴¹ And indeed it is: we have come across defences of ESR along these lines in discussions on various occasions. However, Votsis (2004, ch. 6) provides the only such defence in print.

⁴² When refraining from saying that mathematical objects are the *only* adequate tools for representation, the structural realist is committed to the view that non-structural representations can always be mathematised *without loss* of content.

years, however, Worrall and Zahar (Zahar 2001, Appendix 4) and Worrall (2007, 147–149) have endorsed an approach to structures based on the Ramsey sentence (RS), which, as we have pointed out in Section 2, covers the case of equations because the to-be-Ramsified sentence can be an equation.⁴³ This approach to structures was first introduced into the discussion about SR by Maxwell (1968, 1970, 1971), and is now advocated by many (although not all) ESRists. The approach can be used by proponents of DESR and IESR alike because their positions differ only in what they regard as observable. For this reason we simply talk about ESR in the discussion to follow and take it that what is said about RS equally applies to both approaches.

Assume that the theory we are interested in is expressed in a formal second order language and that it is finitely axiomatisable. These axioms comprise, among other things, the theory's fundamental equations. Now conjoin all axioms of the theory in a single sentence T . Furthermore assume that the predicates of the theory are bifurcated into two classes according to whether or not they are 'epistemically benign'. For the purposes of RS it does not matter how one draws the line between 'benign' and 'problematic' predicates. However, in the debates over ESR it is common to use observability as the criterion and so from now on we associate 'benign' with observation predicates and 'problematic' with non-observation predicates.⁴⁴ Let us denote the non-observation predicates with P 's and the observation ones with Q 's (and, to facilitate notation later on, let P and Q be the set of all P_i and Q_i respectively). Then the theory can be written as a sentence $T(P_1, \dots, P_m, Q_1, \dots, Q_n)$, where m and n are integers. We obtain the theory's Ramsey sentence if we replace the non-observation predicates with variables and then existentially quantify over all of them. In more detail: for all $i=1, \dots, m$, replace all occurrences of P_i with a variable X_i , where the X_i are second order variables ranging over properties, and then existentially quantify over them:

$$T_R := \exists X_1 \dots \exists X_m T(X_1 \dots X_m, Q_1, \dots, Q_n).$$

This is the RS of theory T . Intuitively, the RS of a theory says that there exist *some* unobservable relations in the domain of the theory such that T holds with respect to these unobservable relations and the observable ones. As an example consider the theory $T = \forall x[(P_1x \& P_2x \rightarrow \exists y(Qy))]$, where ' P_1x ' means ' x is a Radium atom', ' P_2x ' means ' x decays radioactively', and ' Qx ' is an observational statement such as ' x is a click in a suitably placed Geiger counter' (Maxwell 1970, 168). The theory's RS is $T_R = \exists X_1 \exists X_2 \forall x[(X_1x \& X_2x \rightarrow \exists y(Qy))]$, which says that there are unobservable properties such that if x instantiates both of them, then there are clicks in a suitably placed Geiger counter.

⁴³ The Ramsey sentence originates in Ramsey (1931). For general discussions see Carnap (1966, Ch. 26), Cei and French (2006), Cruse (2004), Demopoulos (2003), Díez (2005), Lewis (1970), Psillos (2000a, 2000b, 2006b) and Sneed (1971).

⁴⁴ We use the somewhat cumbersome locution 'non-observation predicate' rather than the more common 'theoretical predicate' to avoid confusion. Non-observation predicates are taken to refer to properties in the unobservable domain. There are well known arguments for the conclusion that no bifurcation between observation and non-observation predicates is possible; see for instance Putnam (1962) and Maxwell (1962). At least for the sake of the argument we assume that some division between 'benign' and 'problematic' terms can be drawn.

Before discussing its philosophical significance, we would like to mention four important properties of RS.⁴⁵ First, T and T_R have the same observable consequences: $T \vdash p$ iff $T_R \vdash p$, where p is a sentence couched solely in the observational vocabulary and ‘ \vdash ’ is the second order deduction relation. Second, T_R is a logical consequence of T : $T \vdash T_R$ but not vice versa. Third, two or more RS’s with the same observational consequences cannot be inconsistent with each other. Fourth, first appearances notwithstanding, commitment to unobservables has not been renounced by moving from T to T_R . The RS quantifies over unobservable relations and hence does not eliminate ontological commitment to such relations; what the RS does eliminate is ontological commitment to *specific* unobservable relations.⁴⁶

Many proponents of RS claim that it captures a theory’s full cognitive content; see, for instance, Worrall (2007, 147) and Maxwell (1971, 16). For an ESRist this means that it captures a theory’s empirical and structural content. How plausible is this view? Since, as we have seen above, T_R and T have the same observational consequences, the association of T_R with T ’s *empirical* content seems plausible. But the association of T_R with T ’s *structural* content needs further elucidation. Above we have introduced structures as set theoretical entities, but the RS is a sentence in second order logic. How do the two connect? A plausible first stab at the problem is to say that T_R expresses T ’s structural content in the sense that T_R says that there is a structure S which satisfies T , and if T_R is true then S is instantiated. In that case, S is the structure we are supposed to be realists about in the sense that we believe S to be instantiated in (or to be isomorphic to the structure of) the part of the world described by T .

This is in need of qualification, and, as we shall see, these qualifications lead us into serious difficulties. To explicate the notion that a structure S satisfies T_R we have to introduce a semantics for second order languages such that it incorporates the distinction between observation and non-observation predicates. This requires a slight elaboration of the structures introduced in Section 2. Following Ketland (2004), we introduce a so-called Henkin structure $S = \langle (U_o, U_u), (R_o, R_m, R_u) \rangle$, which we obtain from the ‘standard’ set theoretic structures by bifurcating the domain into observable and unobservable objects (in technical terms, we use a two-sorted domain) and trifurcating the relations into observable, unobservable and mixed.⁴⁷ U_o is the domain of observable objects and U_u the domain of unobservable objects (hence $U = U_o \cup U_u$ is the total domain of S and $U_o \cap U_u = \emptyset$). R_o is the set of relations between observable objects (i.e. relations that hold between elements of U_o); R_u is the set of relations between unobservable objects (i.e. relations that hold between elements of U_u); and R_m is the set of ‘mixed’ relations, ones that hold between observable and unobservable objects (of course, only n -adic relations for $n \geq 2$ can be mixed). Notice

⁴⁵ See English (1973, 458–462), Ketland (2004, 293–294) and Psillos (2006b, Sec. 4).

⁴⁶ Having said this, one can give different interpretations to what the existentially bound variables range over. Carnap thought they ranged over mathematical entities. For more on this see Friedman (2011).

⁴⁷ We omit some technical subtleties here. For details see Ketland (2004). In particular, his definition of a Henkin structure also involves some collection *Rel* of classes and relations on the total domain of S satisfying the comprehension scheme. In what follows we assume that *Rel* contains all relations, i.e. S is ‘full’, which allows us to omit *Rel*. The results reached are then valid for full models.

that this way of introducing observable and unobservable relations is somewhat artificial because it only takes into account the kind of objects between which the relations hold but not the character of the relation itself; we come back to this issue in the next subsection. $S_e = \langle U_o, R_o \rangle$ is the *empirical reduct* (or empirical substructure) of S .

Now consider a two-sorted second order language $L_2(O, M, U)$, where O, M and U are sets of predicates referring to relations in R_o, R_m and R_u respectively.⁴⁸ This language contains two types of individual variables, one ranging over U_o and the other ranging over U_u , and three types of predicate variables: observation predicate variables ranging over R_o , mixed predicate variables ranging over R_m , and non-observation predicate variables ranging over R_u . Then, a second order sentence φ is true in S iff $S \models \varphi$. We now assume that theory T is formulated in $L_2(O, M, U)$. Then there is a question of how to construct the Ramsey sentence, because it is not clear whether or not to Ramsify the mixed relations. Given that mixed relations have an unobservable aspect, which SRists regard as suspect, it seems to be in the spirit of the approach to Ramsify mixed predicates; we come back to this below. Hence, in the language introduced above we have $P = M \cup U$ and $Q = O$.

Some further definitions are needed. Let $S_w = \langle (W_o, W_u), (V_o, V_m, V_u) \rangle$ be the structure of the target domain of T . Then, a structure $S = \langle (U_o, U_u), (R_o, R_m, R_u) \rangle$ is (a) *ontologically correct* iff it is isomorphic to S_w ; (b) *u-cardinality correct* iff U_u and W_u have the same cardinality; (c) *empirically correct* iff its empirical reduct is isomorphic to the empirical reduct of the target domain of T , i.e. iff $\langle U_o, R_o \rangle$ is isomorphic to $\langle W_o, V_o \rangle$. A second order sentence ε is *empirically adequate* iff it has an empirically correct model.

Now assume that T_R is not only empirically correct, but also true; i.e. we assume that there is some structure S such that $S \models T_R$.⁴⁹ This assumption is what distinguishes the SRist from both the instrumentalist, who only believes in the empirical correctness of T_R (or T for that matter), and the realist who believes in the truth of T (rather than only T_R). The crucial question now is: what follows from the truth of T_R about the structure S satisfying T ? The following theorem (which, for want of a better term, we dub ‘Cardinality Theorem’) gives a sobering answer to this question.

Cardinality Theorem (CT): The RS of theory T is true if, and only if, T has a model S (i.e. $S \models T$) which is *u-cardinality correct* and *empirically correct*.

Let us make the implications of this theorem explicit. The left-to-right implication says that all that we can infer from the truth of T_R about the unobservable world is a claim about its cardinality. It is worth pointing out, however, that CT need not determine the exact cardinality of the unobservable world. T need not be categorical and hence can have models of different cardinalities. In this case all we can say is

⁴⁸ This marks a slight departure from the way RS has been introduced above, where we use a one-sorted language (i.e. one with only one type of variable). However, this shift is purely technical in nature and does not affect the main ideas behind the Ramsey sentence. For a discussion of the same issues using a one-sorted logic see Ketland (2009).

⁴⁹ A similar argument can be given on the supposition that T_R is *approximately* true in a suitably qualified sense.

that W_u has the cardinality of *one of* the models of T . And this is the default option since categoricity is the exception rather than the rule. So the situation is in fact even worse than usually envisaged.

The right-to-left implication says that no matter what T claims about unobservable objects, the RS comes out true as long as T is empirically adequate and has a model of the right cardinality; in other words, any claim the RS may make about the existence of unobservable relations or their formal properties is automatically true (or ‘trivially’ true, as the point is often put). Hence, all that is possibly open to discovery concerning the unobservable domain is its cardinality.⁵⁰ This is an undesirable result and has been seen by many as a *modus tollens* against ESR.

One must guard against a common misconstrual of CT. In informal presentations of the material, CT is often paraphrased as the claim that if (i) what an RS says about observables is true and (ii) we quantify over a domain that has the right cardinality, then the RS is a truth of logic (or, in other words, the RS is trivially satisfied). This is wrong. Even if these two conditions are satisfied, the RS can still be false. The RS is guaranteed to be a truth of logic only if the theory *itself* has a model that is empirically correct and whose theoretical domain has the right cardinality, which is stronger than (a) and (b). Examples of cases where (a) and (b) are satisfied and yet the RS of the theory is false are given in Ketland (2009).

The first to arrive at a result of this kind were Demopoulos and Friedman (1985); a formal proof can be found in Ketland (2004).⁵¹ The proof crucially involves a result now known as ‘Newman’s theorem’. In 1928 Max Newman published a critical notice of Russell’s *The Analysis of Matter*, arguing that the ESR claim that we can know only the (abstract) structure of the external world trivialises scientific knowledge. The cornerstone of his argument is the realisation that ‘[a]ny collection of things can be organised so as to have structure W , provided there are the right number of them’ (Newman 1928, 144, cf. 140), where W is an arbitrary structure. More precisely:

Newman’s Theorem (NT): Let C be a collection of objects and let S be a structure whose domain has the same cardinality as C . Then there exists a structure C_S whose domain is C and which is isomorphic to S .

