

Preprint from <http://www.thephilosophyofinformation.net>

This paper has been accepted for publication in

Synthese (Springer)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

It is a publisher's requirement to display the following notice:

The documents distributed by this server have been provided by the contributing authors as a means to ensure timely dissemination of scholarly and technical work on a noncommercial basis. Copyright and all rights therein are maintained by the authors or by other copyright holders, notwithstanding that they have offered their works here electronically. It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may not be reposted without the explicit permission of the copyright holder.

In the case of Springer, it is the publisher's requirement that the following note be added:

“An author may self-archive an author-created version of his/her article on his/her own website and his/her institution's repository, including his/her final version; however he/she may not use the publisher's PDF version which is posted on www.springerlink.com. Furthermore, the author may only post his/her version provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: “The original publication is available at www.springerlink.com.”

A Defence of Informational Structural Realism

Luciano Floridi^{1,2,3}

¹Dipartimento di Scienze Filosofiche, Università degli Studi di Bari; ²Faculty of Philosophy and ³IEG, OUCL, Oxford University.

Address for correspondence: St Cross College, OX1 3LZ, Oxford, UK; luciano.floridi@philosophy.oxford.ac.uk

Abstract

This is the revised version of an invited keynote lecture delivered at the *1st Australian Computing and Philosophy Conference (CAP@AU)*; the Australian National University in Canberra, 31 October – 2 November, 2003). The paper is divided into two parts. The first part defends an informational approach to structural realism. It does so in three steps. First, it is shown that, within the debate about structural realism (SR), epistemic (ESR) and ontic (OSR) structural realism are reconcilable. It follows that a version of OSR is defensible from a structuralist-friendly position. Second, it is argued that a version of OSR is also plausible, because not all *relata* (structured entities) are logically prior to relations (structures). Third, it is shown that a version of OSR is also applicable to both sub-observable (unobservable and instrumentally-only observable) and observable entities, by developing its ontology of structural objects in terms of informational objects. The outcome is *informational structural realism*, a version of OSR supporting the ontological commitment to a view of the world as the totality of informational objects dynamically interacting with each other. The paper has been discussed by several colleagues and, in the second half, ten objections that have been moved to the proposal are answered in order to clarify it further.

Keywords

Epistemic structural realism; informational ontology; levels of abstraction; ontic structural realism; structural realism.

1. Introduction: structural realism and its two challenges

Broadly construed, *structural realism* (SR) argues that the structural properties of reality are knowable in themselves, and hence that, with a bit of luck, one may get them right:

SR) Explanatorily, instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) can be, in the best circumstances, increasingly informative about the *relations* that obtain between the (possibly sub-observable, i.e. unobservable or instrumentally-only observable) objects that constitute the system under investigation (through the observable phenomena).

SR encompasses a variety of positions in epistemology and in the philosophy of science, including, in the past, Carnap's, Cassirer's, Duhem's, Eddington's, Poincaré's, Russell's, Schlick's and Whitehead's.¹ As a whole family of theories, SR is confronted by two major challenges. One is *Newman's problem*.² The other may be referred to as the *ontological problem*.

Newman's problem is epistemological. It contends that what SR has to say about our knowledge of reality is either trivial or false. It is also an *external* challenge, for it is typically deployed to reject any version of SR. I shall not dwell upon Newman's "frying pan or fire" alternative partly because it is a well-known objection to SR, partly because it lies beyond the limited scope of this paper. However, it must be acknowledged that it remains a future concern for any structuralist approach, including the one developed in this paper. At this stage, suffice it to say that, if pressed, I would agree with Worrall and Zahar [2001], French and Saatsi [2004] and Votsis [forthcoming] that arguments based on it are probably much less damaging than detractors of SR appear to believe;³ and that, as Esfeld and Lam [forthcoming] argue, an ontology that does not privilege either relata or relations (the package hypothesis

¹ An earlier assessment of Whitehead's and Russell's SR is Heath [1928]. Russell's and Poincaré's versions of SR have been the most influential so far, thanks to Maxwell [1968], Maxwell [1970a] and Maxwell [1970b]; and to Worrall [1989] and Worrall [1994] respectively. Gower [2000] provides an excellent scholarly analysis of the earlier history of the various structural realisms. Psillos [2000a] reconstructs Carnap's structuralist position with admirable accuracy, in order to show that this version too is subject to roughly the same "Newmanian" difficulties that, according to Psillos, undermine Russell's and Poincaré's positions. French [2003] concerns Eddington's ontic structural realism. French [2001] highlights Cassirer's SR and more generally the Kantian roots of SR.

² Newman's problem (Newman [1928]) originally concerned Russell's version of SR. It was later revived by Demopoulos and Friedman [1985] in terms of the Ramsey-sentence. It has recently been expanded in scope by Psillos [1999], chap. 7 and Psillos [2001], who deploys it to reject any form of SR. It is formalised by Ketland [2004], who supports it, and by Votsis [2003], who seeks to counteract it. On the history of the problem, see also French [2003].

³ See also French's analysis of Eddington's response (to which I am sympathetic) to Braithwaite in French [2003]. I hope to discuss Newman's problem from the perspective of an informational approach to SR in a future paper.

discussed later in this article) may not be subject to Newman's problem because of the concrete nature of the *relata* in question.

The ontological problem is a metaphysical and *internal* challenge. It concerns the ontological commitments of SR. This is the topic of this paper. SR may be convincing, but it leaves unspecified the nature of the *relata* in the structures. As a result, the current debate among structural realists has resuscitated the classic question about the accessibility of the ontic status of objects in themselves. The problem does not have to be restricted to theoretical or sub-observable entities. We shall see that, as the Kantian terminology and the Russellian origins of the approach point out, what is in question is the (knowability of the) ultimate nature of reality. I shall return to this methodological point when discussing objection 3 in section 6.

The ontological problem is explainable by means of a simple formalism. The set-theoretic scheme of a structure S comprises four sets (three, if one treats properties as unary relations):

- 1) a non-empty set O of objects (the domain of S),
- 2) a non-empty set P of first-order, one-place properties of the objects in O ,
- 3) a non-empty set of relations R on O , and
- 4) a possibly empty set T of transitions rules (operations) on O .

This is not the only way of dealing with the concept of structure, and it may not be the best, if compared to the semantic or model-theoretic approach. Chakravartty [2001] has argued that it is questionable whether the model-theoretic approach significantly favours scientific realism rather than any other form of realism. If Chakravartty is right, it follows that the approach may not significantly favour SR either. However, this neutrality thesis can be guaranteed while still endorsing the SR-centric view, supported by e.g. French and Ladyman [2003b] and French and Saatsi [2004], that the semantic (model-theoretic) approach is more suitable for SR than the syntactic (set-theoretic) one. So, the semantic (model-theoretic) approach will be adopted below, when it will become necessary to clarify the different ontological commitments of a theory. Here, however, the set-theoretic approach suffices to formulate the ontological problem in very simple terms: what are the objects in O ? *Epistemic Structural Realism* (ESR) and *Ontic Structural Realism* (OSR) provide different answers.⁴

⁴ The distinction between ESR and OSR, introduced by Ladyman [1998], is now standard, even if the vocabulary adopted may differ, with some authors referring to OSR as “metaphysical” or “ontological” SR. For a recent analysis see Psillos [2004]. Van Fraassen [2006] distinguishes between “moderate structuralism: the theory

ESR takes an *agnostic stance*. Objects can be posited only as *ontic residua*, i.e., what remains in principle unknowable once the knowable structures of reality have been factored out. Quite naturally, Poincaré [1902] used a Kantian vocabulary and spoke of objects in themselves as *noumena*. Ramsey sentences provide a standard (if perhaps unsuccessful by themselves, see French and Ladyman [2003a]) method to circumvent an ontological commitment in favour of these *noumena*.⁵

ESR's agnostic stance may be interpreted as what Lewis [forthcoming] called "Ramseyan humility", although I would opt for some Kantian humility instead. Lewis relied on Ramsey sentences and the multiple realizability of even our hypothetically true and ultimate theory of everything to argue that the intrinsic nature of reality *may be unknowable*, whereas Kant's conclusion was that it is in principle *unknown and unknowable*, owing to the nature of human knowledge, a slightly but significantly different position.⁶ Given its Kantian humility, ESR introduces a restricting clause in its definition of SR:

ESR) Explanatorily, instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) can be, in the best circumstances, increasingly informative *only* about the relations that obtain between the (possibly sub-observable) objects that constitute the system under investigation (through the observable phenomena), but not about the first-order, one-place predicates qualifying *the objects in themselves* (the intrinsic nature of the *noumena*).

Heath [1928] referred to ESR as "bifurcated structuralism" because of its obvious (and typically Kantian) dualism. As French [2003] aptly phrases it: "while the ultimate constituents of the world may be non-structural, all that physics – or science in general – tells us about, and hence all that we can know, on the basis of that physics, is structure" (p. 255).

OSR seeks to overcome the "bifurcation" in ESR by not granting the restricting clause. This can be achieved in at least two ways, one radical, the other more moderate. *Eliminativist*⁷ OSR suggests that:

describes only the structure of a bearer, which has also non-structural features (though science is said not to describe those), radical structuralism: "structure is all there is" [and] in-between structuralism: the structure described by science does have a bearer, but that bearer has no other features at all." This paper falls into the "in-between" category.

⁵ This is a standard approach at least since Maxwell [1970b]. Worrall and Zahar [2001] provide a recent defence. Votsis [2003] reviews the various difficulties involved in using Ramsey sentences in SR. For a defence of a Ramseyan approach see Cruse and Papineau [2002], who are criticised by Newman [2004].

⁶ Langton [2004] provides an enlightening analysis of the implications of Lewis' "Ramseyan humility".

⁷ I'm adapting this distinction from Parsons [2004], who discriminates between eliminativist and non-eliminativist forms of structuralism in the philosophy of mathematics. Parsons' non-eliminativist approach is

EOSR) Ultimately, only structures are primitive and ontologically subsistent. Individual objects are not *residua* but *figmenta*, heuristically more useful but no more existent than Bacon's *idola*.

In EOSR, the epistemological problem (i.e. the bifurcation in ESR) is solved because there is nothing to know about the intrinsic nature of individual objects anyway.

Non-eliminativist OSR holds that:

NOSR) Ultimately, there are entities, but they are not classically re-identifiable individuals; rather, they are themselves structural objects, and in the best cases they can be indirectly denoted (circumscribed) by our models, at least in principle.

Both positions move away from any metaphysical substantialism (there are no permanent and more fundamental, individual “hooks”, substantial atoms, monads, substrata or essential elements), so they may not necessarily be distinguished – Saunders [2003], for example, does not – or may be distinguished but interpreted as interchangeable, as in Morganti [2004]. I shall argue that they are quite different.

Eliminativist OSR has emerged also in connection with the debate on individuality in quantum mechanics (French and Redhead [1988]), and this may explain its extremely revisionist nature. The position avoids the bifurcation in ESR by disposing of any ontology of objects *tout court* (this seems to be the position sometimes advocated by Ladyman [1998] and further defended in French and Ladyman [2003a], but see below). The empirical and common-sensical costs of this strategy are evidently very high. They have been convincingly stressed by Cao [2003], Chakravartty [2003] and Morganti [2004]. True, advocates of eliminativist OSR may rejoin that the metaphysical and epistemological gains are even greater: ontic monism and structural knowledge guarantee that reality is fully knowable in principle. Moreover, there are “thin” versions of mathematical structuralism that can bring some support to eliminativist OSR (Reck and Price [2000] provide a useful overview). However, it is easy to counter that such gains are suspiciously great, for at least three reasons, which may be quickly outlined here.

First, as a form of *syntactic realism*, eliminativist OSR betrays the original Kantian polarization between knowable phenomena and unknowable noumena, which lies at the roots of SR. It starts resembling a metaphysically more problematic form of absolute idealism,

related to Dedekind's “logical structuralism”, as defined by Reck [2003]. It seems to be in line with at least a Neo-Kantian view of mathematical objects, and I would consider it coherent with the theses defended in this paper.

according to which “whatever is real is structural and whatever is structural is real”, to paraphrase Hegel.

Second, eliminativist OSR becomes less clearly distinguishable from a stronger form of scientific realism, which may then be unable to cope with the *Pessimistic Meta-Induction Argument* (see section 2.1).

Finally, without objects/entities, eliminativist OSR triggers a potential regress of structures (see section 3).

For these reasons at least, I agree with Cao [2003] that eliminativist OSR should probably be adopted only as a matter of last resort.

Non-eliminativist OSR is equally anti-substantialist but metaphysically less drastic. It seeks to temper the effects of the bifurcation in ESR by arguing that something can be known about structured entities, namely that they are structural objects. Non-eliminativist OSR is expounded and supported by Chakravartty [1998] and Chakravartty [2003], Dorato and Pauri [forthcoming] and Esfeld [2004]. It seems to be the position defended sometimes by other ontic-structuralists such as Cassirer and French and Ladyman [2003b]. It can find allies among “entity realists” in philosophy of science, and gain support from some “thick” version of mathematical structuralism (see Reck and Price [2000]).⁸ It is the version of OSR that will be discussed in the rest of this article, and henceforth OSR will mean non-eliminativist OSR, unless otherwise specified.

