

IS STRUCTURAL REALISM POSSIBLE?*

Stathis Psillos

Dept of Philosophy & History of Science

University of Athens

37 John Kennedy Str.

16121, Athens, Greece

psillos@netplan.gr

1. Introduction. Structural Realism (SR) is a cluster of attempts to defend a substantive philosophical position concerning what there is in the world and what can be known of it. It is realist because it typically asserts the existence of a mind-independent world, and it is structural because what is knowable of the world is said to be its structure only. As a slogan, the thesis is that knowledge can reach only up to the structural features of the world. This paper aims to unravel and criticise the metaphysical presuppositions of SR. It will question its very possibility as a substantive—and viable—realist thesis.

2. The Upward path. Let the 'upward path' to SR be any attempt to begin from empiricist premises and reach a sustainable realist position. (The term "structural realism" was coined by Grover Maxwell). Arguing against the then dominant claim that only the appearances ("the world of percepts) can be known and that, even if they exist, their "objective counterparts" (the "things-in-themselves") are unknowable, Russell (1919, 61) suggested that "the objective counterparts would form a world having the same structure as the phenomenal world, [a fact which would allow us] to infer from the phenomena the truth of all propositions that can be stated in abstract terms and are known to be true of the phenomena". In his (1927, 226–7), he stressed that only the structure, i.e., the totality of the formal, logico-mathematical properties, of the external world can be known, while all of its first-order properties are inherently unknown. This logico-mathematical structure, he thought, could be legitimately inferred from the structure of the perceived phenomena (the world of percepts). Insofar as this inference is legitimate from an empiricist perspective, then the conclusion, viz., that the unperceived (or unobservable) world has a certain knowable structure, will be legitimate too. But how is this inference possible? Russell rested on the (metaphysical) principle that differences in percepts are brought about

* Thanks to Theodore Arabatzis, Otavio Bueno, Peter Clark, Craig Callender, Chris Daly, Bill Demopoulos, Igor Douven, Steven French, Carl Hoefer, Jeff Ketland, James Ladyman, Thomas Uebel and John Worrall for stimulating comments and discussions. A version of this paper was presented at the annual BSPS conference in Sheffield in July 2000.

by relevant differences in their causes (stimuli). This is a supervenience principle: if two stimuli are identical, then the resulting percepts will be identical. I call this the 'Helmholtz–Weyl' principle, for it was Helmholtz and then Weyl who first enunciated it: "we are justified, when different perceptions offer themselves to us, to infer that the underlying real conditions are different" (cf. quoted by Weyl, 1963, 26). Weyl endorsed it because he thought that it grounds the possibility of knowing something about the "world of things in themselves". Yet, what is known of this world via the Helmholtz–Weyl principle, Russell (and Weyl 1926, 25–6) thought, is its structure. How is that? If we plug in the Helmholtz–Weyl principle the principle of "spatio–temporal continuity" (that the cause is spatio–temporally continuous with the effect), Russell said (1927, 226–7), we can have "a great deal of knowledge as to the structure of stimuli". This knowledge is that "there is a roughly one–one relation between stimulus and percepts", which "enables us to infer certain mathematical properties of the stimulus when we know the percept, and conversely enables us to infer the percept when we know these mathematical properties of the stimulus" (ibid.). The "intrinsic character" of the stimuli (i.e., the nature of the causes) will remain unknown. The structural isomorphism between the world of percepts and the world of stimuli isn't enough to reveal it. But, for Russell, this is just as well. For as he also points out: "(...) nothing in physical science ever depends upon the actual qualities" (1927, 227). Still, he insists, we have known something about the structure of the world (cf. 1927, 254).

There may be good reasons to doubt the Helmholtz–Weyl principle (e.g., that the stimuli overdetermine the percepts). But even if we granted the Helmholtz–Weyl principle, the Russellian argument that we can have inferential knowledge of structural isomorphism between the world of percepts and the world of stimuli relies on a minor miracle. This miracle is necessary because the Helmholtz–Weyl principle is not strong enough on its own to generate the required isomorphism. The determination it dictates is one–way only: same stimuli, same percepts. But for the establishment of isomorphism, the converse of the Helmholtz–Weyl principle—viz., same percepts, same stimuli— is also required. Precisely because Russell doesn't have the converse principle, he talks of "roughly one–one relation". Yet, he has failed to motivate the claim that the relation should be 1–1. (Why can't the same stimuli produce different perceptions at different times, for instance?). Besides, does it make good sense to talk of "roughly one–one relation"? Either it is or it isn't one–one. If it is, we have structure–transference. But if it isn't, we don't.