Intuitively the theorem says that given *any* structure, if collection C has the same cardinality as that structure, then there is system of relations definable over the

⁵⁰ As Demopoulos and Friedman (1985, 633–635) point out, this problem parallels Putnam’s (1978) so-called model theoretic argument.

⁵¹ For further discussions see Demopoulos (2003, 2008). CT is based on a model theoretic notion of empirical adequacy, which is essentially van Fraassen’s (1980, 12). Demopoulos and Friedman seem to phrase their argument in terms of a different notion of empirical adequacy, namely that a theory is empirically adequate iff all consequences of the theory which are couched in a purely observational language are true (1985, 635). As Ketland (2004, 295–6) points out, the latter notion of empirical adequacy is strictly weaker than the former in that the former implies the latter but not vice versa; for this reason he calls the latter ‘weak empirical adequacy’. CT does not hold if empirical adequacy is replaced by weak empirical adequacy. In fact, a theory may be *weakly* empirically adequate and have a u-cardinality correct model, and yet the theory’s RS may be false. Although this problem arises only when T has an infinite model, it is serious problem since most theories involving numbers have infinite models. (Thanks to Jeff Ketland for pointing this out to us.) Demopoulos (2008, 380) agrees, but argues that the model theoretic notion of empirical adequacy was what he and Friedman were referring to in their (1985).

members of C so that C has that structure. For a formal proof see Ketland (2004, 294–5), although the result should be intuitively obvious. From a set theoretic point of view, all that there is to a relation is its extension. So all we have to do in order to define a relation is to put objects in ordered tuples and put these tuples together in sets, which we can always do as long as we have enough objects. This, of course, would not be possible if we had an intensional interpretation of the relations, or required that the relations be in some sense ‘natural’ or ‘physical’.

Following common usage we refer to the challenge posed to ESR by CT or by NT as the ‘Newman Problem’.⁵²

3.4 Getting around the Newman problem

Attempts to get around the Newman Problem fall into two broad categories. Those in the first category hold on to RS (or at least a modified version of it) and deny one of the premises that go into the proof of CT or NT; those in the second category take CT or NT to show that RS is not the right instrument to explicate a theory’s structural content.⁵³ In what follows we consider variants of each of these attempts. It is not our aim here to endorse one solution but rather to reveal the diversity of solutions and the difficulties each faces.

3.4.1 Salvaging the Ramsey sentence

(a) *The construction of RS.* As we already briefly indicated, the above framework of classifying relations into observable, unobservable, and mixed *according to the objects to which they apply* is artificial and has counterintuitive consequences. For instance, on this view there is no single relation ‘larger than’; there are three relations: ‘observably larger than’, ‘unobservably larger than’, and ‘miscellaneously larger than’, ranging exclusively over observables, unobservables and the full domain respectively. Cruse (2005, 564–565) criticises the gerrymandering of relations as unmotivated and urges that we should draw the distinction according to the epistemic status of the relation *itself*: there are relations we can observe (like ‘larger than’), at least in certain circumstances, and ones we cannot observe (like ‘has greater isospin than’). The crucial point is that nothing precludes observable relations from applying to unobservable objects; we can, for instance, meaningfully say that a sulphur hexafluoride (SF_6) molecule has the shape of an octahedron. For this reason, the right way to Ramsify a theory is to divide the relations into observable and unobservable ones along those lines, and then Ramsify the non-observation predicates. Since this leaves standing certain relations that apply to unobservables unramseyfied (i.e. they are not replaced by a predicate variable and

⁵² A note about terminology is in order. The Cardinality theorem is often referred to as ‘Newman’s theorem’, ‘Newman’s result’ or ‘Newman’s objection’. At least from a historical point of view this is incorrect. Newman’s theorem (both as stated here and as presented by Newman himself) does not involve RS, nor does Russell’s (1927) theory against which Newman’s original argument is directed. We turn to the Russell-Newman discussion below.

⁵³ For a detailed discussion of these replies see Ainsworth (2009).

existentially quantified), CT no longer goes through (the proof crucially assumes that observational predicates only apply to observable things).

There is of course a question whether one can draw a meaningful distinction between observable and unobservable relations in this way. Cruse (2005, 567) suggests that one way of doing so that is pertinent to science is to appeal to Charkravarty's (2004) distinction between auxiliary and detection properties. Whether or not this is satisfactory is a question that deserves further discussion; in fact, Cruse himself raises some doubts (2005, 569–771). But even if we answer this question affirmatively, there is a further serious problem. Ainsworth (2008, 138–140) shows that one can prove a theorem very similar to CT: T_R is true iff T has a model S in which all observation predicates have the right extension and which gets the cardinality of the domain of all objects that do not instantiate an observable relation right. This theorem implies that a theory does in fact tell us something about unobservable objects, namely that they instantiate certain observable relations; in this sense we are in fact in a better situation. However, it also implies that we have no knowledge about unobservable relations, either structural or non-structural.

Melia and Saatsi's response departs from the observation that the relations postulated by scientific theories stand in various modal relations to each other: some are counterfactually dependent and some are correlated in a law-like manner (2006, 579). They identify as the source of the problem the inability of extensional second order logic—the logic in which the RS is formulated—to express such relations, and suggest that we augment the language with a number of modal operators expressing modal relations like 'it is physically necessary that' and 'lawfully correlated with' (2006, 281). Scientific theories and their RSs should then be expressed in this stronger language, which no longer allows for a proof of CT. This is an interesting proposal but it raises issues that need further attention. The proposal seems to require the treatment of the extra modal operators as logical primitives, and from what Melia and Saatsi say it is not clear how to motivate and carry out this move.

(b) *Real vs. fictional relations.* As Newman himself points out (1928, 145–147), NT is not a threat if we can distinguish the 'real' relations from those that are merely extensionally defined, which he terms 'fictional'.⁵⁴ Two noteworthy attempts have been made along such lines.

Redhead (2001b, 346) emphasises that unlike T_R , T does refer to specific relations. A specific relation r is 'hypothesised in some explanatory theoretical context so it exists as an ontological posit' and the referent of the relation symbol in T is 'picked out in non-structural terms'. Nevertheless he insists that 'all that we have epistemic warrant for is the second-order structure' and not r in itself (*ibid.*). However, it is not clear in what sense this position still is ESR. On the one hand, if we have the means to pick out referents for a specific hypothesised relation, then we are fully-fledged realists. On the other hand, if we have to remain agnostic about r ,

⁵⁴ After some discussion he concludes that the only way to draw the distinction between real and fictional relations is to align it with the distinction between important and trivial relations. This, however, commits one to admitting importance as an unanalysable primitive, something which, as Newman points out, is absurd.

then it is difficult to see where the epistemic warrant for the second-order structure of r comes from (Psillos 2001b, 369). So it seems that this proposal collapses either into full-fledged realism or scepticism.

In his reply to Putnam's model theoretic argument against metaphysical realism Lewis (1983) puts forward the view that it is a contingent fact which relations are more important—and ultimately natural—than others. Then, so the argument continues, only natural relations should be taken into account when pondering the structure of the world; we need not, strictly speaking, deny that the world instantiates (as NT has it) any relation compatible with its cardinality, but we submit that only natural relations are taken into account when it comes to assessing the claims of a theory (and, in particular, the domain of quantification in RS is to be restricted to natural relations). Melia and Saatsi (2006, 575–576) point out that this suggestion, if sound, would double as a solution to the Newman Problem, but they are quick to point out that it suffers from too many problems of its own to be convincing (and they reach similar conclusions for the other suggestions they consider, namely that relations be intrinsic, qualitative, contingent, or causal). For one, this is just the natural kinds view that realists proffer, and as such it suffers from all the problems that attach to this view, most notably that our views of what counts as a natural kind may change radically over the course of history. Furthermore, there is at least a question about whether one can distinguish between natural and non-natural kinds without appealing to non-structural knowledge of the world. If this is so, as Psillos (1999, 66) claims, then this amounts to a defeat of ESR.

(c) *Reinterpreting the observational content of a theory.* Zahar's reply to Newman's Problem departs from a reinterpretation of CT,⁵⁵ which he takes to be saying that 'only what the Ramsey-sentence asserts *over and above its observational content* is reducible to logic or mathematics' (2001, 239–240, original emphasis; cf. 2004, 10). He then identifies the observational content of a theory with the empirically decidable 'singular atomic sentences' describing observable events that support the theory (2001, 240–241); i.e. sentences like 'this particular piece of copper has a density higher than water'.⁵⁶ He then points out that the RS implies empirical generalisations (i.e. '*all* copper has a density higher than water'), which are stronger than singular sentences. Therefore, he concludes, the RS has content over and above its observational content; yet, this content does not reduce to mathematics or logic, and hence the Newman objection, as he construes it, is wrong.

This response is unconvincing because it equivocates on 'over and above'. On his understanding simple empirical generalisations go beyond the observational content

⁵⁵ Zahar (2001, Appendix IV) is written jointly with Worrall. However, it becomes clear from Zahar (2004) that this response is attributable only to Zahar; we discuss Worrall's response below. A great deal of Zahar's (2001, 2004) and also Worrall's (2007) discussion of RS is concerned with a version of Newman's Problem that interprets the RS as existentially quantifying not only over theoretical but also over observable terms (Zahar 2001, 238). They are, of course, right that this is absurd. However, CT as stated above does not dispense with observational predicates, and so we discuss their arguments only in as far as they address the version of Newman's problem under discussion here.

⁵⁶ The example is Zahar's. Some may object that 'density' is a theoretical term. How this issue is resolved is inconsequential for the current discussion.

of a theory. Even some instrumentalists, however, agree that we can believe in the truth of empirical generalisations; so the difference between singular sentences and empirical generalisations is not what is at stake in debates over ESR. What this debate is about is the question of whether RS asserts anything interesting about the domain of things that are not observable at all (protons, superstrings, etc.); so what the ESRist has to show is that RS goes beyond empirical generalisations! CT shows that RS does so only in a very minimal sense (by implying a claim about the cardinality of the domain of unobservables), and this result is left untouched by quibbles about singular versus general observational statements.

(d) *Renounce referential semantics*. Worrall (2007, 152–4) argues that the entire discussion rests on a faulty premise, namely that we have theory-independent knowledge about unobservable entities or natural kinds (i.e. that we can somehow identify natural kinds without relying on the theories that postulate them). Indeed Worrall goes one step further and argues that for this reason we should give up referential semantics. So far the Newman problem has been discussed in terms of whether quantified variables in the RS come to refer to uniquely determined elements of reality and whether the sentence picks out a uniquely determined structure which can be regarded as the structure of the world. This, thinks Worrall, is starting on the wrong foot. Since, in his view, we have no out-of-theory experience that would allow us to compare theory and reality we should give up referential semantics. Provided we do that, the problems (along with the purported solutions) discussed in this section are simply much ado about nothing.

This proposal raises two questions. The first and obvious question is what should replace referential semantics. How are we to understand a theory's relation to the world if not in terms of reference? Worrall only touches very briefly upon this issue and suggests that usual term-by-term mapping (presupposed by standard referential semantics) should be replaced by the notion that 'the mathematical structure of a theory may globally reflect reality without each of its components necessarily referring to a separate item of that reality' (2007, 154). Much seems to hang on the notion 'globally reflect'. Alas, Worrall remains silent about the details of this notion so we cannot properly evaluate his proposal. Let us just say here that we find it difficult, if not impossible, to conceive of how one could have global structure reflection without also having term-by-term correspondence on some level. Moreover, it is not clear from what Worrall says whether his rejection of referential semantics concerns only theoretical terms. If that is indeed the case, the called-for semantics becomes even more mysterious as it now needs to take some sort of hybrid form.