2. First step: ESR and OSR are not incompatible

Philosophers sympathetic to SR are divided between those who support (various versions of) OSR or ESR and a third party of those who appreciate the advantages and shortcomings of both and choose to remain uncommitted. This section attempts to convince at least this third party that there is a reasonable way of combining the virtues of both brands of SR without incurring most of their costs.

2.1. Indirect knowledge

SR was revived by Worrall [1989] to deal with two problems in the philosophy of science, the *No-Miracles Argument*:

⁸ Dedekind’s “logical structuralism” is probably the version most compatible with the informational approach developed in this paper, see Reck [2003] for a detailed analysis and support.

NMA) (Some form of) realism “is the only philosophy that does not make [the predictive success of] science a miracle” (Putnam [1975], 73);

and the *Pessimistic Meta-Induction Argument*:

PMIA) Since many predictively-successful scientific theories have been discarded in the past, there is some inductive evidence that current theories are likely to be discarded as well, despite their increased success. (Laudan [1981]).

NMA is neutral with respect to the form of realism that it supports. For example, both instrumentalism (understood as a form of realism insofar as it is committed to the existence of an external world that determines the empirical success of scientific theories, more on this in section 6, objection 4) and naïve realism would be plausible candidates. As for SR, in its case NMA leads to the view that the epistemic success of a theory is a function of its being correct about the structure of reality.

PMIA has been questioned on several grounds. Here, it is useful to highlight three general criticisms.

First, *sub specie aeternitatis*, science is still in its puberty, when some hiccups are not necessarily evidence of any serious sickness. What are a handful of centuries compared to the following millennia? From a longer perspective, science may be settling down, and the initial swings may be just a prelude to a more inert state of “structural stability”.⁹ As Hacking [1999] has emphasised, it already seems that it is not so much revolutionary change as normal stability that requires explanation. Of course, the “sub specie aeternitatis” view provides no support for scientific realism, which was the target of PMIA in the first place. The drying up of scientific revolutions may be just a sign of decreasing empirical or intellectual resources. One day, we may be at our wits’ end, or at the end of our technical or financial means. Then the oscillations and vibrations in our scientific theories may be so subtle and infrequent that our thick-skinned metaphysics may not notice them. Indeed, at that point, it may become difficult to talk of bad metaphysics at all. The risk is that then the best theory available may still be the best of a very bad lot, as Van Fraassen has emphasised.

⁹ “Modern mathematical models are ‘structurally stable’, that is, their (qualitative or approximate quantitative) predictions are insensitive to small changes. A circle is not structurally stable, in that a slightly deformed circle is not a circle. But the system of Königsberg bridges retains exactly the same topological structure if its islands are eroded slightly, or its river narrows.” Franklin [1999]

Second, although the *pax scientifica* provides no reason for optimism, it fails to support pessimism too, for the latter is trivially observer-relative. Look on the bright side and you may transform PMIA into an optimistic meta-induction:

OMIA) The *Optimistic Meta-Induction Argument*: since many false scientific theories have been discarded in the past, despite their being predictively-successful, there is some inductive evidence that, if current theories are also false, they will be discarded as well, notwithstanding their increased success.

The half-full glass of OMIA is consistent with NMA and both together are comfortably reassuring. Precisely because there is no symmetry between semantics and pragmatics – truth is predictively-successful but predictive success is not necessarily truth-tracking – in the long run, we, realists, shall not be fooled by currently predictively-successful but actually untruthful mermaids.

Third, SR may endorse OMIA and defuse PMIA by arguing that, although for the sake of argument one may concede that discontinuity in theory change may be radical¹⁰ when non-structural descriptions of the nature of entities are involved, this is counterbalanced by considerable continuity at the structural level.

Now, NMA and PMIA/OMIA are strictly related to two other broad issues in epistemology:

- i) the goal of vindicating at least partly, and at any rate some of, our ordinary ontological commitments, despite
- ii) the problem that we do not have direct or any other privileged epistemic access to reality in itself.

Goal (i) and problem (ii) bear witness to the Kantian roots of SR. Again, this is not a coincidence. Even if nowadays one tends to associate (i) and (ii) more readily with work beginning with Quine [1939] and Sellars [1956] respectively, the Neo-Kantian origins of structuralism can be traced back to a revival of interest in Kant's transcendental idealism at the beginning of the last century (French and Ladyman [2003b], Floridi [2003b]).

Problem (ii) differentiates SR from forms of naïve realism, whereas goal (i) differentiates SR from forms of instrumentalism.

¹⁰ The “radicality” however need not be guaranteed because of recent, convincing defences of Post's *General Correspondence Principle* in French and Kamminga [1993].

Naïve realism relies on a conception of knowledge as a direct relation between model and modelled. Statically and synchronically, it looks robust at a given time because it is rigidly coupled to whatever description of the world is available at that time. However, dynamically and hence diachronically, naïve realism is dreadfully brittle, for it is easily shattered by any epistemic change. This does not have to be a major scientific revolution; it may be just some simple counterexample or a good old-fashioned, skeptical argument. In fact, naïve realism is most informative when it breaks, exactly like the tell-tale glass fixed across cracks to monitor movement in a building.

Various forms of instrumentalism circumvent the risk of a breakdown by avoiding any specific ontological commitment over and above the minimal acceptance of a mind-independent, external reality. This is the (admittedly very deflated) sense in which even an instrumentalist may claim to be a realist (see the 0-order commitment in section 2.4).

Instead, structural realism may seek to combine some richer ontological commitment – a defence of NMA/OMIA and a positive development of (i) – with a degree of resilience to epistemic change, especially in science. It achieves this flexibility not by decoupling knowledge from reality, as instrumentalism does, but by decoupling, within knowledge itself, the descriptions of the knowable structural characteristics of the system from the explanations of its intrinsic properties. Both can still crash, but the structural descriptions are now not cemented to their source and are therefore far more resilient to change than the ontological explanations, which can be left sinking, if necessary. All this comes at a price, which is paid at the checkout represented by problem (ii) above. SR gains epistemic resilience by endorsing a view of knowledge as an *indirect* relation between the epistemic (possibly multi-) agent and the system under analysis, and a somewhat weakened ontological commitment. Let's see how.

2.2. The Method of Levels of Abstraction

In Kant, knowledge of reality is indirect because of the mind's transcendental schematism. After the downfall of Neo-Kantism and Cassirer's and C. I. Lewis' revisions of the transcendental, an approach is needed that is less infra-subjective, mental (if not psychologistic), innatist, individualistic and rigid. In Floridi and Sanders [2004c] and Floridi and Sanders [2004b], the *method of levels of abstraction* (LoAs) has been proposed as a more inter-subjective, socially constructible (hence possibly conventional), dynamic and flexible way to further Kant's approach. This is a step away from *internal realism* (the kinds, categories and structures of the world are only a function of our conceptual schemes), but not

yet a step into *external* or *metaphysical realism* (the kinds, categories and structures of the world belong to the world and are not a function of our conceptual schemes, either causally or ontologically). If necessary, it might be called *liminal realism*, for reasons that will become clearer below.

The method of LoA comes from modelling techniques developed in an area of Computer Science, known as *Formal Methods*, in which discrete mathematics is used to specify and analyse the behaviour of information systems. What follows is rather standard knowledge in computer science, and the interested reader is referred to the relevant literature for further details (see for example De Roever and Engelhardt [1998], Hoare and He [1998]). Before introducing a quick summary of the method of levels of abstraction, an everyday example may be useful.

Suppose we wish to describe the state of a traffic light in Rome. We might decide to consider an *observable*, named *colour*, of type $\{red, amber, green\}$ that corresponds to the colour indicated by the light. This option abstracts the length of time for which the particular colour has been displayed, the brightness of the light, the height of the traffic light, and so on. So the choice of type corresponds to a decision about how the phenomenon is to be regarded. To specify such a traffic light for the purpose of construction, a more appropriate type might comprise a numerical measure of wavelength. Furthermore, if we are in Oxford, the type of colour would be a little more complex, since – in addition to red, amber and green – red and amber are displayed simultaneously for part of the cycle. So, an appropriate type would be $\{red, amber, green, red-amber\}$. What we have just seen is a basic concept of *level of abstraction* (LoA), understood as a finite but non-empty set of observables.

Since the systems investigated may be entirely abstract or fictional, the term “observable” should not be confused here with “empirically perceivable”. Unfortunately, in computer science an *observable* is something different from what is usually meant by the term in philosophy of science, for it may be the specification of some property of a program. It is just an *interpreted typed variable*, that is, a typed variable together with a statement of what feature of the system under consideration it stands for. Being an abstraction, an observable is not necessarily meant to result from quantitative measurement or even perception. It does not have to be a physical magnitude either. Although the “feature of the system under consideration” might be empirical and physically perceivable, an observable might alternatively be a feature of an artefact or of a conceptual model, constructed entirely for the

purpose of analysis. In what follows, whenever the terminology might be confusing I shall specify which sense of “observable” is being used.

No order is assigned to the observables, which are expected to be the building blocks in a theory characterised by their very definition. The definition of observables is only the first step in studying a system at a given LoA. The second step consists in deciding what relationships hold between the observables. This, in turn, requires the introduction of the concept of system “behaviour”.

Not all values exhibited by combinations of observables in a LoA may be realised by the system being modelled. For example, if the four traffic lights in Oxford are modelled by four observables, each representing the colour of a light, the lights should not in fact all be green together (assuming they work properly). In other words, the combination in which each observable is green should not be realised in the system being modelled, although the types chosen allow it. Some technique is therefore required to describe those combinations of observable values that are actually acceptable. The most general method is simply to describe all the allowed combinations of values. Such a description is determined by a predicate, whose allowed combinations of values we call the “system behaviours”. A *behaviour* of a system, at a given LoA, is defined to consist of a predicate whose free variables are observables at that LoA. The substitutions of values for observables that make the predicate true are called the *system behaviours*.

A *moderated LoA* is defined to consist of a LoA together with a behaviour at that LoA. For example, human height does not take arbitrary rational values, for it is always positive and bounded above by (say) nine feet. The variable h , representing height, is therefore constrained to reflect reality by defining its behaviour to consist of the predicate $0 < h < 9$, in which case any value of h in that interval is a “system” behaviour.

Since Newton and Leibniz, the behaviours of the analogue observables, studied in science, have typically been described by differential equations. A small change in one observable results in a small, quantified change in the overall system behaviour. Accordingly, it is the rates at which those smooth observables vary which is most conveniently described. The desired behaviour of the system then consists of the solution of the differential equations. However, this is a special case of a predicate: the predicate holds at just those values satisfying the differential equation. If a complex system is approximated by simpler systems, then the differential calculus provides a method for quantifying the approximation. The use of predicates to demarcate system behaviour is essential in any (nontrivial) analysis of discrete

systems because in the latter no such continuity holds: the change of an observable by a single value may result in a radical and arbitrary change in system behaviour. Yet, complexity demands some kind of comprehension of the system in terms of simple approximations. When this is possible, the approximating behaviours are described exactly, by a predicate, at a given LoA, and it is the LoAs that vary, becoming more comprehensive and embracing more detailed behaviours, until the final LoA accounts for the desired behaviours. Thus, the formalism provided by the method of abstraction can be seen as doing for discrete systems what differential calculus has traditionally done for analogue systems. The use of predicates is essential in subjects like information theory and computer science, where discrete observables are paramount and hence predicates are required to describe a system behaviour. In particular, state-based methods like *Z* (Hayes and Flinn [1993], Spivey [1992]) provide notation for structuring complex observables and behaviours in terms of simpler ones.

The time has come now to combine approximating, moderated LoAs to form the primary concept of the method of abstraction. For a given (empirical or conceptual) system or feature, different LoAs correspond to different representations or views. A *Gradient of Abstractions* (GoA) is a formalism defined to facilitate discussion of discrete systems over a range of LoAs. Whilst a LoA formalises the scope or granularity of a single model, a GoA provides a way of varying the LoA in order to make observations at differing levels of abstraction. In general, the observations at each LoA must be explicitly related to those at the others; to do so, we use a family of relations between the LoAs. More formally we have that a *gradient of abstractions*, *GoA*, is defined to consist of a finite set¹¹ $\{L_i \mid 0 \leq i < n\}$ of moderated LoAs L_i and a family of relations $R_{i,j} \subseteq L_i \times L_j$, for $0 \leq i \neq j < n$, relating the observables of each pair L_i and L_j of distinct LoAs in such a way that:

1. the relationships are inverse: for $i \neq j$, $R_{i,j}$ is the reverse of $R_{j,i}$
2. the behaviour p_j at L_j is at least as strong as the translated behaviour: $P_{R_{i,j}}(p_i) p_j \Rightarrow P_{R_{i,j}}(p_i)$.

The last condition means that the behaviour moderating each lower LoA is *consistent* with that specified by a higher LoA. Without it, the behaviours of the various LoAs constituting a GoA would have no connection to each other. A special case, to be elaborated below in the definition of “nestedness”, helps to clarify the point.

¹¹ The case of infinite sets has application to analogue systems but is not considered here.