Given the importance of the converse of the Helmholtz–Weyl principle for the Russellian argument, does it have any independent motivation? From a realist viewpoint, it should be at least in principle possible that the (unobservable) world has 'extra structure', that is structure not necessarily manifested in the structure of the phenomena. If there is such 'extra structure', the required structural relation between the phenomena and the (unobservable) world should not be isomorphism but embedability, viz., the phenomena are isomorphic to a substructure of the world. But then all the original attraction of the Russellian attempt to motivate a compromise between empiricism and realism is lost. This attraction was that one could stay within the empiricist confines, and yet also claim that the structure of the (unobservable) world can be known (inferred). When isomorphism gives way to the weaker (but more realist in spirit) requirement of embedability, there is no (deductive) constraint any more on what this extra structure of the (unobservable) world might be. The structure of the phenomena no longer dictates it; from the empiricist–cum–structuralist perspective we are more or less free to invest the world with whatever extra structure we please provided that the appearances are embedded in it. So, the attempted structural compromise between an empiricist approach and realism, based on the thought that the structure of the world can be inferred from the phenomena, collapses.

Although Russell remained silent on all this, Weyl endorsed the converse of the Helmholtz–Weyl principle, viz. that "The objective image of the world may not admit of any diversities which cannot manifest themselves in some diversity of perceptions" (1963, 117). So, different stimuli, different percepts. This principle Weyl takes to be "the central thought of idealism" (ibid.), and reckons that science should concede it: "[S]cience concedes to idealism that its objective reality is not given but to be constructed ...". But how can this be asserted a priori? And if it is, as a means to secure a priori the knowability of the structure of the world, it is inconsistent with the realist side of SR. For a realist it should not be a priori false that there is a divergence between the structure of the world and the structure of the phenomena. But Weyl's 'idealist' principle simply blocks the possibility that the world has any 'extra structure' which cannot manifest itself in experience.

To be sure, Russell did stress that "indistinguishable percepts need not have exactly similar stimuli" (1927, 255). That is, he reckoned that the relation between the percepts and the stimuli is one–many and not one–one. Although this is definitely a more reasonable approach, it makes the required structural compromise between

empiricism and realism even more tenuous. Given an one–many relation, the structure of the percepts doesn't determine the domain of the stimuli. But it may not determine its structure either. For if two or more different stimuli can give rise to the same percepts, knowledge of the structure of the percepts doesn't suffice for knowledge of the structure of the stimuli, assuming, as it is natural to do, that these stimuli don't just differ in their "intrinsic nature". They may well differ in other structural characteristics which nonetheless don't surface in any way in the world of phenomena. So, the structural differences of stimuli cannot be inferred from the non–existent corresponding structural differences in the percepts. Hence, the more reasonable one–many approach leaves no room for inferring the structure of the unobservable world from the structure of the phenomena. For all we know, the unobservable world may differ from the world of phenomena not just in its "intrinsic nature", but in its structure too.

In sum, the Russellian 'upward path' to SR faces an important dilemma. Without the converse of the Helmholtz–Weyl principle, it cannot establish the required isomorphism between the structure of the phenomena and the structure of the (unobservable) world. Hence it cannot establish the possibility of inferential knowledge of the latter. With the converse of the Helmholtz–Weyl principle, it guarantees knowledge of the structure of the world, but at the price of conceding—a priori—too much to idealism.

Russell's thesis was revamped by Maxwell—with a twist: Maxwell took the Ramsey–sentence approach to exemplify the proper structuralist commitments. He advocated structural realism as a form of representative realism, where all first–order properties are unknowable and only higher–order properties of things are knowable. As he said: "our knowledge of the theoretical is limited to its purely structural characteristics and (...) we are ignorant concerning its intrinsic nature" (1970a, 188). The structural characteristics (or properties) were taken to be those that are not, or could not be, "direct referents of predicates". To be sure, Maxwell didn't take the so–called structural characteristics to be purely formal. He thought that formal properties, such as transitivity, are purely structural, but added that "not all structural properties are also purely formal [...]; in fact those referred to by scientific theories rarely are" (1970a, 188). Yet, he leaves us in the dark as to what these non–formal structural properties which are referred to by theories are. Saying that they are "always of a higher logical type, they are properties of properties, properties of properties of properties etc." isn't enough to show that they are not purely formal. On the contrary,

in Maxwell's Ramsey-sentence approach to structuralism, where the so-called theoretical properties have been replaced by (second-order) bound variables, there is simply nothing left than formal properties and observable ones. In the Ramsey-sentence approach, all non-observable properties are just formal properties: ultimately, they are just logical relations between the higher-order bound variables. So, when Maxwell says that the Ramsey-sentence "refers by means of description to unobservable intrinsic properties" (whose intrinsic nature is otherwise unknowable), he misses the point. For there is nothing in the Ramsey-sentence itself which tells us that what is referred to by the bound variables are properties (of any sort) of unobservable entities. This last construal, though consistent with the meaning of a Ramsey-sentence, is in no way dictated by it, as Carnap himself noted (cf. Psillos 1999, 53).