The second question is whether one can abandon referential semantics and still be a realist. It is commonly taken to be an essential part of the definition of scientific realism that the theoretical terms of a mature scientific theory usually refer (see, for instance, Psillos 1999, Ch. 12), and so it seems that by giving up reference one is forced to give up realism. Worrall admits that on this conception of realism his position does not qualify as realism (2007, 152–3), but then goes on to reject this conception by denying that reference is an essential part of a definition of realism. Instead he offers two alternative conceptions of realism. We have already seen the first and the problem it faces—this is the appeal to a theory's structure globally

reflecting reality. The second suggestion (made in personal communication) submits that the retention of a theory's RS through theory change is what it *means* for a theory to be (approximately) true—and to be clear, the proposal is not that retention is a good indicator of truth in the sense of correspondence; the proposal is that retention *is* (approximate) truth. This, in our view, is to conflate the reasons in support of a view with the view itself. What lies at the heart of scientific realism is the idea that a theory presents a correct image of how things in the world are; simply surviving the test of time, desirable though it may be, is neither necessary nor sufficient for realism.⁵⁷

3.4.2 Abandoning the Ramsey sentence

So far we assumed that theories are sentences formulated in a second order language. As a consequence SR was construed as a doctrine about the reference of terms, which led to the problems with RS. This has suggested to some that the crucial mistake happened right at the beginning. By construing theories as sentences we have embedded our approach to ESR in the so-called syntactic view of theories, at the heart of which lies the notion that theories are partially interpreted sentences expressed in a formal calculus. This view has been criticised as inadequate as early as the 1960s and has since been given up by most writers on the subject (see Suppe (1977) for a review of this debate). The most widely-held alternative is the so-called semantic view of theories, which posits that theories are families of models. This view comes in various versions which differ in what they regard as models (for a survey see Suppe (1998)). The versions which are relevant to our problem are those that take models to be structures in the above sense.⁵⁸ These structures are then connected to the world by isomorphism, or a generalisation of it.⁵⁹

French and Ladyman suggest that once we abandon the syntactic for the semantic view the Newman Problem no longer arises because this problem is an artefact of the syntactic framework and in particular of its construal of relations in first-order extensional terms (2003a, 33).⁶⁰ At no point, however, do French and Ladyman really elaborate how the Newman Problem can be obviated with the adoption of the semantic approach. Indeed, if, as some authors have argued (e.g. Worrall 1985), the syntactic and semantic formulations of scientific theories are inter-translatable then there is no reason to suppose that a switch from one formulation to the other will make the Newman problem vanish.⁶¹

⁵⁷ See Votsis (2011a) for details why survival through theory change is neither necessary nor sufficient for approximate truth.

⁵⁸ Structuralist versions of the semantic view of theories are put forward, among others, by Da Costa and French (1990, 249) and van Fraassen (1980, 43, 64; 1991, 483; 1995, 6; 1997, 522, 528–9).

⁵⁹ Candidates are, among others, partial isomorphism (French and Ladyman 1999), embedding (Redhead 2001a), and homomorphism (Mundy 1986). For the issues under discussion nothing depends on which of these one chooses.

⁶⁰ As we will see below, French and Ladyman endorse an ontic version of structural realism. What is worth noting here is that this endorsement does not necessitate a switch to the semantic view of theories. Having said that, we know of no ontic structural realist who advocates the syntactic view of theories.

⁶¹ For an argument against the idea that by abandoning extensionalism we can avoid the Newman Problem see Demopoulos and Friedman (1985, 629–30).

Ainsworth (2009, 150–152) argues that French and Ladyman’s suggestion is wrong for two reasons. First, Newman’s original formulation of NT makes no use of a linguistic construal of theories and is in fact based on the view that scientific theories directly specify a structure that represents the structure of the target system. So it is not clear how the move from the syntactic to the semantic view should alleviate the problem. Second, and more importantly, upon closer examination the semantic view seems to run into the same difficulties.⁶² ESR in the semantic version is committed to the claim that the family of structures which makes up the theory contains at least one structure, S , which is empirically correct (in the sense specified above) and isomorphic to the structure in the world. Ainsworth proves that this is equivalent to the following claim: the most we can know about the world is that there is some structure (provided to us by the theory) which is empirically correct and T -cardinality correct (in the above sense). But this is what the theory’s RS affirms, and so we are back to where we started.

There is one structuralist version of the semantic view that does not fall prey to this objection, namely van Fraassen’s (*op. cit.*). But the rescue comes at a price: the abandonment of realism. Van Fraassen’s constructive empiricism posits that science does not aim for the truth; rather it aims at devising empirically adequate theories—i.e. ones that account for all observable phenomena while remaining agnostic about everything that escapes observation. This makes Newman’s problem irrelevant. The same is true of Otávio Bueno’s position (1997; 1999; 2000), which he calls ‘structural empiricism’ and explicitly develops as an extension of van Fraassen’s. With van Fraassen’s position it shares an emphasis on empirical adequacy, which he explains in terms of partial structures. Hence it also avoids the Newman problem, but at the (intentional) cost of abandoning realism altogether.

Russell’s own reaction to the Newman problem involves a climb-down. In a letter sent to Newman shortly after the publication of his critical notice, Russell (1968, 176) acknowledges the consequences of the objection but argues that despite his own pronouncements he had in effect been advocating an impure form of ESR—he cites the assumption that percepts are spatiotemporally continuous with their causes as one such impurity.⁶³ In his subsequent philosophy, for example in his (1948), Russell continued pursuing a structuralist agenda without ever returning to address NT.

Votsis (2003, 886–889) argues that although NT, as reconstructed by Demopoulos and Friedman, may be a serious threat to the RS version of ESR, Russell’s original version has the resources to counter it. The crucial point is that Russell does not use the RS to pick out a structure (in fact, although modern presentations often portray Russell as holding a RS view, neither Russell nor Newman make use of anything like the RS in their arguments). This, Votsis argues, is important because the problem with RS methods is that they pick out a structure ‘indirectly’ as the thing that satisfies a relevant theory; as we saw earlier CT shows that this procedure is too

⁶² A similar argument is put forward by Ketland (2004, 295), and Chakravartty (2001) makes a related point.

⁶³ For a discussion see Votsis (2011b).

permissive. By contrast, Russell asks us to abstract the structure directly from our perceptions. Now, in every situation we experience certain relations but not others. Since in Russell's philosophy specific relations between perceptions are the causal consequences of specific relations between external stimuli, the two being at best isomorphic, the structure we end up with is as much about the one set of relations as it is about the other. Naturally, the chosen structure is satisfied by many other relations. That in itself is not a problem. What matters, stresses Votsis, is that the attribution of a specific structure to a specific physical system is causally and empirically grounded. That is, it is not a matter of saying that *some* relations satisfy the given structure as proponents of the RS do. Rather, it is a matter of pointing out that specific relations between external stimuli satisfy the given structure. These relations can be identified as the causal antecedents for the corresponding specific relations between perceptions that allowed us to abstract the structure in the first place.

This argument is of course only as strong as MR, though this is not contested in NT. A more pertinent problem is the assumption that percepts and their corresponding stimuli have a unique structure, which is at least implicitly assumed in Russell's argument. Frigg (2006) argues that a given object can be ascribed a number of non-isomorphic structures. This is meant to undermine ESR because to be a realist about a structure seems to presuppose that there is a unique structure we believe to be instantiated in the world. Indeed, this seems to be a problem for all realists. One way around this problem is to adopt a perspectivalist approach like Giere's (2006), even though Giere himself is not a structuralist. According to such a view, even though there is one world, different perspectives often yield different representations. So long as there is one true representation for each perspective (and hence many others that are false for *that* perspective), underdetermination will be far from rampant. The hope is that the resulting position amounts to a weak form of realism. This approach can be adopted by, and adapted for, the epistemic structural realist: so long as there is one true structure for each perspective, structural underdetermination will not be abundant. Of course these structures will not typically be true across different perspectives. Even so, the structures that are deemed false (within each perspective) can be collected in a large class that the perspectivalist can parade as the contrast class of what can be known. The question then remains whether this view is sufficiently strong to be dubbed 'realism'.

3.5 Structure versus nature

ESR is sometimes formulated as the claim that all we can know is structure while the nature of things remains hidden from us (see, for instance, the quote from Worrall in Section 3.2.1). Once formulated in this way, the tenability of ESR depends on the tenability of the nature/structure dichotomy, and a refutation of the position takes the form of an argument for the conclusion that nature and structure are either not separable or not neatly separable.

In our view the introduction of natures into the debate is a red herring; natures make no contribution to an adequate articulation of the position, and the vagueness of the concept invites confusion. In this subsection we briefly review some of the

arguments in the ‘nature versus structure’ debate and explain why they remain inconclusive.

Psillos offers the most detailed critique of the nature/structure dichotomy.⁶⁴ He first criticises Worrall for not being clear on ‘what exactly the distinction he wants to draw is’ (1999, 155).⁶⁵ This is true, but given what has been said about structures so far one natural suggestion is that the distinction is best cast as one between an extensional versus an intensional conception of theoretical predicates and of the unobservable properties those predicates are meant to correspond to. Thus understood, the distinction is precise and well-defined, and an answer to Psillos’ two major objections emerges. The first objection is that the nature and structure of an entity form a continuum and therefore cannot be separated, and the second one is that the nature of an entity is no less knowable than its structure (1995, 31–32; 1999, 155–157; 2001a, 20–21). In response to the first objection it can be pointed out that the intension and extension of a theoretical predicate do not form a continuum since there is a clear cut-off point between the two. In response to the second objection it can be pointed out that if the sort of reasons cited earlier are correct—e.g. that the intensions of theoretical predicates are radically overthrown while their extensions are largely preserved—the intension of a theoretical predicate is less knowable than its extension. Similar answers have been given by Votsis (2007), who draws the distinction as that between any description that can be specified up to isomorphism and any description that seeks to go beyond isomorphic specification. Irrespective of how this dispute is resolved the point to emphasise here is that talk of natures, without clarification of what it may mean, is misleading.

Worrall’s examples might suggest another understanding of natures: natures are those things that are thrown overboard in a scientific revolution, and structures are those things that prevail. Obviously this is not a useful way of drawing the distinction since it trivialises ESR (see, for example, Stanford 2007). However, we doubt that this is what Worrall, or any other proponent of ESR, really advocated, since Worrall emphasises that ESR, far from being trivially true, is a position that has to be tested against the history of science.

3.6 The third way on realism

Those who believe that ESR as formulated is rendered untenable by the objections discussed, and yet do not want to revert to conventional realism, may try to find a solution by blending ESR with other kinds of realism. Chakravartty’s *semirealism* aims to reach precisely such a middle ground by bringing together entity realism (ER) and ESR (1998; 2003; 2004; 2007); more specifically, Chakravartty claims that once we give a plausible remoulding of both SR and ER, they actually coincide in what he calls semirealism.

⁶⁴ Further discussion can be found in Ladyman (1998) and van Fraassen (2006); for a rejoinder to Psillos see Elsamahi (2005) and Votsis (2007).

⁶⁵ Stanford (2003a; 2003b) reaches a similar conclusion arguing from the point of view of confirmation theory.

Entity realism (see, e.g., Hacking (1983)) is the claim that while we should generally remain agnostic about the truth of theories, we have good reason to believe in the existence of the entities that occur in them: we should believe, for instance, in the existence of electrons and yet reject the claim that theories about electrons are even approximately true (as the traditional realist would believe). We have good reason to believe in the existence of a certain entity, according to the entity realist, when we can successfully manipulate them to bring about desired effects. The problem with this view is that it seems incoherent to say that we believe in electrons while rejecting theories about electrons, since it is only through these theories that we know what electrons are like and how we can manipulate them; if we reject these theories, electrons could be just about anything.⁶⁶ Chakravartty suggests solving this problem by distinguishing between *detection* properties and *auxiliary* properties. Detection properties are those properties that we detect through our measurement devices; i.e. properties in virtue of which the measurement results become manifest. Properties that do not play a causal role in observation are auxiliary properties. We are then invited to believe what theories say about detection properties, while remaining agnostic about the auxiliary ones.