If one LoA L_i extends another L_j by adding new observables, then the relation $R_{i,j}$ is the inclusion of the observables of L_i in those of L_j and the condition specified above reduces to this: the constraints imposed on the observables at LoA L_i remain true at LoA L_j , where “new” observables lie outside the range of $R_{i,j}$. A GoA whose sequence contains just one element evidently reduces to a single LoA. So our definition of “LoA” is subsumed by that of “GoA”.

The consistency conditions imposed by the relations $R_{i,j}$ are in general quite weak. It is possible, though of little help in practice, to define GoAs in which the relations connect the LoAs cyclically. Of much more use are the following two important kinds of GoA: “disjoint” GoAs (whose views are complementary) and “nested” GoAs (whose views provide successively more information). A GoA is called *disjoint* if and only if the L_i are pairwise disjoint (i.e., taken two at a time, they have no observable in common) and the relations are all empty. It is called *nested* if and only if the only nonempty relations are those between L_i and L_{i+1} , for each $0 \leq i < n-1$, and moreover the reverse of each $R_{i, i+1}$ is a surjective function from the observables of L_{i+1} to those of L_i .

A disjoint GoA is chosen to describe a system as the combination of several non-overlapping components. This is useful when different aspects of the system behaviour are better modelled as being determined by the values of distinct observables. This case is rather simplistic, since the LoAs are more typically tied together by common observations.

A nested GoA (see Figure1) is chosen to describe a complex system exactly at each level of abstraction and incrementally more accurately. The condition that the functions be surjective means that any abstract observation has at least one concrete counterpart. As a result, the translation functions cannot overlook any behaviour at an abstract LoA: behaviours lying outside the range of a function translate to the predicate *false*. The condition that the reversed relations be functions means that each observation at a concrete LoA comes from at most one observation at a more abstract LoA (although the converse fails in general, allowing one abstract observable to be refined by many concrete observables). As a result the translation functions become simpler.




Figure 1 Nested GoA with four Levels of Abstraction

Suppose that constants $\lambda_{red} < \lambda_{red'}$ delimit the wavelength of red, and similarly for amber and green. Then the behaviour of L_1 is simply this predicate with free variable wl

$$(\lambda_{red} \leq wl \leq \lambda_{red'}) \vee (\lambda_{amber} \leq wl \leq \lambda_{amber'}) \vee (\lambda_{green} \leq wl \leq \lambda_{green'}).$$

The sequence consisting of the LoA L_0 and the moderated LoA L_1 forms a nested GoA. Intuitively, the smaller, abstract, type $\{red, amber, green\}$ is a projection of the larger. The relevant relation associates to each value $c: \{red, amber, green\}$ a band of wavelengths perceived as that colour. Formally, $R(colour, wl)$ is defined to hold if and only if, for each $c: \{red, amber, green\}$,

$$colour = c \leftrightarrow \lambda_c \leq wl \leq \lambda_{c'}.$$

Following the technique used to define a nested GoA, it is possible to define less restricted but still hierarchical GoAs. Important examples include tree-like structures, of which our nested GoAs are a special, linear case.

The *method of levels of abstraction* has several advantages. First, and most importantly for our present concerns, it is useful to specify the meaning of “indirect knowledge”¹² in terms of knowledge mediated by a LoA.

Second, specifying the LoA (henceforth, the specification LoA/GoA will be dropped) means clarifying, from the outset, the range of questions that (a) can be meaningfully asked and (b) are answerable in principle. One might think of the input of a LoA as consisting of the *system* under analysis, comprising a set of *data*; its output is a *model* of the system, comprising *information*. The quantity of information in a model varies with the LoA: a lower LoA, of greater resolution or finer granularity, produces a model that contains more information than a model produced at a higher, or more abstract, LoA. Thus, a given LoA provides a quantified commitment to the kind and amount of information that can be “extracted” from a system. The choice of a LoA pre-determines the type and quantity of data that can be considered and hence the information that can be contained in the model. So, knowing at which LoA the system is being analysed is indispensable, for it means knowing the scope and limits of the model being developed.

Third, different LoAs may be fairly compared and ranked depending on their virtues, that is, (a) how well they model the system in question according to what purpose and (b) how well they satisfy a list of explicit modelling specifications, e.g. informativeness, coherence, elegance, simplicity, explanatory power, consistency with the data, predictive power, etc..¹³ So different analyses of the same system can be fairly compared provided that they share comparable LoA.

Finally, by stating its LoA, a theory is forced to make explicit and clarify its ontological commitment, in the following way.

¹² Direct knowledge is to be understood here as typically knowledge of one’s mental states, which is apparently not mediated; indirect knowledge is usually taken to be knowledge that is obtained inferentially or through some other form of mediated communication with the world.

¹³ In Devitt’s terminology, (b) comprises the non-empirical virtues of a LoA and (a), when the purpose is truth, the empirical virtues: “Empirical virtue is a matter of entailing (in conjunction with accepted auxiliaries) observational truths and not entailing observational falsehoods. The nonempirical virtues are explanatory power, simplicity, and the like.” Devitt [forthcoming].

2.3. Ontological commitments and levels of abstractions

We have seen that a model is the output of the analysis of a system, developed at some LoA(s), for some purpose. So a theory of a system comprises at least three components:

- i) a LoA, which determines the range of available observables and allows the theory to investigate the system under analysis and to elaborate
- ii) the ensuing model of that system, which identifies
- iii) a structure of the system at the given LoA.

Let us refer to this as the system-level-model-structure (or SLMS) scheme (see Fig. 2).

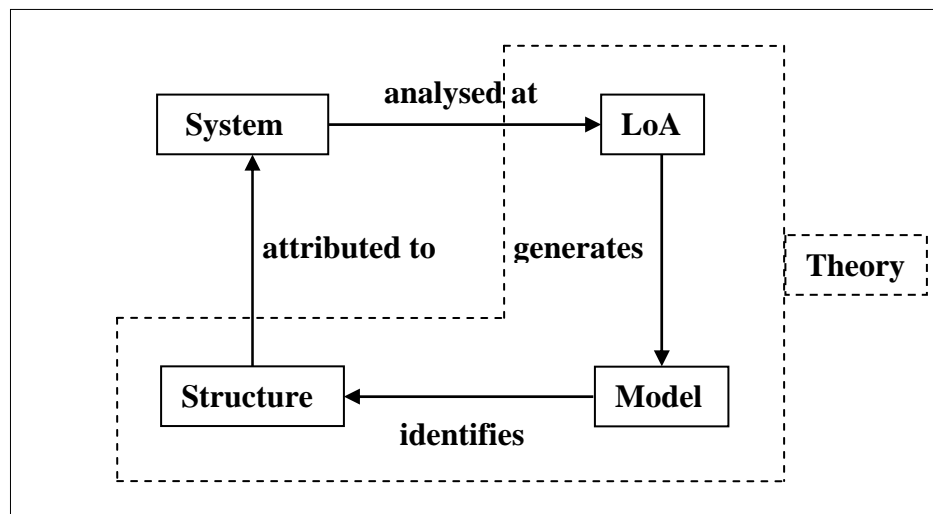


Fig. 2: the SLMS scheme

The structuralist position can now be rephrased thus:

SR*) Explanatorily, instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) *at a given LoA* can be, in the best circumstances, increasingly informative about the relations that obtain between the (possibly sub-observable) objects that constitute the system under investigation (through the observable phenomena).

The ontological commitment of a theory can be clearly understood by distinguishing between a *committing* and a *committed* component, within the SLMS scheme.

A theory commits itself ontologically by opting for a specific LoA, whose application commits the theory to a particular model of the system. The order is purely logical. By adopting a LoA, the theory decides what kind of observables are going to play a role in elaborating the model. In our traffic light example, suppose the LoA commits the theory to

take into account only data relative to colour type. When the LoA generates a model, i.e. when the observables are instantiated, the theory is committed to a particular view of the system. Again, in our example, this might be for instance the specific colours used in the model. According to Votsis [forthcoming], Worrall, for example, seems to hold this much when phrasing SR in terms of “a commitment to structures (including equations) whose observation terms are fully interpreted and whose theoretical terms are implicitly defined through their logical relations with one another and with the observation terms. This just amounts to the Ramsey sentence approach to theories.”

To summarise, by accepting a LoA a theory commits itself to the existence of certain types of objects, the types constituting the LoA (by trying to model a traffic light in terms of three colours one shows one’s commitment to the existence of a traffic light of that kind, i.e. one that could be found in Rome, but not in Oxford), while by endorsing the ensuing models the theory commits itself to the corresponding tokens (by endorsing a particular model, which is the outcome of the interpretation of the data at the chosen LoA, one commits oneself to that model, e.g. one now cannot have a fourth phase when amber and green are on at the same time). Fig. 3 summarises this distinction (note that, for the sake of simplicity the term “theory” is the dotted line that comprises, as above, LoA, model and structure).




Fig. 3: the SLMS scheme with ontological commitment

The distinction just introduced can now be used to reconcile ESR and OSR.

2.4. How to reconcile ESR and OSR

What scope of ontological commitment can a LoA provide to a theory? There is virtually no upper limit. Depending on the LoA, a theory may end up by talking of witches and neurotransmitters, zodiac signs and the Morning Star, ether and white swans, phlogiston and hadrons. However, there is a lower limit, a minimalist approach to how little a theory may safely assume to be there (in the system itself) for the theory (and for the history of knowledge in general and of science in particular) to be consistent and meaningful. ESR endorses such *minimalism*. It argues that, in the best circumstances, a theory is justified in adopting a LoA that commits it (the theory) ontologically to a realist interpretation of the structural properties of the system identified by the model that has been produced by the theory at the chosen LoA. This ontological commitment is minimal in that it concerns knowledge of the first-order structural properties (the non-empty set P of first-order, one-place properties of the objects in O , see section 1) of the system under investigation (more on this in section 6, objection 5). It is, in other words, the sort of commitment that is deemed safer by all supporters of structural realism, since it is based on the existence of certain relations, identified by the model as implemented by the system, rather than of particulars in themselves, with specific properties within the system.

ESR is arguably correct. On the one hand, the adoption of any LoA supporting a degree of ontological commitment less minimal than the epistemic-structuralist one endorsed by ESR seems metaphysically risky and suspicious. This is the point usually made by supporters of ESR: it is better to limit one's ontological commitment to the existence and knowability of (properties of) relational structures. Any other commitment to the knowability of (properties of) the relata in themselves seems unnecessary, and not backed up by a general conception of knowledge, understood as an indirect relation with the world, even if this were not explicitly interpreted in terms of a LoA methodology. In short, Ockham's razor applies and one tries to keep one's commitment as minimal as possible. But could the commitment be even more minimalist? No, because there is no ontological commitment more minimalist than the one supported by ESR – which holds that structures are knowable – but still less minimalist than the extremely deflated commitment of the instrumentalist, who may be happy

to concede just the existence of an external world. Basically, there is no logical space for manoeuvre here between ESR and instrumentalism.¹⁴

An important consequence is that one may believe that, if ESR is correct, this is all that can be said about the issue and there is no logical space left for OSR either.

The impression is widespread. The current internal debate among structural realists has been developing on the assumption that, if there is some room for manoeuvre, it must be found at expense of one of the two positions. It is held that (versions of) OSR and ESR are incompatible, one of the two must go and, it is often argued, this is not ESR. Indeed, since ESR is the most popular variety of SR, Morganti [2004], for example, defends it negatively, by showing that the only SR-friendly alternative to ESR, namely OSR, is untenable. The strategy becomes a case of friendly fire when Newman's problem is further deployed to reject ESR as well, and hence to discard SR altogether. It is a strategy we shall encounter again in section 6, objection 1.

This attack on OSR is, however, misguided. ESR and OSR work at separate LoAs, so they are not mutually exclusive. As far as a first-order analysis is concerned, SR justifies only one kind of ontological commitment, the one endorsed by ESR; but at a *derivative* or *metatheoretical* level of analysis, OSR correctly argues in favour of an economical view of objects, in the following way.

Consider again the SLMS scheme. The assumption is that there is no *direct* knowledge of the intrinsic nature of the entities in themselves that constitute the system under investigation. However, once a theory has ontologically committed itself to the structural properties of the system, one is entitled to infer *indirectly* that, whatever the system and its components (i.e., the objects or *relata*) may be in themselves, they must be such as to allow the theory to model at least some of their structural/relational properties. This is a transcendental feedback: what are the conditions of possibility of knowledge (the knowledge offered by the theory) of the structural properties of the system?

On the one hand, all realists agree that there is a mind-independent, external reality addressed by our knowledge, yet this is a red herring. We have seen that this ontological commitment is even poorer or more basic than the one supported by the epistemic structuralist. It is the sort of uninteresting commitment that allows a reasonable conversation

¹⁴ Of course ESR may be more or less "lite", depending on the scope of the relations taken into account, but this does not change the sort of ontological commitment in question, which still favours relations instead of relata.

among the various forms of realism (naïve, structural, scientific, even the instrumentalist, in the sense specified above, and so forth). It is not what is in question here.

On the other hand, at the metalevel level of analysis just assumed, one is exploring whether there might be a justifiable, higher LoA, that commits the theory to some kind of interpretation of the sort of objects/*relata* that support the structural properties of the system identified by the model that has been produced by the theory at the lower LoA. The question is what makes ESR possible. Obviously, the epistemic structuralist simply has nothing to offer here, for her concern is with a primary commitment. The ontic structuralist, on the contrary, can argue in favour of a minimalist approach with respect to the *relata*. And this seems a reasonable step to take. The LoA that one may justifiably adopt at this level is one that commits the theory to an interpretation of the objects/*relata* as themselves structural in nature. I shall say more about this in the next two sections, but let me first clarify in what this higher-level commitment consists.

At the first-order level (Fig. 4), we have seen that one may reasonably adopt a transcendental approach to knowledge – as a LoA-mediated epistemic relation with reality – and use Ockham’s razor as a methodological safety measure to limit the *number* of (types of) components a theory should be ontologically committed to by its LoAs. In the best circumstances, first-order LoAs should ontologically commit a theory at most to the existence of the (type of first order) structures identified by its models.

At a second-order level, one re-adopts a transcendental approach to what makes possible the previous first-order knowledge of structures. In order not to “revert to the sin of transcendental metaphysics”, as Quine [1992] nicely phrases it, one applies a qualitatively modified version of Ockham’s razor, which now suggests, still methodologically, keeping the *nature* of the objects/*relata* as simple as possible: *entia non sunt complicanda praeter necessitatem*, to use some medieval Latin. According to this new safety principle, it is reasonable to assume that 0-order *relata* are structural entities, for this is all the theory needs in order to justify its previous commitment in favour of first-order relations. A higher degree of commitment (to things/*relata*) – achieved through the reflection on the condition of possibility of the previous, minimalist commitment (to structures/*relations*) – corresponds to a lower degree of “stuff” assumed to be out there in the world (*relations* are first-order, *relata* are 0-order). Since the two ontological commitments occur at different levels there is no incompatibility and hence no objection. Fig. 4 summarises the new analysis.

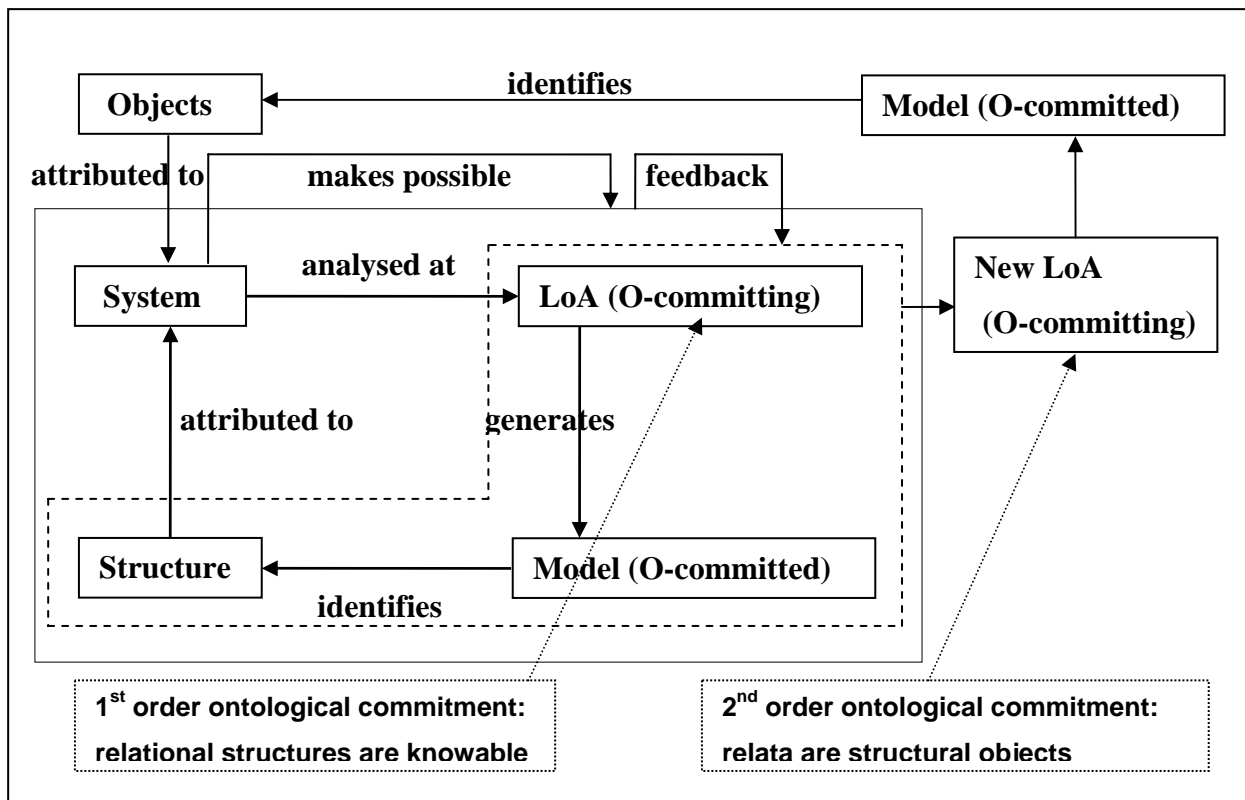


Fig. 4: the SLMS scheme with ordered ontological commitments

3. Second step: *relata* are not logically prior to all relations

The opponent of OSR may not be satisfied. Eliminating the apparent inconsistency between ESR and OSR is only a first step. It does not yet show that OSR is *plausible* because there is a direct and quite forceful objection to OSR that still needs to be neutralized. This is our second step.

Relations (structures) require *relata* (structured/able objects), which therefore cannot be further identified as relations (structures) without running into some vicious circularity or infinite regress. Yet this is precisely what OSR appears to be forced to argue, if the very idea of structural objects is supposed to make sense.¹⁵

As usual, one may bite the bullet and talk of an infinite regress of structures. Note that the point here may be not that objects may be structural, but that, even if there are objects, they are unreachable, hidden behind an infinite series of structural layers, or as the lower limit of an infinite series of structures. In either case, it may be turtles all the way down, as

¹⁵ The objection is expounded in Psillos [1999], Psillos [2001], and Psillos [2004]; its strength is fully acknowledged in French and Ladyman [2003b]; Morganti [2004] reviews some replies.

Saunders [2003] wittily puts it. This is certainly a solution. Whether it is satisfactory is another matter. Instead of tackling the problem, we run away from it, and run forever. True, we are not caught by our opponents, but we do not defeat them either. For those who do not like draws, or who run out of conceptual breath when asked to run *ad infinitum*, there is an alternative.

Admittedly, *external* relations usually require *relata*. Distance and speed are two good examples. However, *internal* relations constitute their *relata* for what they are. “Married” comes easily to one’s mind: John and Mary are husband and wife only because of their mutual relation. More formally, if an individual x has a predicate φ which is such that, by virtue of having φ , x necessarily has a relation R to at least another individual y , then R is an internal relation of x .¹⁶ Could this provide the right approach?

Not yet. Unfortunately, as E. G. Moore showed, internal relations seem to supervene on their *relata* and further qualify them. Mary and John did not come into existence by getting married; they only acquired a new contingent property, their marital status, after their wedding. If one wants to defend the logical priority of internal relations over their *relata*, then one must show much more, namely that the essential properties of the objects in question depend on some more fundamental internal properties. And this is arguably very difficult.

Difficult but not impossible. For there is a significant exception, a case that is ontologically more fundamental than the case in which the essence of the *relata* is in question. This is the (internal) relation of *difference*, which constitutes its *relata*.

The relation of difference seems a precondition for any other relation and hence for any process of knowledge. *Relata* as merely *differentiated* entities and nothing else (at least not yet) are possible only because there is a relation of initial, ontological difference in the first place. They can be understood as epistemically (still) indistinguishable but (already) ontically distinctly-existing *differentiae*, to follow French [2003]. For consider what a completely undifferentiable entity x might be. It would be one unobservable and unidentifiable at any possible LoA. Modally, this means that there would be no possible world in which x would exist. And this simply means that there is no such x . Call this the modified Leibnizian principle of *the necessary inexistence of the logically undifferentiable*. The principle still says very little about the nature of the objects/*relata* in question, but one may

¹⁶ This was the basis of the classic debate on internal vs. external relations between Bradley and Moore, with the former arguing that ultimately all relations are necessarily internal, and the latter trying to show that this position can be defended only on the basis of a confusion concerning the scope of the modalities involved.

argue that there is finally a clear sense in which *relata* are logically preceded by a fundamental internal relation of difference. Difference is our *Ur*-relation.¹⁷

Well, not so fast. Describing the relation of difference as an *Ur*-relation may make things clearer but it does not make them less controversial. On the contrary, the argument just sketched would be much weaker if it did not rely on a classic fallacy, namely an unqualified use of universal quantification. One considers Everything (Bradley would have called it the Absolute) and wonders what this Everything would be like if all differences were erased, i.e. if Everything were identical to Everything else. A similar universe would be a universe in which no individual could exist and nothing could happen, for example, although it would not have to be an empty universe. Imagine a toy universe constituted by a two-dimensional, boundless, white surface. Anything like this toy universe is a paradoxical fiction that only a sloppy use of logic can generate. For example, where is the observer in this universe? Would the toy universe include (at least distinguishable) points? Would there be distances between these points? The answers should be in the negative, for this is a universe without relations. Perhaps it is not true that there are *relata* that are logically priori to their relations, but we should be wary of the fallacious nature of any unqualified universalisation.

There seems to be some truth in the last reply, leading to a sort of truce. Nobody should quantify without restriction or qualifications, no matter whether on *relata* or on relations. At the fundamental level where *relata* appear as bare *differentiae de re*, it makes little sense to talk of logical priority. Like the two playing cards that can stand up only by supporting each other or not at all, ultimately the relation of difference and the *relata* it constitutes appear logically inseparable. Difference and the differentiated are like the two sides of the same sheet of paper: they come together or not at all. So it is more reasonable to accept that, ultimately, basic entities and structures, *relata* and relations, simply co-exist as a package: they make each other possible and one cannot choose which one to have, it is all or nothing.

So the conclusion is perhaps weaker. It may not be true that a relation of difference is logically prior to the *relata* that it differentiates. Luckily, we can leave the question unanswered because the package hypothesis is all that is needed; it is already enough for our strategy that the logical rank of the most basic relation of difference is not necessarily higher

¹⁷ Note that this is the other side of the Fregean problem of identity statements discussed in *Sinn und Bedeutung*.

than the logical rank of the most basic *relata*. It is not true that all *relata* are logically prior to relations. This much has been granted and this much suffices.

Unsurprisingly, this was also Eddington's conclusion: "The relations unite the *relata*; the *relata* are the meeting points of the relations. The one is unthinkable apart from the other. I do not think that a more general starting-point of structure could be conceived" (Eddington [1928], 230–231, quoted in French [2003], 233). It makes sense. In terms of naïve ontology, we may define roads as what connects roundabouts, yet the latter are not less real (and dangerous, or safe, and so forth) just because they are the meeting points of the roads they connect. On a more sophisticated and much smaller scale, Esfeld [2004] has built on quantum entanglement to argue in favour of "replacing [in quantum physics] a metaphysics of intrinsic properties with a [moderate] metaphysics of relations. [...] A metaphysics of relations is often dismissed out of hand, for it seems to be paradoxical. It seems that (a) relations require *relata*, that is, things which stand in the relations, and that (b) these things have to be something in themselves, that is, must have intrinsic properties over and above the relations in which they stand. However, a metaphysics of relations merely has to reject the second part of this claim: one can maintain that (a) relations require *relata*, that is, things which stand in the relations, but that (b*) these things do not have any intrinsic properties that underlie the relations in which they stand. There is nothing paradoxical in this latter claim." The idea seems to be that *relata* without relations are possible but intrinsically anonymous. I shall return to this point in the conclusion. If even a weaker position such as Eddington's – the package hypothesis – is convincing, and I think it is, OSR should be granted full plausibility.¹⁸

4. Third step: the concept of a structural object is not empty

Plausible, however, does not mean applicable. Even if ESR and OSR are compatible, and even if it is true that, ultimately, structured *relata* are not necessarily logically prior to all structural relations, one still lacks a clear grasp of what these structural objects might be like, even if indirectly or metatheoretically. What if they are merely conceivable fictions or just *impossibilia*? In either case, they would be perhaps harmless but certainly inapplicable speculations, philosophically useless.

¹⁸ For a defence of a similar moderate approach to the *relata* vs. relations issue in the philosophy of physics, see Esfeld and Lam [forthcoming], who argue that "objects and relations (structure) are on the same ontological footing, with the objects being characterized only by the relations in which they stand".

They are not. A straightforward way of making sense of these structural objects is as *informational objects*, that is, as cohering clusters of *data*, not in the alphanumeric sense of the word, but in an equally common sense of *differences de re*, i.e. mind-independent, concrete points of lack of uniformity. Some clarifications are in order.