This fix of entity realism also leads the way for a remoulding of ESR. Deviating from standard ESR, semirealism claims that what is preserved through theory change is not only what is contained in the RS—the theory’s empirical and structural content—but also the interpretation of those theoretical terms that refer to detection properties. Now, ESR *thus* understood is indeed indistinguishable from entity realism as construed above because both views basically advocate belief in what theories say about observables and detection properties, while suggesting agnosticism about claims that go beyond that.

This raises at least two questions. The first is whether semirealism really is, as Chakravartty seems to suggest at times, a kind of ESR. Semirealism claims that we have full knowledge of unobservable detection properties and that this is non-structural knowledge. At the same time we are told not to take seriously those parts of the theory that are not concerned with detection properties because these do not typically survive theory change; so those parts of a theory provide no knowledge at all. At least on the face of it, there does not seem to be a domain of which we have purely structural knowledge—indeed not even impure structural knowledge of the kind advocated by Russell—and so one wonders why this position should count as a kind of ESR. Of course, this objection can easily be evaded if one concedes that semirealism merely inherits a number of features from ESR without itself being a version of ESR, and there is in fact some indication to believe that Chakravartty favours this view.

The second and more pressing worry concerns the question of how we identify detection properties. Chakravartty’s answer is that we turn to the theory’s equations and ask what they demand as an interpretation (1998, 396). What Chakravartty has in mind seems to be interpretation relative to a certain causal context. The context fixes a minimal interpretation, and then the maxim is to renounce add-ons. To be

⁶⁶ Despite their professed aversion towards theory, entity realists make allowances for some, low-level, theory. Hacking, for example, appeals to ‘low-level causal properties’, which, no matter how much glazing he puts on them, are simply theoretical properties.

precise, interpretations that are either only indirectly connected or completely unconnected to detection practices are eliminated in favour of those that are so connected. Discussing the case of light he warns that ‘[t]o suppose that a direction of propagation is furthermore a direction in the ether is to go beyond what is minimally required to give an interpretation of this particular set of [i.e. Fresnel’s] equations’ (2007, 53). He goes on to say that the existence of unobservable properties that are first-order⁶⁷ and intrinsic like the intensities of the incident, reflected and refracted beams gives the requisite minimal interpretation to Fresnel’s equations. The problem with this suggestion is that Chakravartty appeals to precisely the same properties that other structural realists employ to support their own view; e.g. Worrall (1989) can be read as broadly construing these as observational properties. Even if these properties are best construed as unobservable, as Chakravartty insists, it is not clear why we should also regard them as first-order and intrinsic. In sum, both the epistemic structural realists and the semi-realists need to motivate why their view is better positioned to account for the success enjoyed by theories.

4 Ontic structural realism

Structural realism took a different turn at the end of the 1990s when James Ladyman (1998) and Steven French (1998) suggested construing SR as an ontological rather than an epistemological position. In Ladyman’s words,

‘we should seek to elaborate structural realism in such a way that it can diffuse problems of traditional realism [...] This means taking structure to be primitive and ontologically subsistent.’ (Ladyman 1998, 420)

This position has been dubbed *ontological structural realism* (OSR).⁶⁸ Its core is the belief that structures are ontologically basic: what originally appeared as the structure of something with unknown qualitative features is actually all that there is to nature. In other words, OSR denies the presumption that objects and their qualitative features are ontologically basic and that structures, so far as they exist, depend on the existence of these objects and their properties. OSR reverses this order of dependence and sees structures as ontologically basic, and posits that in so far as objects and properties exist, their existence depends on the existence of a structure. OSR agrees with ESR that we can have knowledge only of structural aspects of the world, but it sees the reason for this in the fact that structure is all that there is or at least that it is all that is ontologically basic and that therefore there is nothing else to know.

This central idea has been developed in different ways, and there is controversy over the correct formulation of the position. In Section 4.1 we discuss different versions of OSR, and in Section 4.2 we turn to arguments that have been offered in its support.

⁶⁷ In this context ‘first-order’ means that a property is not a property of another property but of an object. For instance ‘red’ is first order; ‘being a colour’ is second order.

⁶⁸ A similar position has been suggested by Dipert (1997), but he does not use the term ‘structural realism’.

4.1 Different versions of OSR

Radical OSR (ROSR) takes the above doctrine at face value and asserts that, literally, structure is all that there is, and we should not be fooled into believing that structures have, or even need, ‘carriers’ like objects. The world is one large structure: there is nothing that is *structured*.

Although some statements of OSR *prima facie* seem to come close to asserting ROSR, and they certainly have been understood by some commentators in that way, proponents of OSR do not seem to entertain this position seriously. And this better be so. In the eyes of most if not all commentators ROSR is downright unintelligible. Recall what has been said about structures in Section 2: they consist of dummy objects and relations which are defined as sets of ordered tuples of such dummy objects. It just does not seem to make sense to claim that the concrete physical world literally *is* nothing but a structure *thus defined*.

A more sophisticated reading of ROSR replaces identity by supervenience: the physical world supervenes on a structure in the same way in which the mind is claimed to supervene on the brain. But this position runs up against the problem that structures do not seem to be sufficient to fix non-structural properties—as is required by supervenience—because different non-identical systems can have the same structure. This has been noted in a different context by Mary Hesse (1963) who points out that different material systems can instantiate the same formal relations and then refers to these systems as standing in *formal analogy* to each other. For instance, the elongation of a lead ball bouncing up and down on a metal spring and the voltage over a condenser obey the same equation and, in that sense, have the same formal structure even though the systems are materially very different.

The proponents of ROSR could respond that this is an artefact of us focussing only on very small parts of the world, and that the problem would go away if we focussed on the structure of the entire universe. But why should it? Even if we consider a structure vastly larger than a laboratory system, it is still true that different combinations of non-structural properties can have the same structure. The same holds for the entire universe, though the reasoning in that case would be counterfactual, e.g. different combinations of non-structural properties could have the same structure as the entire universe. Assuming that current physics is the correct theory of the world, physical symmetries provide us with relevant scenarios. Take the CPT theorem, which (roughly) says that every Lorentz invariant quantum field theory is CPT invariant, meaning that if you change the direction of time as well as charge and parity of every entity in the world, then the equations of the theory are preserved. Since—on the doctrine at stake—equations provide us with structure, the structure is preserved under a CPT transformation. But the world is a different one: what was positively charged is now negatively charged. So we have a qualitatively different world with the same structure.

There does not seem to be a way around the fact that structures (and with them structural properties) are abstract with respect to physical properties and relations, and that physical properties therefore are neither identical with nor supervenient on structural properties. Essentially this point has been made by Busch (2003, 220), Cao (2003b, 59), Frigg (2006, 55), Psillos (2006a, 362–65) and Simons (2002, 38). For this reason ROSR seems to be a non-starter.

Eliminative OSR (EOSR) is the position of French (1998; 1999; 2003) and Ladyman (1998; 2001), and French and Ladyman (2003a; 2003b); the term ‘eliminative OSR’ is due to Psillos (2001a). EOSR holds that relations, not objects, are ontologically fundamental. This commits EOSR to the view that relations do not need relata between which they hold: relations exist, and can be ontologically fundamental, without there being relata. For this reason objects become otiose, and are eliminated from the basic ontology: ‘objects’ are only places in a relational structure (in as far as they are something at all) and should not be taken ontologically seriously. This can mean two things. The most common version of EOSR regards independently existing individual objects as nothing more than convenient fictions and denies their existence. An attenuated version of EOSR acknowledges the existence of objects, but denies that they are individuals (we say more about this position below).

The essential difference between ROSR and EOSR lies in their treatment of relations. While ROSR insists on an extensional treatment of relations that is standard in mathematical logic, EOSR allows for relations that have intensions; that is, fundamental relations can be relations like ‘being larger than’ rather than only ‘standing in a transitive relation’. In other words, relations are interpreted.⁶⁹ This is how EOSR avoids the criticisms that have been levelled against ROSR and accounts for the ‘physical content’ of the fundamental structure. While in the light of the problems of ROSR this is certainly the right move, it is worth emphasising that this marks a radical departure from the original ideas of structural realism, which were squarely based on an extensional understanding of structures.⁷⁰

The main criticism of EOSR (one that also applies to ROSR) is that one cannot have relations without relata (Busch 2003, 213; Cao 2003b; Chakravartty 1998, 399; Lyre 2004, 667; Psillos 2001a): relations cannot be real unless there are things that are related. In Busch’s words, ‘the very idea of structure presupposes some elements that go together to make up that structure. A relation might take *anything* as its relata, but it always takes *something*.’ (*ibid.*, original emphasis) Hence EOSR is simply incoherent.

Proponents of EOSR can respond to this objection in three ways. The first response, due to Ladyman and Ross (2007), points out that most relations do in fact have relata, but upon closer analysis these relata turn out to be structures themselves. What appears to be an object standing in certain relations to other objects is in fact, at a more basic level, itself a structure. The problem with this response is that it obviously gets us into an infinite regress. If we stop at some level, then the incoherence objection can be made again: the basic structure we have identified has to be built up of objects, and so there are not only relations after all. To undercut this objection proponents of EOSR have to respond that this argument relies on a premise that they deny, namely that there is a fundamental level. Every time we think we have identified an object, we just need to look closer and then the object

⁶⁹ This is explicit, for instance, in Ladyman and Ross (2007, 128).

⁷⁰ In the terms introduced in Section 2 one could say that EOSR uses a concrete rather than an abstract structure as far as relations are concerned. There is a question about whether such structures are allowed to contain properties, i.e. monadic relations, or whether all relations have to be polyadic. Esfeld (2004), and Esfeld and Lam (2008) seem to suggest that monadic properties should be ruled out. For a discussion of this issue see Ainsworth (2008, 162–9).

turns out to be a structure, and so on ad infinitum. The world is structure all the way down. This stands in stark contrast to the pervasive intuition that, if there is an ontological hierarchy of levels in nature, then there is a fundamental level. There is therefore at least a question about whether a conception of reality with no fundamental level can be given.⁷¹

The second response is that the ‘no relations without relata’ objection is fuelled by an extensional understanding of relations, on which the idea of relations without relata is indeed incoherent. However, once we take into account that EOSR is not committed to an extensional understanding of relations, the conceptual revisions needed to make the proposal fly are less dramatic than it first appears. In fact, once we adopt a substantial understanding of relations, EOSR becomes almost equivalent to the so-called bundle theory of objects (Morganti 2004, 97).⁷² The problems of EOSR then become variants of the problems of the bundle theory (see Armstrong 1989 for a discussion of these); in particular, the ‘no relations without relata’ problem becomes a variant of the issues surrounding non-instantiated properties and the question of how a bundle can be formed without there being a substratum in which properties are instantiated. While these are serious metaphysical issues, they arguably don’t add up to an incoherency. What is required to make EOSR palatable is a cogent response to the objections faced by the bundle theory of objects.⁷³

The third response, drawn from French and Krause (1995; 2006, Ch. 7), eschews commitment to individuals without renouncing objects. They do so by employing so-called quasi-set theory, a non-classical set theory at the heart of which lies the notion that there are two kinds of atoms (objects) rather than only one (as in standard set theory). The atoms of the first kind are like the atoms of standard set theory and represent individual objects. The difference lies in the second kind of atoms, which are objects but not individuals in that identity statements simply cannot be formed of them (statements like ‘ $x=y$ ’ are not well formed when x and y are taken to denote atoms of the second kind). There is then an indistinguishability relation that can hold between pairs of atoms of the second kind. It follows that atoms of the second kind can be indistinguishable without being identical.