In its simplest form, a datum can be reduced to just a lack of uniformity, that is, a binary difference, like the presence and the absence of a black dot, or a change of state, from there being no black dot at all to there being one. Admittedly, a datum is usually classified in our folk ontology of chairs and pens as the relatum exhibiting the anomaly. This is often so because the relatum exhibiting the anomaly is perceptually more conspicuous or subjectively less redundant than the other relatum, seen as a background condition. However, the relation of difference is binary and symmetric. In the example, the white sheet of paper is not just the necessary background condition for the occurrence of a black dot as a datum; it is a constitutive part of the datum itself, together with the fundamental relation of inequality that couples it with the dot. In this specific sense, nothing is a datum *per se*, without its counterpart, just as nobody can be a wife without there being a husband. It takes two to make a datum. So, ontologically, data (as still unqualified, concrete points of lack of uniformity) are purely relational entities. Of course, from a structural perspective, they remain unknowable in themselves. In this technical sense, one never deals epistemically with pure data but only with somewhat interpreted data (in semantics, this is often referred to as *content*, see Floridi [2005b]). Now, cohering clusters of data as relational entities (differences *de re*) are the elementary *relata* we are looking for in our modified version of OSR. Thus, the structuralism in question here is based on relational entities (understood structurally) that are *particular*, not on “patterns” that are abstract and universal (this distinction is carefully analysed by Reck and Price [2000]; it is parallel to the difference between Poincaré’s and Russell’s approaches to structures respectively). These relational entities are unknowable not because they are somehow mysteriously unreachable, but because their epistemic malleability is *inexhaustible* (I shall return to this point in section 7).

All this is not entirely a novelty. As French [2001] points out, in the context of the philosophy of physics, “if we want to continue to talk, in everyday language, about electrons as objects - because we lack the logico-linguistic resources to do otherwise – then we can do so ‘only indirectly’, ‘... not insofar as they themselves, as individuals, are given, but so far as they are describable as ‘points of intersection’ of certain relations’ (Cassirer, *ibid.*). And this relational conception of an object is taken straight from Kant himself: ‘All we know in matter

is merely relations ... but among these relations some are self-subsistent and permanent, and through these we are given a determinate object.’ (Kant, quoted on p. 182)”.

The informational view of structural objects can be further clarified by considering the following objection. What could an informational version of OSR look like, not just in some logically possible way, but “metaphysically seriously”, as Bloomfield [2005] quite rightly requires when chastising Chalmers’ modal laxity? After all, if informational/structural objects are mere logical *possibilia* this explains why one can do some fancy mathematics with them, but not much more. However, serious metaphysics should be *applicable* to the real world and *scalable* (more on this shortly) from sub-observables to observables, that is, it needs to be *usable* when dealing with the macroworld of everyday life and experience, so that it may become *useful* in other philosophical contexts, not only in epistemology, but also in ethics, in the philosophy of mind, or in the philosophy of language. In a nutshell, the objection is that an ontology of structural objects as informational entities may be logically viable, but seems to lack both seriousness and scalability (see below).

Answering this twofold objection requires a clarification, and I hope the reader won’t mind if we borrow some more technical vocabulary from computer science.

We have seen that an ontology is a conceptualization of a system under analysis, made possible (the conceptualization, that is) by a LoA, and that it is embedded in the model of the system. An ontology may satisfy different specifications, such as experiential or commonsensical realism and resilience to scientific tests and discoveries. In particular, from a methodological point of view, the *generality* of an ontology is a function of its *portability*, *scalability* and *interoperability*.¹⁹

Portability refers to the ease with which a piece of software or a file format can be “ported”, i.e., made to run on a new platform and/or to compile with a new compiler. For example, assembly code is the least portable type of language, since it is specific to one particular (family of) processor(s). Using this concept for our theoretical purposes, it is clear that the lower the LoA, the less portable the corresponding model is. This means that an ontology is increasingly portable the more general it can be. The most portable ontology would be one that could be made to “run” on any possible world. This is what Aristotle meant by a general metaphysics of being *qua* being. The portability of an ontology is a function of its importability and exportability between theories even when they are disjoined (the LoAs

¹⁹ The following definitions are standard and only slightly adapted from FOLDOC, <http://foldoc.doc.ic.ac.uk/>

have no observables in common). Imagine an ontology that successfully accounts for the natural numbers and for natural kinds. I shall return to this point in section 6, when discussing objection 10.

Scalability refers to how well a solution to some problem will work when the complexity or magnitude of the problem increases. In our context, the scalability of an ontology is a function of its successful applicability independently of the vertical granularity and the horizontal scope of hierarchically nested or overlapping theories (the LoAs share at least one typed variable) and their corresponding systems. Imagine an ontology that successfully accounts not only for Schrödinger's cat but also for the atomic particles dangerously decaying in its proximity.

Finally, *interoperability* refers to the ability of software and hardware on multiple machines from multiple vendors to communicate and interact successfully. *Mutatis mutandis*, the interoperability of an ontology is a function of its capacity of allowing interactions between different theories. Imagine an ontology that successfully accounts for a system modelled as a brain and as a mind.

Obviously, these three specifications are interdependent. High portability, scalability and interoperability guarantee increasing context-independence and seamless integration of models into a unified world-view. An ontology that is highly portable, scalable and interoperable is an ontology that can be applied to a variety of possible worlds, at a variety of degrees of complexity and that allows a variety of theories that adopt it to interact with each other successfully. We are now ready to define a specific metaphysics as an ontology with fixed degrees of portability, scalability and interoperability: there is Descartes' metaphysics and Spinoza's, Leibniz's and Hume's, David Lewis' and Armstrong's, and each will have its own degrees of portability, scalability and interoperability.

It has become common to acknowledge that both physics (Saunders [2003], Chakravartty [2003]) and mathematics (Brink and Rewitzky [2002], Shapiro [2004]) underdetermine metaphysics. This may be cause for despair in some skeptical quarters, for some sort of "anything goes" thesis may seem inevitable; but it may also be interpreted, more optimistically, as a green light for a purely philosophical approach. I share the latter view. Metaphysical questions are intrinsically open (cannot be answered either by experiments or by calculations, to put it bluntly, they are a matter of rational exchange of arguments) and, as far as structural realism is concerned, evaluating the ontological commitment of a theory is a philosophical business. So let us return to the general characterization of a metaphysical

position. A metaphysics may be criticised, philosophically, for being *local* whenever its degrees of portability, scalability and interoperability are just local maxima. For example, a Cartesian metaphysics is notoriously undermined by its poor degree of interoperability: the mind/body dualism generates a mechanistic physics and a non-materialist philosophy of mind that do not interact very well. Leibniz's metaphysics of monads is not easily scalable (it is hard to account for physical macro-objects in its terms). Indeed, a metaphysics may be so local as to be peripheral, i.e. usable only within a possible world that is constituted by it. Imagine the ontology of *The Lord of the Rings*.

Now, as a metaphysics, one of the greatest strengths of OSR is its portability between physical and mathematical theories. This is to be expected, given its connection with the philosophy of quantum physics. Mathematics is the common language that helps to explain structural objects group-theoretically, as in Weyl and Cassirer, for example, or set-theoretically – in terms of partial structures, as in Ladyman [1998] and in Da Costa and French [2003], or of quasi-structures, as in Krause [2005] – or categorically, as suggested again by Krause [2005]. To use Cassirer's phrase, it has been a long time since mathematics moved from substance to function.

Moreover, the portability of OSR between physical and mathematical theories helps to explain why mathematics is so successfully applicable to a structurally-interpreted physical world. Indeed, the degree of portability is such that even the problems are similar. As Reck and Price [2000] have highlighted, there is a variety of structuralisms in the philosophy of mathematics and Carter [forthcoming] shows that some of the basic issues arising in the development of structuralist philosophies of physics, especially those concerning the ontology of individuality and the semantic of re-identification, bubble up in the structuralist philosophies of mathematics as well.

All this reinforces the point that an ontology of structural objects is possible (read: logically consistent), but it does not yet remove all doubts about its seriousness, since OSR may still be metaphysically too local and lack scalability.

This is where an informational approach turns out to be a further strength. An ontology of structural objects is portable to computer science as well, and while mathematics guarantees consistency, computer science facilitates scalability. Simply put, micro- and macro-objects are analysable as informational objects that satisfy OSR specifications. To see more precisely how, let us turn to an instructive example provided by the methodology of *Object Oriented Programming* (OOP; Rumbaugh [1991]).

OOP is a method of programming that changed the approach to software development. A program used to be viewed as an algorithmic procedure that takes input data, processes it, and produces output data. The difficulty was then represented by the elaboration of the algorithmic process. OOP shifted the focus from the logic procedures, required to manipulate the objects, to the objects that need to be manipulated. Consider a pawn in a chess game. Its identity is not determined by its contingent properties as a physical body, including its shape and colour. Rather, a pawn is a well-defined cluster of specific *states* (properties like white or black, and its strategic position on the board) and determined *behavioural rules* (it can move forward only one square at a time, but with the option of two squares on the first move; it can capture other pieces only by a diagonal, forward move; and it can be promoted to any piece except a king when it reaches the opposite side of the board), which in turn are possible only in relation to other pieces and the logical space of the board. For a player, the actual pawn is only a placeholder, whereas the real pawn is an “informational object”. It is not a material thing but a set of typed variables, using the LoA terminology, or a mental entity, to put it in Berkeley’s terms, or an entity constituted by a bundle of properties, to use a Humean expression. Its existence and nature is determined by the differences and nomological relations that characterise the game of chess. The physical placeholder can be replaced by a cork without any semantic loss at the LoA required by the game. Indeed, a player may not even need a physical placeholder at all. Now in OOP, *data structures* (e.g. the pawn’s property of being white) and their *behaviour* (programming code, e.g. the pawn’s power to capture pieces only by moving diagonally forward) are packaged together as informational objects. Discrete objects are self-contained collections of data structures and computational procedures. They are grouped in a hierarchy of classes (e.g. pawns), with each class inheriting characteristics from the class above it (e.g. all pieces but the king can be captured, so every pawn can be captured). A class is a named representation for an abstraction, where an abstraction is a named collection of attributes and behaviour relevant to modelling a given entity for some particular purpose at a certain LoA. The routines or logic sequences that can manipulate the objects are called *methods*. A method is a particular implementation of an operation, i.e. an action or transformation that an object performs or is subject to by a certain class. Objects communicate with each other through well-defined interfaces called *messages*.

There is no need to add further details. Clearly, OOP provides us with a rich description of informational objects that can be used to *conceptualize* a structural object as a combination of both data structure and behaviour in a single entity, and a system as a

collection of discrete structural objects. Given the flexibility of the conceptualization, it becomes perfectly possible, indeed easy, to approach the macroworld of everyday experience in a structural-informational way.

All this does not mean that we can overstep the limits of our structural knowledge of reality set by ESR. Recall that all we are assuming to be in the world are informational objects as the conditions of possibility of those structures that our first-order LoAs allow us to know and to attribute to the world in the first place. Thus, OOP is not a philosophical ontology but a practical example of a valuable methodology that can clarify the nature of our ontological components and prove their scalability.

OOP is a good example of how the concept of informational object provides a useful way to conceptualise precisely the relata described as structural objects by OSR. Once again, it is reassuring to see that all this too is not a complete novelty. Dennett [1991], for example, argues that macro-objects are patterns and that the existence of a pattern as a real object depends on the usefulness of the theories – especially their explanatory power and reliability in making predictions – that admit that pattern in their ontology. One may say that Informational Structural Realism introduces a more flexible and powerful methodology to clarify and make precise Dennett’s “pattern-like” objects.²⁰ A computational approach (e.g. the methodology of Object Oriented Programming) can do for the scalability of a structural ontology what category or group theory can do for its portability. So ISR is perfectly applicable.

5. Informational structural realism

The time has come to summarise the proposed solution to the ontological challenge to SR introduced in section 1.

Informational structural realism (ISR) is a version of SR. As a form of realism, ISR is (0-order) committed to the existence of a mind-independent reality addressed by and constraining our knowledge. Like ESR, ISR supports LoAs that carry a minimal ontological commitment in favour of the structural properties of reality. Like OSR, ISR also supports LoAs that carry a reflective, equally minimal, ontological commitment in favour of structural objects. Unlike ESR and OSR, ISR supports an informational interpretation of these structural objects. The second commitment (in favour of structural relata) is justified by epistemic

²⁰ I owe this insight to Wallace [2003].

reasons. We are allowed to commit ourselves ontologically to whatever minimal conception of objects is useful to make sense of our first commitment in favour of structures. The first commitment answers the question “what can we know?”; the second commitment answers the question “what can we safely assume to be there?”. We are now ready for a definition:

ISR) Explanatorily, instrumentally and predictively successful models (especially, but not only, those propounded by scientific theories) at a given LoA can be, in the best circumstances, increasingly informative about the relations that obtain between the (possibly sub-observable) informational objects that constitute the system under investigation (through the observable phenomena).

A significant consequence of ISR is that, as far as we can tell, the ultimate nature of reality is informational, that is, it makes sense to adopt LoAs that commit our theories to a view of reality as mind-independent and constituted by structural objects that are neither substantial nor material (they might well be, but we have no need to suppose them to be so) but informational. This allows IRS to reconcile two metaphysical views in the philosophy of information. One is Wiener’s, who thought that “information is information, not matter or energy. No materialism which does not admit this can survive at the present day” (Wiener [1961], 132). The other is Wheeler’s, who proposed the “it from bit” thesis: “It from bit. Otherwise put, every ‘it’ – every particle, every field of force, even the space-time continuum itself – derives its function, its meaning, its very existence entirely – even if in some contexts indirectly – from the apparatus-elicited answers to yes-or-no questions, binary choices, bits. ‘It from bit’ symbolizes the idea that every item of the physical world has at bottom – a very deep bottom, in most instances – an immaterial source and explanation; that which we call reality arises in the last analysis from the posing of yes-no questions and the registering of equipment-evoked responses; in short, that all things physical are information-theoretic in origin and that this is a participatory universe.” (Wheeler (1990), 5).

When Cassirer talked about structuralism, he had in mind, like Kant and Russell earlier and Maxwell later, a full-blooded ontology of objects as structural entities (see for example the quotation from French in the previous section). ISR seeks to fulfil this role. Of course, if an ontology of informational objects is scalable to the macroworld, then it should be applicable to the folk ontology of cars and chairs and usable within other philosophical contexts. There is no space to show here how well ISR can pass this test, but the interested reader is referred to some initial work in epistemology (Floridi [2004]), computer ethics

(Floridi [2003a]) and the philosophy of mind (Floridi [2005a]), where ISR has already been shown to be a flexible and fruitful approach.

Second Part

6. Ten objections and replies

The case in favour of ISR can now be reinforced by clearing the ground of some potential perplexities and objections.²¹

6.1) One might wonder whether reconciling ESR and OSR really has any purpose. The objection is that ESR may be an appealing position, but it faces a serious objection (articulated by Psillos and Papineau), namely

- i) if restriction of beliefs to structural claims is really no restriction at all, then
- ii) ESR collapses into full-blown realism; therefore
- iii) ESR cannot provide a notion of continuity in science that avoids the pessimistic meta-induction; however
- iv) we want to retain the notion that structure is what is preserved successfully across otherwise radical theory change; so
- v) we had better move to a notion of ontological structural realism in which structures directly represent reality; and we can be right about those structures. But then,
- vi) from the beginning OSR conceded that the argument against ESR was compelling and sought to develop a notion of realism about structure that did not fall prey to the same argument; so
- vii) why should OSR be interested in reconciling its view with ESR, when it has already been conceded that the latter is untenable?
- viii) Conversely, if it is felt that adopting ESR is inconsistent with adopting OSR, then one first needs to overcome the objection that made it necessary to move to OSR in the first place.

The answer is that the objection is valid but unsound, as it starts from a false premise:

- 1) in order to hold the inference (i) \rightarrow (ii), the objection needs to be serious about “collapse”, that is, it must mean that ESR is at least logically equivalent to full-blown realism (FBR), so that (ii) = (ESR \leftrightarrow FBR), which means that
- 2) (i) \rightarrow (ESR \leftrightarrow FBR); but then
- 3) if (ESR \leftrightarrow FBR) is false, then

²¹ In this section I have summarised several questions and objections that I have received during and after the conference, when the paper circulated among several colleagues. I take full responsibility for their specific formulation, although many are as close as possible to the original wording.

- a) either (i) is true, in which case the inference in (2) is invalid and anything can trivially follow from it;
- b) or (i) is false, but then anything can trivially follow from (i) itself;

in either case, the whole argument is valid but unsound. So the point becomes considering whether $(ESR \leftrightarrow FBR)$ is false.

It is. If properly constructed, ESR cannot become equivalent to FBR both for theoretical reasons – for example, ESR rejects the possibility of knowledge of the intrinsic nature of objects in themselves, which is supported by FBR, so the two forms of realism have different scopes, with only ESR implementing “Kantian humility” – and for metatheoretical reasons – since, if ESR could be reduced to FBR, then, *mutatis mutandis*, objections such as Newman’s could be extended to FBR as well, an obvious *reductio*.

Once $(i) \rightarrow (ii)$ is dismissed, it becomes clear that (iii) cannot be held, since it depends on (ii). As for the inference $(iv) \rightarrow (v)$, it may be sound, but it is also uninteresting unless it is further constrained by (iii), that is, unless (iii) were true and hence $((iii) \wedge (iv)) \rightarrow (v)$ were the case non-trivially, but we have seen that (iii) cannot be granted.

A supporter of the objection may retort that (vi)-(viii) still apply, independently of whether they follow from the premises, which could be mended or replaced.

The reply is that, once we take a more sceptical view of (2), the rest of the alleged reconstruction of the original dialectic between ESR and OSR becomes historically inaccurate and conceptually suspicious.

Regarding the historical inaccuracy, one needs to remember that OSR has emerged within the context of the debate on individuality in quantum mechanics (French and Redhead [1988]) and that it is actually eliminativist OSR that runs the risk of becoming indistinguishable from a stronger form of scientific realism, which may then be unable to cope with the *Pessimistic Meta-Induction Argument* (see sections 1 and 2.1). So point (vi) cannot be granted as a faithful description of the current internal debate between ESR and OSR (see for example Morganti [2004]). Such debate is much more reliably interpreted by e.g. Ladyman [1998] and Votsis [forthcoming] in terms of a discussion concerning the ontological implications of SR. Moreover, even if (vi) were actually granted *de facto* in the literature, no structural realist should grant it *de jure*, because we have seen that $(ESR \leftrightarrow FBR)$ is unjustified, and it appears even more evidently so from a SR-friendly perspective. Concerning (vii), there are no more reasons to raise the rhetorical question, since ESR is far from

untenable, given that the inference in (2) is either invalid or valid but trivial. As for (viii), the whole point of section 2 is to show that ESR and OSR are not mutually inconsistent but may be reconciled, a point that now may be better appreciated as quite interesting. The objection should address not the value of the endeavour but, in this case, the validity of the arguments.

Regarding the suspicious nature of the objection, this is the interesting aspect mentioned above. The objection (mistakenly) assumes (2) to be uncontroversial in order to build a Trojan horse, as we have seen in section 1.1. However, the goal of the previous reconstruction is not really to wonder whether reconciling ESR and OSR is to any purpose – from an SR-friendly perspective it obviously is, witness the current debate – but rather to move an implicit attack against SR as a whole: for if ESR is untenable, then SR can be reduced to OSR, but then SR can be more easily defeated by showing that OSR is equally untenable because of its metaphysical consequences. Now, from a structuralist perspective, ESR is the strongest position, when it comes to resisting objections from without; but OSR is the most satisfactory, when it comes to settling internal debates. As I made explicit at the beginning, ISR is not an attempt to support SR directly, but to show that supporters of SR could successfully combine, by means of ISR, the strengths of both ESR and OSR in order to counterbalance any *divide et impera* strategy.

6.2) It may be objected that the definition of ESR, offered in section 1, is inconsistent with e.g. Grover Maxwell's conception of structural realism, for the latter would allow one-place predicates qualifying objects, if they were suitably perceptual, that is, it is only theoretical terms with dubious status – for the empiricist – that have to be replaced in the Ramsey sentence.

This objection provides the opportunity to clarify two points. First, the definition of ESR in section 1 is explicitly based on the neo-Kantian versions of SR endorsed by Cassirer, Russell and Poincaré, among others. The approach is arguably justified. Bitbol [2002], French [2001], Gower [2000] and Ryckman [2003], for example, provide well-supported reconstructions of structural realism from a transcendental perspective. The latter equates the nature of things-in-themselves with the first-order properties (monadic or unary relations) qualifying noumena. This is why the definition explicitly specifies that by “*first-order, one-place predicates qualifying the objects in themselves*” one is referring to “*the intrinsic nature of the noumena*” (note that the objection overlooks all the specifications in italics). In other words, the relations that are said to be unknowable, in the definition of ESR, are not the n-

places relations subject to Ramseyfication mentioned by the objection, which are perfectly acceptable. This is the second point. The relations mentioned by the objection are the equivalent of phenomenal relations, in Kantian terms. So the definition is consistent with Maxwell's version of SR as well. This will become even more obvious in the reply to the next objection.

6.3) The previous objection leads to a further concern regarding the mixing of issues specific to philosophy of science and the epistemology of scientific theories with general issues in epistemology and metaphysics. This seems an unjustified confusion. For example, someone who may have genuine worries specifically about the epistemology of scientific theories may have no worries about epistemology in the broad sense, considering that it is only when things become unobservable that we begin to run into difficulties, and this seems a perfectly respectable position.

The objection is correct insofar as it argues that issues specific to the different fields in question do not have to be brought together, but it is mistaken in inferring that therefore they cannot or should not be combined, and that doing so may not be helpful but only confusing. Threading together issues in different contexts is not a problem *per se*. On the contrary, in the previous pages the strategy has been intentionally adopted in order to provide interdisciplinary depth and strength to the case for ISR. One can cast light on problems that have several dimensions only by analysing them from different perspectives, while leaving anyone free to maintain a more localised position, should she wish to do so. In the reply to objection (6.2), for example, we have seen that it is important to understand carefully when one is speaking about relations as unknowable predicates of noumena/relata or as ordinary predicates in the language of a scientific theory or in any ordinary experience. In this respect, the new example provided by objection (6.3) is equally enlightening. From a Kantian perspective, it is simply not true that we begin to run into philosophical difficulties only when unobservables are in question, not least because the scope of "unobservable" is much wider. It is a scope inherited by several structural realists. For example, following Russell's version of structural realism, and hence assuming a Neo-Kantian perspective, Maxwell [1968] remarks that "all of the external world, including even our own bodies, is unobserved and unobservable" (152). This reinforces the reply to the second objection, if necessary: Maxwell's conception of SR allows one-place predicates qualifying objects only insofar as the objects are not unobservable in the strong sense just seen, not in the sense that they are

suitably perceptual, as alleged by the objection. Poincaré and Worrall endorse the same perspective: unobservable means noumenal, that is, epistemically inaccessible.²² One may disagree with them. Indeed, throughout the article I have been careful not to adopt the same conceptual vocabulary, preferring to use sub-observable to refer to the usual unobservables as ordinarily understood in philosophy of science (this is also the reason why it is better to keep the comment on Maxwell’s understanding of “observable” separate from the previous objection). Yet one must also concede that, following their line of reasoning, it is natural to merge issues in epistemology with issues in the philosophy of science, metaphysical issues with issues concerning the ontological commitments of scientific theories. This is the Kantian approach that has been intentionally followed in this paper as well, where scientific knowledge is taken to be only knowledge of a special kind.

6.4) A slightly different objection, related to the previous one, concerns the use of the term “instrumentalism”, which seems unclear. At points, instrumentalism is said to be a form of realism, but that is flatly incompatible with the notion of realism that is pertinent in philosophy of science (realism about scientific theories or entities). Evidently, a more general sense of realism is thus supposed to be in use.

The objection is correct: instrumentalism is said to be a form of realism in three cases, and in each case what is at stake is a very basic, minimal form of realism. However, this sense is neither unclear nor unusual. An instrumentalist may wish to claim that she is a realist insofar as she accepts “the ontological thesis” of realism, namely that there is a mind-independent world. That this is a fairly common way of looking at instrumentalism is shown, for example, by Psillos [2000b], who, in a carefully argued defence of scientific realism, suggests, correctly, that what distinguishes scientific realism from instrumentalism is not “The Metaphysical Thesis: the world has a definite and mind-independent structure” (which one must assume they both share, though Psillos is not explicit about this) but rather “The Semantic Thesis: scientific theories should be taken at face-value”.

6.5) The “Ockham” motivation for the move to structural objects at the second level is a crucial part of the argument in section 2.4, but it needs stronger motivation. The first level LoA is agnostic about the *relata* of its relations – it says that they are just things standing in

²² On this point see Votsis [forthcoming].

some relations. Why is it then *simpler* to construe them structurally? At the first level, we have objects bearing relational properties; it is not clear that objects with monadic properties are a different *type* of objects, or that monadic predicates are more complicated than relational predicates.

The objection can be answered by dispelling two potential confusions on which it is based. The first is that, by talking of first- vs. second-order LoAs one may be talking about first- vs. second-order predicates (monadic or not). This is not the case, and indeed in 2.2 the precise formulation says “*derivative or metatheoretical*”. The second-order LoA, at which one may argue in favour of a structural understanding of *relata*, is a meta-level: it is related to the first-order LoA as a sentence *p* in English may be related to one in Italian that *p* is explaining. It has nothing to do with the sense in which first- and second-order is (justifiably but inappropriately) used in the objection. So the issue here is not monadic vs. non-monadic predicates. The “Ockhamist strategy” is first applied at the “object-level” (nothing to do with “level of objects”, of course), where ESR is shown to provide the best alternative, and then is re-applied at a meta-level, where OSR is shown to provide the best alternative because it is the one that assumes the less demanding ontology for the *relata* in question. The distinction clarifies why the argument is explicitly described as *transcendental*, a label that would make no sense if first- and second-order were understood as the objection suggests, instead of levels of reflection.

The second potential confusion concerns the value of the “Ockhamist strategy”. When, by reflecting on the conditions of possibility of ESR, we come to consider the possibility of OSR, we still do so economically, but the nature of the question has changed by moving up one level of reflection: instead of asking “what can we know?”, a question answered by ESR, we now ask “what is the least that we can safely assume to be there in order for our acceptance of ESR to make sense?”. The latter is a question that pertains to OSR, which offers a structuralist analysis of entities as the least demanding approach consistent with the acceptance of ESR. Both strategies adopt Ockham’s razor, but one in order to constrain one’s commitment to the *number* of types of things in the world that are knowable (the doxastic scope discussed in the reply to objection 6.1), and hence defend ESR; the other in order to constrain one’s commitment to the possible *nature* of things in the world that would be consistent with the previous commitment, and hence defend OSR as supporting ESR.