This approach qualifies EOSR in an important respect: in French and Krause’s framework there are objects, but these need not be individuals: some are, others are not. We refer to this position as *attenuated EOSR* to set it apart from unqualified EOSR, which denies the existence of objects *tout court*, rather than only denying that objects are individuals.

Attenuated EOSR marks a less radical break with traditional metaphysics than other versions of OSR also in another respect. Even though the framework makes room for non-individual objects, it also contains individuals and it does not prejudge which parts of reality are best described as containing objects of one kind or the

⁷¹ For a general discussion of this point (not geared towards OSR) see Schaffer (2003).

⁷² The two are ‘almost’ equivalent because traditionally bundle theories focus on monadic properties, while OSR places the emphasis on polyadic properties, i.e. proper relations. However, one can have a bundle theory that rests only on polyadic properties, and that most bundle theorists don’t discuss this possibility is an accident of history and not indicative of an intrinsic limitation of that approach. For a metaphysics of relations see Mertz (2003).

⁷³ For a discussion of the bundle theory with a special focus on tropes see French (2001, 21–22).

other. French and Krause make it clear that their focus is on QM and they leave it open whether or not objects outside the quantum realm are individuals.

Before turning to other versions of OSR, we would like to mention briefly a further family of criticisms that has been levelled at EOSR. Individual objects, so the argument goes, play a crucial role in theories of causation, and in theories of change more generally.⁷⁴ But since unqualified EOSR does not recognise objects, it cannot account for such a notion of causation (Chakravartty 2003; Psillos 2001a; b). French (2006b) discusses these objections and points out that they rely on specific views of causation and change that, first, are not uncontroversial, and, second, the OSRist need not accept. One suggestion that has recently come out of the OSR camp is that deflated accounts of causation may be more suited for the bare-bones metaphysics endorsed by them. After all, as they often complain, full-bodied metaphysical accounts of causation are out of place in the context of modern physics.

Other versions of OSR take a more reconciliatory stance on objects than EOSR. Esfeld (2004) argues for the view that objects are individuals but that they do not have any intrinsic properties; he dubs this view ‘moderate OSR’. On this view objects are like Lockean substances: they are featureless substrata whose only function is to instantiate relations of a structure. Whether one can make sense of the notion of something that is not anything intrinsically is a time-honoured philosophical problem; for a survey see Robinson (2004).

The challenge for moderate OSR is to explain what confers individuality upon these objects. Appealing to intrinsic properties is obviously not an option. Another suggestion would be to say that they possess primitive thisness or haecceity. However, this suggestion is too close to a traditional, non-structuralist ontology to be attractive to moderate OSR. In keeping with the spirit of OSR, proponents of moderate OSR have to argue that the identity of objects is conferred upon them by the structure of which they are part. This runs counter to the usual conception of individuality, which sees individuality as intrinsic. Moderate OSRists need to show that there is a workable conception of non-intrinsic individuality, which is, as we see below when discussing quantum mechanics, not a straightforward matter. Esfeld and Lam (2008, 32–34) assert that relations are sufficient to formulate identity conditions, but do not provide an explicit formulation of such conditions. For a discussion of this issue see also Ladyman (2005, 2007).

Lyre (2011) has argued for a view diametrically opposite to moderate OSR: objects are not individuals but they have intrinsic properties; he dubs this view ‘intermediate OSR’. According to Lyre, we can identify these properties by looking for what remains invariant under symmetry transformations. For example, electric charge is an intrinsic property because it remains invariant under the QED gauge group $U(1)$. It is not entirely clear what motivates the view that invariants are really best conceived of as intrinsic properties, and that these properties are the ones that we are to regard as real. What makes invariant properties so special compared to non-invariant ones? Perhaps what Lyre has in mind is that an invariant of a physical system is possessed independently of any relations the physical system has to other things, but this could be true of other intrinsic properties too. A further issue is that

⁷⁴ For a discussion of the modal aspects of structures see Esfeld (2009).

some symmetries pertain to entire systems rather than the parts of it that we traditionally regard as individuals. There is at least a question about what kind of intrinsic properties result from such invariances.

4.2 Arguing for OSR

There are two groups of arguments in support of OSR. Arguments belonging to the first group claim that OSR avoids conceptual problems facing rival views. The arguments in the second group focus on particular pieces of physics and submit that these require a radical revision of traditional ontology.

In the first group two arguments stand out. The first is that OSR is not subject to Newman's Objection (Ladyman 1998; Ladyman and Ross 2007). As we have seen above, at least some versions of OSR are not committed to an extensional understanding of relations, but it is exactly this construal of relations that lies at the heart of the proof of Newman's Theorem. Ladyman and Ross are correct in pointing out that shifting to a substantive understanding of relations at least *prima facie* undercuts the proof of Newman's Theorem and hence makes the problem go away. In fact, reverting to a substantive understanding of relations was Newman's own solution to his problem (see Section 3.4.1).

The second argument in the first group is that OSR is the only metaphysical position that leaves no gap between ontology and epistemology because it does not postulate the existence of things about which nothing can be known. This is a desirable metatheoretical feature, which tells in favour of these positions (Esfeld 2004; Esfeld and Lam 2008).⁷⁵ But why should there be no gap between epistemology and ontology? Why should there be only things we can know about? What underlies this argument is some version of Occam's razor, and whether or not one should accept such a principle or its application in this specific case is an issue that demands further scrutiny. A pertinent side remark here is that the epistemic version of structural realism *need not* postulate the existence of entities or properties about which we can have no knowledge. In other words, ESR can simply remain agnostic about anything that is beyond our grasp. It is thus far from clear whether OSR enjoys a metaphysical advantage over ESR (thus understood).

The main tenor of arguments belonging to the second group is that OSR is forced upon us by modern physics, since many parts of modern physics are in fact incompatible with traditional metaphysical views. The focus of the discussion has been on quantum mechanics, quantum field theory, space-time physics, and the use of group theory. Many of these arguments draw on a detailed analysis of the physical theories at hand and defy summary in simple terms. For reasons of space we here restrict attention to quantum mechanics (QM).⁷⁶

⁷⁵ Strictly speaking Esfeld and Lam make this argument for their version of OSR, but it is obvious that the argument applies to other versions as well.

⁷⁶ Quantum field theory is discussed in Cao (1997, 2003a), Lyre (2004), Kantorovich (2003) and Saunders (2003a). Group theory is the focus of French (1998, 1999, 2000) and Castellani (1998). Space-time theories are discussed with a special focus on OSR in Dorato (2000), Esfeld and Lam (2008), Ladyman (2001), Pooley (2006), Stachel (2002) and Saunders (2003c, 2003d). Lorentz's theory of electrons is discussed in Cei (2005). Tegmark (2008) argues that OSR is the only consistent way of believing the existence of the external world.

Everyday objects like chairs, houses, and bicycles are commonly held to be individuals in that we can refer to them with singular terms (we can name them, or refer to them with indexical expressions or definite descriptions) and meaningfully compare them to other objects. The same is typically true of elementary particles in classical physics.⁷⁷ In QM the notion of an individual becomes problematic, and many have argued that QM in fact entails that quantum particles are not individuals.

To assess this claim we need to have a principle of individuality, i.e. a criterion for when two objects are distinct. The principle used in the debates over QM is Leibniz's principle, which individuates objects in terms of their properties. This principle, also referred to as the *principle of the identity of indiscernibles* (PII) : $\forall x\forall y[(\forall P)(Px \leftrightarrow Py) \rightarrow (x = y)]$, where x and y are first order variables (ranging over objects) and P is a second order variable (ranging over properties).⁷⁸ In words, PII says that if objects x and y share all properties, then they are identical; or conversely, if two objects are not in fact one and the same, then there must be at least one property that one possesses and the other does not. The crucial question is what kind of properties P ranges over. Typically three options are discussed: (a) only monadic properties; (b) monadic and relational properties excluding spatial properties; (c) monadic and relational properties including spatial properties. Since the first two are not sufficient to determine individuality even in classical mechanics only the third option is interesting in what follows. The difference between classical and quantum physics lies in how they fare with respect to the third option: PII in the third version is true in classical mechanics (since particles are impenetrable), but it seems to fail in QM.

The argument for this conclusion goes as follows.⁷⁹ Consider two quantum particles that have the same intrinsic non-state dependent properties (mass, charge, ...) and assume that they are, at least initially, considered individuals being labelled '1' and '2'. The problem we have to solve now is to distribute these two particles on two quantum states by $|a\rangle$ and $|b\rangle$. We then write $|a\rangle|b\rangle$ to indicate that the first particle is in state $|a\rangle$ and the second is in state $|b\rangle$. So one would expect that the two-particle system can be in four different states: $|a\rangle|a\rangle$, $|a\rangle|b\rangle$, $|b\rangle|a\rangle$, and $|b\rangle|b\rangle$. While this would be correct if the particles were classical particles, QM tells us that the system can actually only be in one of the following states:

- (1) $|a\rangle|a\rangle$
- (2) $1/\sqrt{2} (|a\rangle|b\rangle + |b\rangle|a\rangle)$
- (3) $1/\sqrt{2} (|a\rangle|b\rangle - |b\rangle|a\rangle)$
- (4) $|b\rangle|b\rangle$

Furthermore, these states are not available to all particles. There are two fundamental kinds of particles: fermions and bosons. Particles cannot change their status over time: if a particle is a fermion (or boson) at one point in time it is fermion

⁷⁷ However, notice that classical particles can be regarded as indistinguishable (Saunders 2006, 52).

⁷⁸ Sometimes Leibniz's law is more broadly associated with the conjunction of PII and its converse, called 'the principle of the indiscernibility of identicals', $\forall x\forall y[(x = y) \rightarrow (\forall P)(Px \leftrightarrow Py)]$.

⁷⁹ This set-up is introduced in French and Redhead (1988). French (2006a) provides a comprehensive and accessible survey of the discussion.

(or boson) forever. It follows from QM that fermions can only be in state (3), while bosons can be in states (1), (2) and (4).

French and Redhead (1988, 238–242) show that particles in one of these states have the same (state-dependent) monadic and relational properties in the following sense:⁸⁰

Monadic: For every property P and every value p of P , the probability that a measurement of property P on particle 1 has outcome p is equal to the probability that a measurement of property P on particle 2 has outcome p .

Relational: For every two properties P and Q and all values p and q of these properties, the probability that a measurement of P on particle 1 has outcome p given that a measurement of Q on particle 2 has outcome q is equal to the probability that a measurement of P on particle 2 has outcome p given that a measurement of Q on particle 1 has outcome q .

These results are derived using standard QM, and hence from the point of view of standard QM the two particles are indistinguishable. Since position is a state-dependent property in QM, the results extend to spatial properties and hence QM particles are indistinguishable also with respect to their spatial properties.

From this it follows that the antecedent of PII is true. But we don't want to conclude that the two particles are identical: clearly we have a two-particle system and not a one-particle system. So PII is false and cannot be used to confer individuality upon quantum particles. This leaves us with two options. Either we can uphold PII as a criterion of individuality and conclude that quantum particles are not individuals; or we can overthrow PII and claim that quantum particles are individuals but they are so on grounds other than that they satisfy PII, e.g. they could have primitive thisness or haecceity. As French and Redhead point out, the latter option does *not* stand in conflict with QM. It is thus available to those who have no metaphysical quibbles with primitive thisness, haecceity or any other PII-violating notion of individuality.

How does this argument, if correct, support OSR? This is the place to guard against a common misunderstanding. It would *seem* that supporters of EOSR should uphold PII and embrace the conclusion that quantum particles are non-individuals. This seems to be exactly the result needed to support attenuated EOSR (which postulates an ontology of non-individual objects), and it also seems to be grist to the mill of unqualified EOSR since individuality is an essential aspect of our conception of an object and so a tempting way to resolve the problem of having objects that fail to be individuals is to get rid of objects altogether.