6.6) Concerning the modified Leibnizian principle of the necessary inexistence of the logically undifferentiable, isn't there another counter-example to this proposed principle? What if we consider the world as a whole? There is nothing for it to be differentiated from, but it still exists.

The objection is correct in pointing out that a principle that held the necessary inexistence of the *undifferentiated* would be irreparably mistaken, for the counterexample would indeed apply. Yet this is irrelevant, for the principle endorsed in the paper (as the objection itself indicates) holds the necessary inexistence of the *logically undifferentiable*. So the objection does not apply, since a world W as a whole would still be logically differentiable, even if it were actually undifferentiated, and it would be so not only with respect to any of its parts, but also to its negative counterpart, to a whole $\neg W$. The principle says that if something is logically undifferentiable, i.e. if it is logically impossible to differentiate it from anything else, then that something does not exist. This seems as safe (and perhaps uncontroversial) as the indiscernibility of the identicals, not as controversial as the identity of the indiscernibles.

6.7) Concerning the toy universe, why can't we think of the white surface itself as an individual?

Perhaps we could, but so far nobody has been able to solve the problem of finding a principle of identification for something that, by definition, is undifferentiated, like our white surface. This is one of the main arguments used by Quine against second-order logic: we have no right to treat e.g. colours as individuals because we have no criteria to identify them as such. Things are even worse in our toy universe, where there is not even an external "place" from which an observer can wonder about the identity criterion. So the reply is: no individuality without identifiability, and the toy universe is by definition constructed as unidentifiable.

6.8) The discussion concerning the concept of datum is problematic. We are told that "[A] datum can be reduced to just a lack of uniformity, that is, a binary difference". Now, it seems that two distinct ideas are being run together here. First, a datum involves (i) a difference between (at least) two possible states of an entity and (ii) a background condition of readability. So a datum involves distinguishable (binary) differences, not just difference per se. But if so, it looks less clear that data will be the suitable entities for basic structural entities

in OSR. The notion of datum outlined in the paper seems to overplay the condition of readability as constituting the datum; this cannot be right.

The objection presents an epistemological conception of datum, based on “readability”, which would be acceptable in other contexts, but that is inappropriate here in order to understand the argument developed in favour of OSR and its basic structural entities. The concept of data employed here is as mere *differentiae de re*, that is, mind-independent, concrete points of lack of uniformity in the fabric of being. These *differentiated* entities are epistemically (still) virgin but (already) ontically distinctly-existing. It is not a matter of vocabulary. Clearly, we are dealing with a stage that logically precedes the one described by the objection. The *relata/data*, as structural entities, do not yet require any background condition to be read, for they do not need to be read at all in order to subsist. If they did, OSR would be a version of absolute idealism, that is, it would reduce the existence of the external world to its epistemic condition of “readability”, something that is explicitly rejected. Indeed, at this initial stage, we do not even know whether the *relata* as differences *de re* may be so synchronically, in which case one could imagine them as two entities, or diachronically so, in which case one could imagine them as two states of the same entity. We can only guess, for the assumption is that these *data/relata* are noumenal, and we do not have access to them. This sense of datum as difference or inequality *de re* provides the condition of possibility of the sort of epistemological datum that is then described by the objection, but that is a further step. A step criticised by the following objection.

6.9) If a datum is a difference, then a datum is an abstract thing, as a difference is an abstract, rather than a concrete thing. But this raises a severe problem, if data are supposed to be the basic objects in our structural realism (assuming that this is supposed to be structural realism as a possible position in philosophy of science). For the structural objects need to be *concrete* things in the world (even if they are things structurally conceived), as they are the things that other concrete things are made from. You cannot make concrete things out of abstract things, so informational objects do not seem to be viable candidates for the objects in an ontological structural realism fit for the philosophy of science.

The reply to the objection can be articulated in three steps. First, it is true that no concrete things can come out of purely abstract things, at least not without presupposing some metaphysical superpower that science (and its philosophy) had better leave alone, if they can. However, no reason has been offered to justify the view that data, understood as *differentiae*

de re, may not be as concrete as one's definition of "concrete" may allow. Indeed, we have seen that objection (6.8) goes exactly in the opposite direction, for it stresses that it is *insufficient* to consider data as concrete differences, since, still according to objection (6.8), their conception must also include an epistemic component ("readability"). Now, the correct position is somewhere in between: as far as the argument in favour of ISR is concerned, data are neither purely epistemic (abstract) entities, as suggested by objection (6.9), nor ontic (concrete) *differentiae de re* inseparably coupled to some epistemic component, as suggested by objection (6.8). They are (or need to be treated as) ontic (concrete) differences that are *then* epistemically exploitable as resources, by agents like us, for their cognitive processes. After all, one should recall that the data/differences in question can be concrete because we do not have to assume something as radical and problematic as Leibniz's conception of monads: Eddington's package hypothesis (Ur-relata + Ur-relations) is sufficient to support OSR.

Second, it should therefore be clear that the interpretation of structural objects as informational objects is not meant to replace an ontology of concrete things with one of virtual entities. As specified, OOP provides us with an interesting example of how we may *conceptualise* structural objects and make sense of their ontology.

Third, by talking of concrete *differentiae de re*, and conceptualising them as data structures and hence as informational objects, we are defending a version of structural realism that supports at least an irreducible, fundamental dualism, if not a Leibnizian pluralism, as the correct description of the ultimate nature of reality.

6.10) It is difficult to relate the comments on ontological commitment in section 2.3 to the comments in section 4. If a type of ontological commitment is associated with a particular choice of Level of Abstraction, then how is it possible to port ontology from one Level of Abstraction to a disjoint one? Moreover, the treatment of informational objects and ontology makes it hard to see what work the appeal to Object Oriented Programming is doing.

The objection raises several issues. First, the connection between 2.3 and 4 is the following: it is possible to reconcile ESR and OSR by showing that they work at different LoAs, but OSR still faces the problem of providing an ontology that may be applied both to sub-observable and to observable objects, if possible. In 4, it is shown that, by translating structural objects into informational objects, this can be achieved in a way that is both efficient and elegant.

Second, the same question about portability could be asked in computer science, where the answer is that a piece of software p is portable from platform A to platform B if there is a compiler for p both for A and for B , even when A and B are completely different (disjoint). An ontology is portable from a LoA_x to another LoA_y (even when they are disjoint) if there is one or more LoA_z that can translate the ontology from one to the other. For example, it has been argued that Kripke's view of rigid designators is based on an ontology that is "portable" to Aristotle's view of substances.

Third, the appeal to OOP is that it provides an intuitive example of a well-defined method to clarify how one may conceptualise structural objects as informational objects. A unified ontology, scalable from the sub-observable to the observable, has an added value, with respect to localist ontologies: it does not have to assume some radical dichotomy between fields of knowledge. We have rejected the Greek (large scale) distinction between sublunar and celestial physics/ontology, and it would be better not to reintroduce it at the (small scale) level of observable vs. unobservable worlds. ISR does support a Kantian-style dualism between knowable and unknowable, and at least a dualist if not a pluralist metaphysics, but not a dualism in terms of ultimately irreconcilable ontologies for utterly different fields of knowledge.

Finally, a further general advantage of an informational approach to structural realism is that one is able to practice what one preaches, that is, use informational concepts to develop and defend an informational approach.

7. Conclusion

At least since Plato's images of the line and of the cave, philosophers have often relied on spatial analogies to explain their theories. References to rooms are particularly popular. Sextus Empiricus thought that we are like people in a dark room, searching for gold or shooting at a target: no matter how long the search or the shooting proceeds, it is pointless because in principle there is no way to establish whether any of us has found a nugget or hit the mark (*Outlines of Pyrrhonism* I.52 and II.325). Turing used different rooms for his famous test. And Searle devised a Chinese room for his counterexample. I shall rely on their examples and suggest a double box analogy to illustrate ISR. But first, a final bit of terminology.

In software engineering, *black-box* refers to a test-design method that focuses on testing functional or behavioral requirements of a program. The methodology treats the *analysandum* as a completely opaque and closed system, avoiding using explicit knowledge

of its internal nature or structure to understand the way it works. The opposite methodology is known as *white-box* test design. This allows one to “look inside” the system, and it focuses on using specific and detailed knowledge of the program code to guide the selection of test data. A *grey-box* approach is one that allows only a partial view of the internal properties of the system.

According to ISR, any white-box approach to reality is excluded in principle, given the LoA-mediated nature of knowledge. Forget about getting out of Plato’s cave. There is no God’s eye perspective from without. We look at the world as if we were in Sextus’ dark room. This is the first box. We are inside it, but our goal is not mimetic, nor are our (often *causal*) *interactions* with the furniture in the room unidirectional, as Sextus assumed. Unlike Sextus’, ours is only a grey-box. In the best cases, it allows the reconstruction of the structural properties relating the furniture of the room, i.e., our informational objects. These are our second kind of boxes. As in Turing’s test, they are black-boxes, not directly knowable, but “epistemically interactable” through LoAs. Sometimes, we can indirectly capture their nature by observing their behaviour and mutual interactions, but we do not know their intrinsic properties. How we relate them and use them to build other black-boxes is our responsibility. This is the right context in which to talk about a demiurgic power (Floridi and Sanders [2004a]). ISR takes our epistemic goal to be *constructionist* (mind, not *constructivist* in any psychologistic or “sociologistic” sense), not mimetic. Knowledge is not a matter of either (a) discovering and describing, or (b) inventing and constructing, but of (c) designing and modelling reality, its features and behaviours into a meaningful world as we experience it. And one may design and model successfully even in the dark. Intelligibility is the name of the epistemic game, and humanity tries to achieve it at any cost, even when this means distorting reality in the most absurd way, from the conception of a flat earth placed at the centre of the universe to the interpretation of natural forces and events as anthropomorphic divinities or to the assumption of calories, phlogiston and luminiferous ether. Since we wish to devise an intelligible conceptual environment for ourselves, we do so not by trying to picture or photocopy whatever is in the room (mimetic epistemology),²³ but by interacting with it as a *resource* for our semantic tasks, interrogating it through experience, tests and experiments. Reality in itself is not a source but a resource for knowledge. Structural objects (clusters of

²³ Suárez [2003] provides a series of arguments and references against what I have called here mimetic epistemology.

data as relational entities) work epistemologically like *constraining affordances*: they allow or invite certain constructs (they are *affordances* for the information system that elaborates them) and resist or impede some others (they are *constraints* for the same system), depending on the interaction with, and the nature of, the information system that processes them. They are exploitable by a theory, at a given LoA, as input of adequate queries to produce information (the model) as output. This epistemic malleability of reality as a resource seems to be what Chakravartty [2003] defines as the “dispositional nature” of structural objects²⁴ and Saunders [2003] calls their “heuristic plasticity”. Transforming constraining affordances into information need not be a metaphysically violent business (as Bacon thought it might), if reality in itself is indeed indeterminate (Rosen and Smith [2004]) or if we are ready to be led by it insofar as it is determinate. From this perspective, semantic concerns (most importantly reference, representation and truth²⁵) belong to the relation among models, that is, among outcomes of LoAs (Kant’s phenomenal world of experience), not to the relation between models and reality in itself.²⁶

It turns out that we are like Turing’s interrogator, since the model of investigation is erotetic: we have indirect (LoA-mediated) access to reality and can query it as a database. Bacon and Galilei shared a similar view. But since our task is not to find out who is who, we resemble Searle in his Chinese room: we get the input on one side and output information on the other. The difference, in this case, is that we have some understanding of the rules of the game. It makes little sense to ask whether the factual information we gain resembles its source. The Parthenon is as concrete and objective as anyone may wish it to be, but it does not represent marble. Knowing reality is interpreting it constructively, not portraying it passively.

So the basic idea behind ISR is quite simple: black-boxes inside a grey-box. The last specification to be added is that these qualifications are LoA-dependent, in the same way as the distinction between being a system and being a component or unit of a system is. A black-box may be opened, but opening it transforms it into a grey-box, in which more black-boxes may be found. Whether *ad infinitum* we simply cannot know. It might be Russian dolls (informational objects) all the way in.²⁷

²⁴ Apparently David Lewis held a similar view, see Langton [2004].

²⁵ Van Fraassen [forthcoming] seems to be sharing a similar view on truth as a relation between models.

²⁶ This seems to me the point of convergence between ISR and the position expressed by French and Saatsi [2004].