However, supporters of EOSR have argued in a different way (see, for instance, Ladyman 1998, 420–421). Rather than embracing PII and some kind of non-individuality, they have placed emphasis on the fact that QM fails to determine whether or not quantum particles are individuals since QM is compatible with either option. So we are in a situation of underdetermination of the metaphysics by the

⁸⁰ QM only provides probabilities for a property to obtain. So we here extend the above statement of PII to include the probabilistic case: if for all properties the probability to be instantiated in x equals the probability to be instantiated in y , then x and y are identical.

physics. This is perceived to be a problem, and, so the argument goes, the best solution to the problem is to eliminate those entities whose characteristics are underdetermined, namely objects or at least individuals.⁸¹

This seems to be a curious form of argument. While it is certainly true that problems concerning the nature of objects can go away by eliminating objects from one's ontology, it is not clear why this is not just a case of burying one's head in the sand. Moreover, why should the fact that physics leaves open metaphysical questions be a cause for concern to begin with? For one, why would one expect physics to completely determine the answers to what are genuinely metaphysical questions? Physics can of course constrain the space of acceptable solutions, and this is why contemporary metaphysicians pay so much attention to science, but it is perhaps unrealistic to expect science to single out one view as the only correct one. For another, such underdetermination is rather common. For instance, physics leaves it open whether properties are universals, tropes, nominalistic classes, or yet something else. Should we take this as an invitation to eliminate properties from our ontology? It seems that at most underdetermination lends support to agnosticism between the two views of particles and, hence, agnosticism about OSR (see Cao 2003a; Chakravartty 2003; Morganti 2004 and Saunders 2003b).⁸²

This argument is open to three further lines of criticism. The first is that there actually is no underdetermination because the option of regarding quantum particles as non-individuals vanishes under closer inspection. Saunders (2003d, 2006) points out that we have overlooked irreflexive relations like 'having opposite direction of each component of spin to ...', which render two fermions in an entangled state discernible, and hence their individuality can be grounded in PII after all. As Saunders himself notes, irreflexive relations of this kind fail to make bosons distinguishable, and hence a grounding of the individuality of bosons in PII cannot be supported.⁸³

The second counter is due to Morganti (2008, 2009c) who argues that the choice of PII as a principle of individuality is ill-motivated, both from a philosophical and from a naturalistic point of view, since it neither does justice to an empiricist standpoint (as is often claimed) nor does it sit well with the practice of science. He claims that a position that takes individuality to be primitive fares better on both counts, and goes on to develop an ontology of tropes that provides what seems to be a viable theory of identity (2009a; d). Hence, again, the conclusion is that underdetermination vanishes since non-individuality is not really an option.

The third argument takes issue with the use of standard QM in the derivation of the above result. Standard QM, so the argument goes, is well known to be seriously

⁸¹ For more on how this kind of underdetermination relates to other more traditional kinds of underdetermination see French (2011).

⁸² There are cases where this form of argument has perhaps more currency. Take the question of the existence of any supernatural being. In the absence of evidence supporting such a being's existence many people would simply eliminate it from their ontology instead of merely be agnostics about it. The question then becomes whether individuals and/or objects are similarly eliminable.

⁸³ See also Ketland (2006). The claim about bosons is disputed in Muller and Seevinck (2009). Muller (2011) also argues that far from fearing the weak discernibility of elementary particles, advocates of OSR should embrace it, because, according to him, it points to a relational understanding of particles that lends credence to OSR. For further discussions of the individuality in QM see Morganti (2009a, b, e).

defective from a conceptual point of view. It is therefore at the very least too hasty to draw robust metaphysical conclusions from such a theory. Other approaches to QM should also be considered when discussing questions surrounding individuality. Van Fraassen (1985) examines the issue within the modal interpretation of QM, and Brown et al. (1999) focus on Bohmian mechanics.

In sum, the jury still seems to be out on whether modern physics really favours OSR over its more traditional rivals. Moreover the question remains whether OSR, and ESR for that matter, can give an adequate account of the ontology and epistemology of other sciences. The bulk of the literature on SR has thus far focussed on modern (and in particular fundamental) physics. This is no accident of history. A structuralist analysis of scientific theories usually departs from those theories' mathematical formalism, and formalisation is the hallmark of modern physics. Therefore SR seems to be at odds with less formal sciences such as biology or the social sciences. This has led some critics to claim that SR is a philosophy with little, if any, relevance outside the province of physics (Gower 2000; Newman 2005).

There are several options for the OSRist. First, the OSRist can be a reductionist about all sorts of other ontologies claiming that the only ontology that really matters is the one at the fundamental physical level. Second, the OSRist can adopt some form of local realism and restrict OSR to fundamental physics, and accept that other levels may have dissimilar ontologies. Third, the OSRist can accept the ontological autonomy of higher levels of organisation but try to show that even there all that exists is structure (in which case her underdetermination arguments must be shown to apply at these levels too).

Some SRists have recently begun tackling the challenge posed by the 'other' sciences. Kincaid (2008), Ladyman and Ross (2007), and Ross (2008) discuss SR in the social sciences and economics and argue that SR ideas in fact not only square well with these sciences, but can also shed light on old problems. These papers mark the beginning of the discussion, but it remains to be seen whether convincing SR philosophies of these sciences can be developed.

5 Conclusion

Structural realism is a vibrant research programme in philosophy. As we have seen, it faces important challenges, but it also promises to have a bright future. In this concluding section, we would like to draw up a list of the most pressing of those challenges in the hope that we can encourage targeted research on them. In no specific order of importance these are the following. First, there is no escape from the challenge posed by the history of science. If history is going to prove to be an ally of the SRists, or at least not a foe, then a more thorough investigation needs to be conducted into the conditions under which structure is preserved. Second, underdetermination arguments are often a plague for realist views. Yet, as we saw earlier, SRists try to turn versions of such arguments to their advantage. The epistemological and ontological underdetermination of objects may be a blessing to SRists, but underdetermination of the structure itself (like the ones we saw in section 3) seems to be a curse. A well-developed form of SR must address the

spectre of underdetermination in its various manifestations. Third, SR cannot continue ignoring sciences other than physics. A cogent SR approach towards chemistry and biology as well as the social sciences and the humanities must be articulated. Fourth, the Newman problem and related issues need a secure resolution. Fifth, the subtleties of the relationship between ESR and OSR (in their various guises) need to be explored. One pertinent question in this context is whether our epistemic commitments must be in perfect alignment with our metaphysical ones, as OSR commands. If not, as ESR decrees, how much misalignment is allowed and why? Sixth, SRists need to settle on the right tools. This includes the right notion of structure, the right semantics, the right notion of structure preservation, etc. Seventh, due justice must be given to modality and causality, either by dismissing them as mere unscientific fancy or by working out how, if at all, they cohere with the structuralist standpoint. Eighth, it usually goes without saying, but given the sorry state of parts of the literature we cannot afford to do so: The discussion must be kept rigorous. This means providing crisp formulations of the positions and the existing arguments instead of the more-often-than-desired dependence on hand-waving.

In writing this critical survey we hope to have provided the reader with a bountiful resource that will help her get to grips with the foregoing list of challenges. The easy part is over. Let's get started on the hard one.

Acknowledgements We would like to thank Peter Ainsworth, Anjan Chakravartty, Steven French, Jeff Ketland, James Ladyman, Matteo Morganti, F.A. Muller, John Worrall and two anonymous referees for helpful discussions and/or comments on earlier drafts of the paper. Roman Frigg wishes to acknowledge financial support from the Descartes Centre at the University of Utrecht, where part of the paper has been written, and of the Grant FFI2008-01580 of the Spanish Ministry of Science and Innovation. Ioannis Votsis similarly wishes to acknowledge financial support from the German Research Foundation (Deutsche Forschungsgemeinschaft) as well as from the Center for Philosophy of Science at the University of Pittsburgh, where he was a visiting fellow in the Fall term 2010.

References

- Ainsworth, P. (2008). *Structural realism: A critical appraisal*. PhD Thesis, University of London.
- Ainsworth, P. (2009). Newman's objection. *British Journal for the Philosophy of Science*, 60, 135–171.
- Armstrong, D. (1989). *Universals: An opinionated introduction*. Boulder: Westview.
- Boolos, G., & Jeffrey, R. (1989). *Computability and logic* (3rd ed.). Cambridge.
- Boyd, R. (1985). Lex orandi est lex credendi. In P. Churchland & C. A. Hooker (Eds.), *Images of science* (pp. 3–34). Chicago: University of Chicago Press.
- Brading, K., & Landry, E. (2006). Scientific structuralism: Presentation and representation. *Philosophy of Science*, 73, 571–581.
- Brown, H., Sjöqvist, E., & Bacciagaluppi, G. (1999). Remarks on identical particles in de Broglie-Bohm Theory. *Physics Letters A*, 251, 229–235.
- Bueno, O. (1997). Empirical adequacy: A partial structures approach. *Studies in History and Philosophy of Science*, 28A(4), 585–610.
- Bueno, O. (1999). What is structural empiricism? Scientific change in an empiricist setting. *Erkenntnis*, 50(1), 59–85.
- Bueno, O. (2000). Empiricism, scientific change and mathematical change. *Studies in History and Philosophy of Science*, 31(2), 269–296.
- Busch, J. (2003). What structures could not be. *International Studies in the Philosophy of Science*, 17, 211–225.
- Cao, T. Y. (1997). *Conceptual development of 20th century field theories*. Cambridge: Cambridge University Press.