²⁷ Many people have helped me to shape and sharpen the ideas presented in this paper. I am grateful to John Weckert and the Centre for Applied Philosophy and Public Ethics (CAPPE) for his invitation to deliver the

References

- Bitbol, M. 2002, "Jean-Louis Destouches' Transcendental Structuralism", *One-Day Workshop "Structuralism in Twenty-Century Physics"*, University of Leeds, June 14.
- Bloomfield, P. 2005, "Let's Be Realistic About Serious Metaphysics", *Synthese*, 144(2), 69-90.
- Brink, C., and Rewitzky, I. 2002, "Three Dual Ontologies", *Journal of Philosophical Logic*, 31(6), 543-568.
- Cao, T. Y. 2003, "Structural Realism and the Interpretation of Quantum Field Theory", *Synthese*, 136(1), 3-24.
- Carter, J. forthcoming, "Identity between and Individuating Objects in Structures: A Problem for Structuralism?" *Synthese*.
- Chakravartty, A. 1998, "Semirealism", *Studies in History and Philosophy of Science*, 29(3), 391-408.
- Chakravartty, A. 2001, "The Semantic or Model-Theoretic View of Theories and Scientific Realism", *Synthese*, 127(3), 325–345.
- Chakravartty, A. 2003, "The Structuralist Conception of Objects", *Philosophy of Science*, 70(5), 867 - 878.
- Cruse, P., and Papineau, D. 2002, "Scientific Realism without Reference" in *The Problem of Realism*, edition, edited by M. Marsonet (Aldershot: Ashgate Publishing Company),
- da Costa, N. C. A., and French, S. 2003, *Science and Partial Truth: A Unitary Approach to Models and Reasoning in Science* (Oxford: Oxford University Press).
- de Roeper, W.-P., and Engelhardt, K. 1998, *Data Refinement : Model-Oriented Proof Methods and Their Comparison* (Cambridge: Cambridge University Press).

keynote lecture at the *1st Australian Computing and Philosophy Conference (CAP@AU)*; the Australian National University in Canberra, 31 October – 2 November, 2003), on which this article is based. I am also grateful to Olga Pombo, for the opportunity to deliver a series of lectures on the philosophy of information at the Universidade Nova de Lisboa in March 2003, which provided the first basis for this paper; and to Jim Moor, for the opportunity of discussing a revised version of it at a seminar at Dartmouth College in April 2004. In each case, I benefited enormously from the audience's feedback. Gianluca Paronitti and Matteo Turilli made several useful comments on a previous version. I am not sure I have answered all their insightful questions. I was influenced by Erich Reck's work on mathematical structuralism and he kindly helped me with his papers. Paul Oldfield carefully copyedited the final manuscript. As usual, my intellectual debt to Jeff Sanders is such that he is virtually a co-author, although I remain the only one of us responsible for any remaining mistake. The anonymous referees of the journal provided a most useful feedback. Too often this sort of acknowledgement is taken to be a mere formality. It is not. I am truly grateful for the referees' time and careful reading of the paper and for their exceptionally sharp comments and high level of scholarship. The paper is much better thanks to them.

- Demopoulos, W., and Friedman, M. 1985, "Critical Notice: Bertrand Russell's the Analysis of Matter", *Philosophy of Science*, 52, 621-639.
- Dennett, D. 1991, "Real Patterns", *Journal of Philosophy*, 87, 27–51. Reprinted in D. Dennett, *Brainchildren* (London: Penguin 1998) pp. 95–120.
- Devitt, M. forthcoming, "Scientific Realism" in *The Oxford Handbook of Contemporary Analytic Philosophy*, edition, edited by Frank Jackson and Michael Smith (Oxford: Oxford University Press),
- Dorato, M., and Pauri, M. forthcoming, "Holism and Structuralism in Classical and Quantum General Relativity" in *Structuralism and Quantum Gravity*, edition, edited by S. French (Oxford: Oxford University Press),
- Eddington, A. 1928, *The Nature of the Physical World* (Cambridge: Cambridge University Press).
- Esfeld, M. 2004, "Quantum Entanglement and a Metaphysics of Relations", *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, 35, 625-641.
- Esfeld, M., and Lam, V. forthcoming, "Moderate Structural Realism About Space-Time", *Synthese*.
- Floridi, L. 2003a, "On the Intrinsic Value of Information Objects and the Infosphere", *Ethics and Information Technology*, 4(4), 287-304.
- Floridi, L. 2003b, "The Renaissance of Epistemology: 1914-1945" in *Cambridge History of Philosophy 1870-1945*, edition, edited by Thomas Baldwin (Cambridge: Cambridge University Press),
- Floridi, L. 2004, "On the Logical Unsolvability of the Gettier Problem", *Synthese*, 142(1), 61-79.
- Floridi, L. 2005a, "Consciousness, Agents and the Knowledge Game", *Minds and Machines*, 15(3-4), 415-444.
- Floridi, L. 2005b, "Is Information Meaningful Data?" *Philosophy and Phenomenological Research*, 70(2), 351-370.
- Floridi, L., and Sanders, J. W. 2004a, "Internet Ethics: The Constructionist Values of Homo Poieticus" in *The Impact of the Internet on Our Moral Lives*, edition, edited by Robert Cavalier (New York: SUNY),

- Floridi, L., and Sanders, J. W. 2004b, "Levellism and the Method of Abstraction", *IEG Research Report 22.11.04*, Oxford University, <http://web.comlab.ox.ac.uk/oucl/research/areas/ieg>.
- Floridi, L., and Sanders, J. W. 2004c, "The Method of Abstraction" in *Yearbook of the Artificial. Nature, Culture and Technology. Models in Contemporary Sciences*, edition, edited by M. Negrotti (Bern: Peter Lang), Preprint from <http://www.wolfson.ox.ac.uk/~floridi/papers.htm>.
- Franklin, J. 1999, "Structure and Domain-Independence in the Formal Sciences", *Studies in History and Philosophy of Science*, 30(4), 721-723.
- French, S. 2001, "Symmetry, Structure and the Constitution of Objects", *Symmetries in Physics, New Reflections: Oxford Workshop, January 2001*, Oxford
- French, S. 2003, "Scribbling on the Blank Sheet: Eddington's Structuralist Conception of Objects", *Studies in History and Philosophy of Modern Physics*, 34, 227–259.
- French, S., and Kamminga, H. (ed.) 1993, *Correspondence, Invariance and Heuristics. Essays in Honour of Heinz Post* (Dordrecht: Kluwer).
- French, S., and Ladyman, J. 2003a, "The Dissolution of Objects: Between Platonism and Phenomenalism", *Synthese*, 136(1), 73-77.
- French, S., and Ladyman, J. 2003b, "Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure", *Synthese*, 136(1), 31–56.
- French, S., and Saatsi, J. 2004, "Realism About Structure: The Semantic View and Non-Linguistic Representations", *Proceedings Philosophy of Science Assoc. 19th Biennial Meeting - PSA2004: PSA 2004 Symposia*.
- 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. 1999, *The Social Construction of What?* (Cambridge, Mass.: Harvard University Press).
- Hayes, I., and Flinn, B. 1993, *Specification Case Studies* 2nd ed (New York ; London: Prentice Hall).
- Heath, A. E. 1928, "Contribution to the Symposium 'Materialism in the Light of Scientific Thought'", *Proceedings of the Aristotelian Society, Supplement*, 8, 130–142.
- Hoare, C. A. R., and He, J. 1998, *Unifying Theories of Programming* (London: Prentice Hall).

- Ketland, J. 2004, "Empirical Adequacy and Ramsification", *British Journal for Philosophy of Science*, 55(2), 287-300.
- Krause, D. 2005, "Structures and Structural Realism", *Logic Journal of the IGPL*, 13(1), 113-126.
- Ladyman, J. 1998, "What Is Structural Realism?" *Studies in History and Philosophy of Science*, 29A(3), 409-424.
- Langton, R. 2004, "Elusive Knowledge of Things in Themselves", *Australasian Journal of Philosophy*, 82(1), 129-136.
- Laudan, L. 1981, "A Confutation of Convergent Realism", *Philosophy of Science*, 48, 19-48.
- Lewis, D. forthcoming, "Ramseyan Humility" in *The Canberra Plan*, edition, edited by David Braddon-Mitchell and Robert Nola (Preprint in Lewis (2001), *Ramseyan Humility*, University of Melbourne Philosophy Department Preprint 1/01.
- Maxwell, G. 1968, "Scientific Methodology and the Causal Theory of Perception" in *Problems in the Philosophy of Science*, edition, edited by Imre Lakatos and Alan Musgrave (Amsterdam: North-Holland Publishing Company), 148-160.
- Maxwell, G. 1970a, "Structural Realism and the Meaning of Theoretical Terms" in *Analyses of Theories, and Methods of Physics and Psychology*, edition, edited by Stephen Winokur and Michael Radner (Minneapolis: University of Minnesota Press), 181-192.
- Maxwell, G. 1970b, "Theories, Perception and Structural Realism" in *Nature and Function of Scientific Theories*, edition, edited by Robert Colodny (Pittsburgh: University of Pittsburgh Press), 3-34.
- Morganti, M. 2004, "On the Preferability of Epistemic Structural Realism", *Synthese*, 142(1), 81-107.
- Newman, M. 2004, "Ramsey-Sentence Realism as an Answer to the Pessimistic Meta-Induction", *Proceedings Philosophy of Science Assoc. 19th Biennial Meeting - PSA2004: PSA 2004 Contributed Papers*.
- Newman, M. H. A. 1928, "Mr. Russell's "Causal Theory of Perception"", *Mind*, 37, 137-148.
- Parsons, C. 2004, "Structuralism and Metaphysics", *The Philosophical Quarterly*, 54(214), 56-77.
- Poincaré, H. 1902, *Science and Hypothesis* (repr. New York: Dover, 1952).
- Psillos, S. 1999, *Scientific Realism : How Science Tracks Truth* (London: Routledge).
- Psillos, S. 2000a, "Carnap, the Ramsey-Sentence and Realistic Empiricism", *Erkenntnis*, 52, 253-279.

- Psillos, S. 2000b, "The Present State of the Scientific Realism Debate", *British Journal for Philosophy of Science*, 51 (Special Supplement), 705-728.
- Psillos, S. 2001, "Is Structural Realism Possible?" *Philosophy of Science (Supplement)*, 68(3), 13-24.
- Psillos, S. 2004, "The Structure, the Whole Structure and Nothing but the Structure?" *Proceedings Philosophy of Science Assoc. 19th Biennial Meeting - PSA2004: PSA 2004 Symposia*.
- Putnam, H. 1975, "What Is Mathematical Truth?" *Mathematics, Matter and Method, Philosophical Papers, Volume 1*, 60-78. (Cambridge: Cambridge University Press)
- Quine, W. V. 1939, "A Logistical Approach to the Ontological Problem", *Fifth International Congress for the Unity of Science*, Cambridge, MA, repr. in *The Ways of Paradox and Other Essays* (New York: Random House, 1966), 64-9.
- Quine, W. V. 1992, "Structure and Nature", *Journal of Philosophy*, 89(1), 5-9.
- Reck, E. H. 2003, "Dedekind's Structuralism: An Interpretation and Partial Defense", *Synthese*, 137(3), 369-419.
- Reck, E. H., and Price, M. P. 2000, "Structures and Structuralism in Contemporary Philosophy of Mathematics", *Synthese*, 125(3), 341-383.
- Rosen, G., and Smith, N. J. J. 2004, "Wordly Indeterminacy: A Rough Guide", *Australasian Journal of Philosophy*, 82(1), 185-198.
- Rumbaugh, J. 1991, *Object-Oriented Modeling and Design* (Englewood Cliffs ; London: Prentice-Hall International).
- Ryckman, T. A. 2003, "Surplus Structure from the Standpoint of Transcendental Idealism: The "World Geometries" Of Weyl and Eddington ", *Perspectives on Science*, 11(1), 76-106.
- Saunders, S. 2003, "Structural Realism, Again", *Synthese*, 136(1), 127-133.
- Sellars, W. 1956, "Empiricism and the Philosophy of Mind" in *Minnesota Studies in the Philosophy of Science*, edition, edited by H. Feigl and M. Scriven (Minneapolis, MN: University of Minnesota Press),
- Shapiro, S. 2004, "Foundations of Mathematics: Metaphysics, Epistemology, Structure", *The Philosophical Quarterly*, 54(214), 16-37.
- Spivey, J. M. 1992, *The Z Notation : A Reference Manual* 2nd ed (New York ; London: Prentice-Hall).

- Suárez, M. 2003, "Scientific Representation: Against Similarity and Isomorphism", *International Studies in the Philosophy of Science*, 17(3), 225-244.
- Van Fraassen, B. C. forthcoming, "Representation: The Problem for Structuralism", *Philosophy of Science*.
- Van Fraassen, B. C. 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, Supplement*, 70(5), 879-890.
- Votsis, I. forthcoming, "Dispelling Certain Misconceptions About Structural Realism".
- Wallace, D. 2003, "Everett and Structure", *Studies in History and Philosophy of Modern Physics*, 34, 87–105.
- Wiener, N. 1961, *Cybernetics, or Control and Communication in the Animal and the Machine* (New York: John Wiley).
- Worrall, J. 1989, "Structural Realism: The Best of Both Worlds?" *Dialectica*, 43, 99-124.
- Worrall, J. 1994, "How to Remain (Reasonably) Optimistic: Scientific Realism and The "Luminiferous Ether"" in *Psa 1994*, edition, edited by David L. Hull, Micky Forbes, and Richard M. Burian (East Lansing, MI: Philosophy of Science Association), 334-342.
- Worrall, J., and Zahar, E. 2001, "Appendix Iv: Ramseyfication and Structural Realism" in *Poincare's Philosophy: From Conventionalism to Phenomenology*, Chicago and La Salle (IL): Open Court),