- Cao, T. Y. (2003a). Structural realism and the interpretation of quantum field theory. *Synthese*, 136(1), 3–30.
- Cao, T. Y. (2003b). Can we dissolve physical entities into mathematical structures? *Synthese*, 136(1), 57–71.
- Carnap, R. (1928). *Der Logisches Aufbau der Welt*. Berlin: Schlachtensee Weltkreis-Verlag.
- Carnap, R. (1966). *An introduction to the philosophy of science*. New York: Dover. 1995.
- Castellani, E. (1998). Galilean particles: An example of constitution of objects. In E. Castellani (Ed.), *Interpreting bodies: Classical and quantum objects in modern physics* (pp. 181–194). Princeton: Princeton University Press.
- Caws, P. (2000). *Structuralism. A philosophy for the human sciences*. Humanity Book: Amherst/NY.
- Cei, A. (2005). Structural distinctions: entities, structures and changes in science. *Philosophy of Science*, 72(Supplement), 1385–1396.
- Cei, A., & French, S. (2006). Looking for structure in all the wrong places: Ramsey sentences, multiple realisability, and structure. *Studies in History and Philosophy of Science*, 37, 633–655.
- Chakravartty, A. (1998). Semirealism. *Studies in History and Philosophy of Science*, 29A(3), 391–408.
- Chakravartty, A. (2001). The semantic or model-theoretic view of theories and scientific realism. *Synthese*, 127, 325–345.
- Chakravartty, A. (2003). The structuralist conception of objects. *Philosophy of Science*, 70(5), 867–878.
- Chakravartty, A. (2004). Structuralism as a form of scientific realism. *International Studies in the Philosophy of Science*, 18, 151–171.
- Chakravartty, A. (2007). *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge: Cambridge University Press.
- Creath, R. (1998). Carnap, Rudolf. In E. Craig (Ed.), *Routledge Encyclopedia of Philosophy*. London: Routledge. Retrieved from <http://www.rep.routledge.com/article/DD012SECT1>.
- Cruse, P. (2004). Scientific realism, Ramsey sentences, and the reference of theoretical terms. *International Studies in the Philosophy of Science*, 18, 133–149.
- Cruse, P. (2005). Empiricism and Ramsey's account of theories. In H. Lillehammer & H. Mellor (Eds.), *Ramsey's legacy*. Oxford: Oxford University Press.
- Da Costa, N., & French, S. (1990). The model-theoretic approach in the philosophy of science. *Philosophy of Science*, 57, 248–265.
- Daston, L., & Galison, P. (2007). *Objectivity*. New York: Zone Books.
- Demopoulos, W. (2003). On the rational reconstruction of our theoretical knowledge. *British Journal for the Philosophy of Science*, 54, 371–403.
- Demopoulos, W. (2008). Some remarks on the bearing of model theory on the theory of theories. *Synthese*, 164, 359–383.
- Demopoulos, W., & Friedman, M. (1985). Critical notice: Bertrand Russell's *The Analysis of Matter*: Its historical context and contemporary interest. *Philosophy of Science*, 52, 621–639.
- Descartes, R. (1641). In S. Tweyman (Ed.), *Meditations on first philosophy*. London: Routledge. 1993.
- Diez, J. (2005). 'The Ramsey sentence and theoretical content. In M. J. Frápoli (Ed.), *F. P. Ramsey: Critical reassessments* (pp. 70–103). London and New York: Continuum.
- Dipert, R. (1997). The mathematical structure of the world: the world as graph. *Journal of Philosophy*, 94(7), 329–358.
- Dorato, M. (2000). Substantivalism, relationism and structural spacetime realism. *Foundations of Physics*, 30(10), 1605–1628.
- Duhem, P. ([1914] 1991). *The aim and structure of physical theory*. Princeton: Princeton University Press.
- Elsamahi, M. (2005). A critique of localised realism. *Philosophy of Science*, 72(Supplement), 1350–1360.
- English, J. (1973). Underdetermination: Craig and Ramsey. *Journal of Philosophy*, 70, 453–462.
- Esfeld, M. (2004). Quantum entanglement and a metaphysics of relations. *Studies in History and Philosophy of Modern Physics*, 35, 601–617.
- Esfeld, M. (2009). The modal nature of structures in ontic structural realism. *International Studies in the Philosophy of Science*, 23, 179–194.
- Esfeld, M., & Lam, V. (2008). Moderate structural realism about space-time. *Synthese*, 160, 27–46.
- Floridi, L. (2008). A defence of informational structural realism. *Synthese*, 161, 219–253.
- French, S. (1998). On the withering away of physical objects. In E. Castellani (Ed.), *Interpreting bodies: Classical and quantum objects in modern physics* (pp. 93–113). Princeton: Princeton University Press.
- French, S. (1999). Models and mathematics in physics: The role of group theory. In J. Butterfield & C. Pagonis (Eds.), *From physics to philosophy* (pp. 187–207). Cambridge: Cambridge University Press.
- French, S. (2000). The reasonable effectiveness of mathematics: Partial structures and the application of group theory to physics. *Synthese*, 125, 103–120.

- French, S. (2001). *Symmetry, structure and the constitution of objects*, unpublished manuscript available on the Pittsburgh Archive (<http://philsci-archive.pitt.edu>), ID Code 327.
- French, S. (2003). Scribbling on the blank sheet: Eddington's structuralist conception of objects. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 34(2), 227–259.
- French, S. (2006a). Identity and individuality in quantum theory. *Stanford Encyclopedia of Philosophy* (URL=<http://plato.stanford.edu/>), Spring 2006 Edition.
- French, S. (2006b). Structure as a weapon of the realist. *Proceedings of the Aristotelian Society*, 106(1), 169–187.
- French, S. (2011). Metaphysical underdetermination: why worry? *Synthese*, 180, 205–221.
- French, S., & Krause, D. (1995). A formal approach to quantum non-individuality. *Synthese*, 102, 195–214.
- French, S., & Krause, D. (2006). *Identity in physics: A historical philosophical, and formal analysis*. Oxford: Clarendon.
- French, S., & Ladyman, J. (1999). Reinflating the semantic approach. *International Studies in the Philosophy of Science*, 13, 103–121.
- French, S., & Ladyman, J. (2003a). Remodelling structural realism: quantum physics and the metaphysics of structure. *Synthese*, 136, 31–56.
- French, S., & Ladyman, J. (2003b). The dissolution of objects: Between platonism and phenomenalism. *Synthese*, 136, 73–77.
- French, S., & Ladyman, J. (2011). In defence of ontic structural realism. In A. Bokulich & P. Bokulich (Eds.), *Scientific structuralism* (Boston Studies in the Philosophy and History of Science) (pp. 25–42). Springer.
- French, S., & Redhead, M. (1988). Quantum physics and the identity of indiscernibles. *British Journal for the Philosophy of Science*, 39, 233–246.
- Friedman, M. (2011). Carnap on theoretical terms: Structuralism without metaphysics. *Synthese*, 180, 249–263.
- Frigg, R. (2006). Scientific representation and the semantic view of theories. *Theoria*, 55, 49–65.
- Giedymin, J. (1982). *Science and convention: Essays on Henri Poincaré's philosophy of science and the conventionalist tradition*. Oxford: Pergamon.
- Giere, R. (2006). *Scientific perspectivism*. Chicago: University of Chicago Press.
- Gower, B. (2000). Cassirer, Schlick and “Structural” realism: the philosophy of the exact sciences in the background to early logical empiricism. *British Journal for the History of Philosophy*, 8(1), 71–106.
- Hacking, I. (1983). *Representing and intervening. Introductory topics in the philosophy of natural science*. Cambridge: Cambridge University Press.
- Hellman, G. (1989). *Mathematics without numbers: Towards a modal-structural interpretation*. Oxford.
- Hellman, G. (1996). Structuralism without structures. *Philosophia Mathematica*, 3, 100–123.
- Hellman, G. (2001). Three varieties of mathematical structuralism. *Philosophia Mathematica*, 9, 184–211.
- Hesse, M. (1963). *Models and analogies in science*. London: Sheed and Ward.
- Hodges, W. (1997). *A shorter model theory*. Cambridge: Cambridge University Press.
- Howson, C. (2000). *Induction: Hume's problem*. Oxford: Clarendon.
- Hume, D. ([1739] 1975). In L. A. Selby-Bigge & P. H. Niddich (Eds.), *A treatise of human nature*. Oxford: Clarendon Press.
- Kantorovich, A. (2003). The priority of internal symmetries in particle physics. *Studies in History and Philosophy of Modern Physics*, 34, 651–675.
- Ketland, J. (2004). Empirical adequacy and Ramsification. *British Journal for the Philosophy of Science*, 55(2), 287–300.
- Ketland, J. (2006). Structuralism and the identity of indiscernibles. *Analysis*, 66, 303–315.
- Ketland, J. (2009). Empirical adequacy and ramsification, II. In A. Hieke & H. Leitgeb (Eds.), *Reduction - abstraction - analysis*.
- Kincaid, H. (2008). Structural realism and the social sciences. *Philosophy of Science*, 75(Supplement), 720–731.
- Ladyman, J. (1998). What is structural realism? *Studies in History and Philosophy of Science*, 29, 409–424.
- Ladyman, J. (2001). Science, metaphysics and structural realism. *Philosophica*, 67, 57–76.
- Ladyman, J. (2005). Mathematical structuralism and the identity of indiscernibles. *Analysis*, 65, 218–221.
- Ladyman, J. (2007). On the identity and diversity of individuals. *Proceedings of the Aristotelian Society*, Supplementary Volume LXXXI, pp. 23–43.

- Ladyman, J., & Ross, D. (2007). *Every thing must go: Metaphysics naturalised*. Oxford: Oxford University Press.
- Landry, E. (2007). Shared structure need not be shared set-structure. *Synthese*, 158(1), 1–17.
- Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, 48, 19–49.
- Leitgeb, H., & Ladyman, J. (2008). Criteria of identity and structuralist ontology. *Philosophia Mathematica*, 16, 388–396.
- Lewis, D. (1970). How to define theoretical terms. *Journal of Philosophy*, 67(13), 427–446.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61, 343–377.
- Locke, J. (1690). In P. H. Nidditch (Ed.), *An essay concerning human understanding*. Oxford: Oxford University Press. 1975.
- Lyre, H. (2004). Holism and structuralism in U(1) Gauge theory. *Studies in History and Philosophy of Modern Physics*, 35, 643–670.
- Lyre, H. (2011). Is structural underdetermination possible? *Synthese*, 180, 235–247.
- Magnus, P. D., & Callender, C. (2004). Realist ennui and the base rate fallacy. *Philosophy of Science*, 71, 320–338.
- Maxwell, G. (1962). The ontological status of theoretical entities. In H. Feigl & G. Maxwell (Eds.), *Scientific explanation, space, and time, vol. 3, Minnesota studies in the philosophy of science* (pp. 3–15). Minneapolis: University of Minnesota Press.
- Maxwell, G. (1968). Scientific methodology and the causal theory of perception. In I. Lakatos & A. Musgrave (Eds.), *Problems in the philosophy of science*. Amsterdam: North-Holland Publishing Company.
- Maxwell, G. (1970). Structural realism and the meaning of theoretical terms. In S. Winokur & M. Radner (Eds.), *Analyses of theories, and methods of physics and psychology* (pp. 181–192). Minneapolis: University of Minnesota Press.
- Maxwell, G. (1971). Theories, perception and structural realism. In R. Colodny (Ed.), *Nature and function of scientific theories* (pp. 3–34). Pittsburgh: University of Pittsburgh Press.
- Melia, J., & Saatsi, J. (2006). Ramseyfication and theoretical content. *British Journal for the Philosophy of Science*, 57(3), 561–585.
- Mertz, D. (2003). An instance ontology for structures. *Metaphysica*, 1, 127–164.
- Mill, J. S. ([1874] 2008). *A system of logic ratiocinative and inductive*. New York: Cosimo Inc.
- Morganti, M. (2004). On the preferability of epistemic structural realism. *Synthese*, 142(1), 81–107.
- Morganti, M. (2008). *Identity, individuality and the ontological interpretation of quantum mechanics*. PhD Thesis, University of London.
- Morganti, M. (2009a). A new look at relational holism in quantum mechanics, forthcoming in *Philosophy of Science* (Supplement).
- Morganti, M. (2009b). Individual particles, properties and quantum statistics, forthcoming *Proceedings of the 2007 Founding Conference of the EPSA*, Springer.
- Morganti, M. (2009c). Weak discernibility, quantum mechanics and the generalist picture, forthcoming in *Facta Philosophica*.
- Morganti, M. (2009d). Tropes and physics. *Grazer Philosophische Studien*, 78, 85–105.
- Morganti, M. (2009e). Inherent properties and statistics with individual particles in quantum mechanics. *Studies in History and Philosophy of Modern Physics*, 40, 223–231.
- Muller, F. A. (1998). *Structures for everyone*. Amsterdam.
- Muller, F. A. (2010). The characterisation of structure: Definition versus axiomatisation. In F. Stadler et al. (Eds.), *The Philosophy of Science in a European Perspective, vol. 1* (pp. 399–416). Dordrecht: Springer.
- Muller, F. A. (2011). Withering away, weakly. *Synthese*, 180, 223–233.
- Muller, F. A., & Seevinck, M. (2009). Discerning elementary particles. *Philosophy of Science*, 76, 179–200.
- Mundy, B. (1986). On the general theory of meaningful representation. *Synthese*, 67, 391–437.
- Musgrave, A. (1988). The ultimate argument for scientific realism. In R. Nola (Ed.), *Relativism and realism in sciences* (pp. 229–252). Dordrecht: Kluwer.
- Newman, M. (1928). MR. Russell's "causal theory of perception". *Mind*, 37, 137–148.
- Newman, M. (2005). Ramsey sentence realism as an answer to the pessimistic meta-induction. *Philosophy of Science*, 72, 1373–1384.
- Poincaré, H. ([1905] 1952). *Science and hypothesis*. New York: Dover.
- Poincaré, H. ([1913] 1946). The Value of Science, translated by George B. Halsted. In H. Poincaré (Ed.), *The Foundations of science: Science and hypothesis, the value of science, and science and method*. Lancaster: The Science Press.

- Pooley, O. (2006). Points, particles and structural realism. In D. Rickles, S. French, & J. Saatsi (Eds.), *Structural foundations of quantum gravity* (pp. 83–120). Oxford: Oxford University Press.
- Post, H. (1971). Correspondence invariance and heuristics. *Studies in the History and Philosophy of Science*, 2(3), 213–255.
- Psillos, S. (1995). Is structural realism the best of both worlds? *Dialectica*, 49, 15–46.
- Psillos, S. (1999). *Scientific realism: How science tracks truth*. London: Routledge.
- Psillos, S. (2000a). Carnap, the Ramsey-sentence and realistic empiricism. *Erkenntnis*, 52(2), 253–279.
- Psillos, S. (2000b). An introduction to Carnap’s “Theoretical Concepts in Science”, together with Carnap’s ‘Theoretical Concepts in Science. *Studies in History and Philosophy of Science*, 31, 151–172.
- Psillos, S. (2001a). Is structural realism possible? *Philosophy of Science*, 68(Supplement), S13–S24.
- Psillos, S. (2001b). Author’s response to symposium on *Scientific Realism: How Science Tracks Truth*. *Metascience*, 10(3), 366–371.
- Psillos, S. (2006a). The structure, the whole structure and nothing but the structure? *Philosophy of Science*, 73(2006), 560–570.
- Psillos, S. (2006b). Ramsey’s Ramsey-sentences. In M. C. Galavotti (Ed.), *Cambridge and Vienna: Frank P. Ramsey and the Vienna Circle* (pp. 67–90). Berlin: Springer.
- Putnam, H. (1962). What theories are not. In E. Nagel, A. Tarski & P. Suppes (Eds.), *Logic, methodology, and philosophy of science*. Reprinted in: Hilary Putnam: *Mathematics, Matter, and Method. Philosophical Papers, Volume I*. Cambridge 1975, pp. 215–227.
- Putnam, H. (1978). Realism and reason. In *Meaning and the moral sciences* (pp. 123–140). London: Routledge and Kegan Paul.
- Quine, W. (1968). Comment’ in the discussion section of Maxwell’s ‘scientific methodology and the causal theory of perception. In I. Lakatos & A. Musgrave (Eds.), *Problems in the philosophy of science*. Amsterdam: North-Holland Publishing Company.
- Quine, W. (1969). *Ontological relativity and other essays*. New York: Columbia University Press.
- Quine, W. (1992). Structure and nature. *Journal of Philosophy*, 89(1), 5–9.
- Quine, W. (1998). *From stimulus to science* (2nd ed.). Cambridge: Harvard University Press.
- Ramsey, F. (1931). Theories. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other essays* (pp. 101–125). London: Routledge and Kegan Paul.
- Redhead, M. (2001a). The intelligibility of the universe. In A. O’Hear (Ed.), *Philosophy at the new millennium* (pp. 73–90). Cambridge: Cambridge University Press.
- Redhead, M. (2001b). ‘Quests of a realist’, review article of Stathis Psillos’s scientific realism: how science tracks truth. *Metascience*, 10(3), 341–347.
- Resnik, M. (1997). *Mathematics as a science of patterns*. Oxford: Oxford University Press.
- Rickart, C. (1995). *Structuralism and structure. A mathematical perspective*. Singapore: World Scientific.
- Roberts, B. (forthcoming). Group structural realism. *British Journal for the Philosophy of Science*.
- Robinson, H. (2004). Substance. *Stanford Encyclopedia of Philosophy* (URL=<http://plato.stanford.edu/>), Fall 2004 Edition.
- Rosner, J. (1996). Quine’s global structuralism. *Dialectica*, 50, 235–242.
- Ross, D. (2008). Ontic structural realism and economics. *Philosophy of Science*, 75(Supplement), 732–743.
- Russell, B. (1912). *The problems of philosophy*. Oxford: Oxford University Press.
- Russell, B. (1918). *Mysticism and logic and other essays*. London and New York: Longmans, Green.
- Russell, B. (1919). *Introduction to mathematical philosophy*. London: George Allen and Unwin.
- Russell, B. (1927). *The analysis of matter*. London: George Allen & Unwin.
- Russell, B. (1948). *Human knowledge: Its scope and limits*. London: George Allen and Unwin.
- Russell, B. (1968). *The autobiography of Bertrand Russell* (Vol. 2). London: George Allen and Unwin.
- Saatsi, J. (2005). Reconsidering the Fresnel-Maxwell case study. *Studies in History and Philosophy of Science*, 36, 509–538.
- Salmon, W. (1994). Comment: Carnap on realism. In W. Salmon & G. Wolters (Eds.), *Logic, language and the structure of scientific theories*. Pittsburgh: University of Pittsburgh Press.
- Saunders, S. (2003a). ‘Critical notice’: Tian Yu Cao’s “the conceptual development of 20th century field theories”. *Synthese*, 136(1), 79–105.
- Saunders, S. (2003b). Structural realism, again. *Synthese*, 136(1), 127–133.
- Saunders, S. (2003c). Indiscernibles, general covariance, and other symmetries. In A. Ashtekar, D. Howard, J. Renn, S. Sarkar, & A. Shimony (Eds.), *Revisiting the foundations of relativistic physics: Festschrift in honour of John Stachel* (pp. 151–173). Dordrecht: Kluwer.
- Saunders, S. (2003d). Physics and Leibniz’s principles. In K. Brading & E. Castellani (Eds.), *Symmetries in physics: Philosophical reflections* (pp. 289–307). Cambridge: Cambridge University Press.
- Saunders, S. (2006). Are quantum particles objects? *Analysis*, 66, 52–63.

- Schaffer, J. (2003). Is there a fundamental level? *Noûs*, 37, 498–517.
- Schurz, G. (2009). When empirical success implies theoretical reference: a structural correspondence theorem. *British Journal for the Philosophy of Science*. doi:10.1093/bjps/axn049.
- Shapiro, S. (1983). Mathematics and reality. *Philosophy of Science*, 50, 523–548.
- Shapiro, S. (1991). *Foundations without foundationalism. A case for Second-Order Logic*. Oxford.
- Shapiro, S. (1997). *Philosophy of mathematics*. Oxford: Structure and Ontology.
- Shapiro, S. (2000). *Thinking about mathematics*. Oxford.
- Simons, P. (2002). Candidate general ontologies for situating quantum field theory. In M. Kuhlmann, H. Lyre, & A. Wayne (Eds.), *Ontological aspects of quantum field theory* (pp. 33–52). Singapore: World Scientific.
- Sneed, J. (1971). *The logical structure of mathematical physics*. Dordrecht: Reidel.
- Solomon, G. (1989). An addendum to Demopoulos and Friedman (1985). *Philosophy of Science*, 56, 497–501.
- Solomon, G. (1990). What became of Russell's Relation-Arithmetic? *Russell*, 9(2), 168–173.
- Stachel, J. (2002). "The relations between things" versus "the things between relations": The deeper meaning of the hole argument. In D. Malament (Ed.), *Reading natural philosophy: Essays in the history and philosophy of science and mathematics* (pp. 231–266). Chicago and LaSalle: Open Court.
- Stanford, K. (2003a). No refuge for realism: Selective confirmation and the history of science. *Philosophy of Science*, 70, 913–925.
- Stanford, K. (2003b). Pyrrhic victories for scientific realism. *Journal of Philosophy*, 100(11), 553–572.
- Stanford, K. (2007). *Exceeding our grasp: Science, history, and the problem of unconceived alternatives*. Oxford: Oxford University Press.
- Stump, D. (1989). Henri Poincaré's philosophy of science. *Studies in History and Philosophy of Science*, 20(3), 335–363.
- Suárez, M. (2003). Scientific representation: Against similarity and isomorphism. *International Studies in the Philosophy of Science*, 17, 225–244.
- Suppe, F. (Ed.) (1977). *The structure of scientific theories*. Urbana and Chicago.
- Suppe, F. (Ed.) (1998). Scientific theories. In E. Craig (Ed.), *Routledge Encyclopaedia of Philosophy*. London.
- Tegmark, M. (2008). The mathematical universe. *Foundations of Physics*, 38, 101–150.
- Van Fraassen, B. (1980). *The scientific image*. Oxford: Clarendon.
- Van Fraassen, B. (1985). Statistical behaviour of indistinguishable particles: Problems of interpretation. In P. Mittelstaedt & E. W. Stachow (Eds.), *Recent developments in quantum logic* (pp. 161–187). Mannheim.
- Van Fraassen, B. (1991). *Quantum mechanics*. Oxford: Oxford University Press.
- Van Fraassen, B. (1995). A philosophical approach to foundations of science. *Foundations of Science*, 1, 5–9.
- Van Fraassen, B. (1997). Structure and perspective: Philosophical perplexity and paradox. In M. L. Dalla Chiara et al. (Eds.), *Logic and scientific methods*. Dordrecht: Kluwer.
- Van Fraassen, B. (2006). Structure: its shadow and substance. *British Journal for the Philosophy of Science*, 57(2), 275–307.
- Votsis, I. (2003). Is structure not enough? *Philosophy of Science*, 70(5), 879–890.
- Votsis, I. (2004). *The epistemological status of scientific theories: An investigation of the structural realist account*. PhD Thesis, London School of Economics, http://www.votsis.org/PDF/Votsis_Dissertation.pdf.
- Votsis, I. (2005). The upward path to structural realism. *Philosophy of Science*, 72(5), 1361–1372.
- Votsis, I. (2007). Uninterpreted equations and the structure-nature distinction. *Philosophical Inquiry*, 29(1–2), 57–71.
- Votsis, I. (2011a). Structural realism: Continuity and its limits. In A. Bokulich & P. Bokulich (Eds.), *Scientific structuralism* (Boston Studies in the Philosophy and History of Science) (pp. 105–117). Springer.
- Votsis, I. (2011b). How not to be a realist or why we ought to make it safe for closet structural realists to come out. In D. Rickles (Ed.), *Structure, objects and causality* (Western Ontario Series in Philosophy of Science). Dordrecht: Kluwer ch. 3.
- Votsis, I. (2011c). The prospective stance in realism. *Philosophy of Science*.
- Williams, J. (2005). *Understanding poststructuralism*. Bucks: Acumen.
- Worrall, J. (1982). Scientific realism and scientific change. *Philosophical Quarterly*, 32(128), 201–231.
- Worrall, J. (1984). An unreal image. *British Journal for the Philosophy of Science*, 34(1), 65–80.
- Worrall, J. (1989). Structural realism: The best of both worlds? In D. Papineau (Ed.), *The philosophy of science*. Oxford: Oxford University Press. 1996.

- Worrall, J. (1994). How to remain (reasonably) optimistic: Scientific realism and the “luminiferous ether”. In D. Hull, M. Forbes, & R. M. Burian (Eds.), *PSA 1994, vol. 1* (pp. 334–342). East Lansing: Philosophy of Science Association.
- Worrall, J. (2007). Miracles and models: Why reports of the death of structural realism may be exaggerated. In A. O’Hare (Ed.), *Philosophy of science (Royal Institute of Philosophy 61)* (pp. 125–154). Cambridge: Cambridge University Press.
- Worrall, J., & Zahar, E. (2001). Ramseyfication and structural realism. Appendix IV in Zahar E. *Poincaré’s Philosophy: From Conventionalism to Phenomenology*. Chicago and La Salle (IL): Open Court.
- Zahar, E. (1996). Poincaré’s structural realism and his logic of discovery. In J.-L. Greffe, G. Heinzmann, & K. Lorenz (Eds.), *Henri Poincaré: Science and philosophy*. Berlin: Academie Verlag and Paris: Albert Blanchard.
- Zahar, E. (2001). *Poincaré’s philosophy: From conventionalism to phenomenology*. Chicago and La Salle (IL): Open Court.
- Zahar, E. (2004). Ramseyfication and structural realism. *Theoria*, 19, 5–30.