While Russell takes the structure of the (unperceivable) world to be isomorphic to the structure of the phenomena, Maxwell takes the phenomena to be the observable part of the Ramsey-sentence of our theory of the world. The structure of the phenomena is not, on Maxwell's view, isomorphic to the structure of the world. Instead, (descriptions of) the phenomena are embedded in logico-mathematical structures. This difference might be thought enough to guarantee that Maxwell's path to SR is viable. But it isn't. For two reasons. First, unless we smuggle in a suspect requirement of direct acquaintance (or intuitability), it isn't clear why the first-order properties of unobservable entities are unknowable. They are, after all, part and parcel of their causal role. So, if all these entities are individuated and become known via their causal role, there is no reason to think that their first-order properties, though contributing to causal role, are unknowable. Second, the most fatal objection to the Russellian programme, viz. Newman's (1928) claim that knowledge of structure is empty since it just amounts to knowledge of the cardinality of the domain of discourse, applies to Maxwell's thesis no less. This objection has been recently discussed a lot, so I'll simply summarise it: even when the domain of the structure of the stimuli is fixed, a relational structure on this domain can always be defined (cardinality permitting) in such a way as to guarantee isomorphism between the structure of the percepts and the structure of the stimuli. Hence, the only information encoded in the claim of structural isomorphism is that the domain of the stimuli has a certain cardinality (cf. Demopoulos & Friedman 1986; Psillos 1999, 61-9). Maxwell's Ramsey-style approach to SR fares no better than Russell's vis-à-vis the Newman objection. It can easily be shown that without further (non-structural) constraints on the range of the second-order variables of the Ramsey-sentence, if the

Ramsey-sentence is empirically adequate, then it is guaranteed—by logic alone—to be true (cf. Psillos 1999). Hence, the claim that the Ramsey-sentence captures the structure of the (unobservable) world is empty, unless—to repeat—non-structural constraints are placed on the range of its variables.

To sum up, the 'upward path' to SR is not viable. Structural isomorphism cannot be inferred without contentious metaphysical assumptions—which compromise the realist side of the wedding, anyway. And, even if it could be, it would be an empty claim (the Newman objection).¹

3. The Downward Path. Let the 'downward path' to SR be any attempt to start from realist premises and construct a weaker realist position. It has found its most prominent exponent in Worrall's (1989). Take realism to incorporate two conditions: the Independence Condition (IC), viz. that there is a mind-independent world (that is a world which can be essentially different from whatever is constituted by our conceptual capacities and our capacities to perceive). And the Knowability Condition (KC), viz., that this mind-independent world is knowable. Call a thesis which subscribes to both conditions Metaphysical Realism (MR).² Being realism, SR should subscribe to at least IC. SR then, has to be different from MR in its approach to KC. In particular, SR differs from MR in that it puts a principled restriction on what aspects of the mind-independent world can be known. The claim is that only the structure of the world can be known. Note that all that MR says is that the world is knowable. It doesn't imply anything specific about what aspects of the world are knowable. It may be an overly optimistic thesis that everything in the world is knowable. So, let's see what principled restrictions SR might want to impose on KC. Taking account of the SR claim that only the structure of the world can be known, two options are available. Either there could be something else—call it X—of the world, which however cannot be known, or there is nothing else in the world to be known. On the first disjunct, the restriction imposed by SR is epistemic. Call this version of SR, Restrictive Structural Realism (RSR). On the second disjunct, the restriction is ontological: there is nothing other than structure that can be known, because there is nothing else. Call this version of SR Eliminative Structural Realism (ESR). So it appears that the stronger version of realism which SR tries to surpass is

1 Hochberg (1994) has claimed that a) Russell came to defend a version of "hypothetico-scientific realism" and b) an intensional understanding of relations is enough to dispel Newman's challenge. I think that a') even if (a) is correct, hypothetico-scientific realism is different from structural realism and b') an appeal to intentions may be enough to answer the Newman challenge only at the price of abandoning pure (extensional) structuralism (cf. Psillos, 1999, 67–9).

2 For a discussion of these two theses cf. Psillos (forthcoming).

the claim that there is more to the world than its structure, and this more—the X—can be known. RSR grants that there is this extra X, but denies its knowability, whereas ESR eliminates this extra X altogether. But what is this extra X?

Let's think in terms of a metaphor. Take MR to assert that the world forms a gigantic (interpreted) 'natural' graph, where there are entities and their properties and relations to which they stand among themselves and higher-order properties and relations. It is an essential part of the metaphor that all these entities are definite: there are objective similarities and differences among things which constitute a fact-of-the-matter as to what these entities are and how they relate to each other. KC implies that all these are in principle knowable (as a contingent matter-of-fact). What would then the parts of the graph which SR either intends to restrict epistemically or eliminate altogether? Here are the options about the supposed unknowable X: the objects (individuals); the (first-order) properties; the relations; the higher-order properties of properties and relations. Each level of abstraction creates a version of SR. The claim that we can know the structure of the graph is, therefore, ambiguous. It may mean that:

(A): We can know everything but the individuals that instantiate a definite structure.

Or, (B): We can know everything except the individuals and their first-order properties.

Or (C): We can know everything except individuals, their first-order properties and their relations.

Notice that as we went down the line, we relegated more and more things to the 'unknowable extra X'. Where exactly do we draw the line? The issue is that, with the possible (but uninteresting) exception of (A), by leaving out entities we leave out structure as well.

Take RSR to mean (A). Then, RSR implies that a proponent of MR should endorse the following thesis: if there were two interpreted structures which were exactly alike in all respects except their domains of discourse, there would still be a fact of the matter as to which of those is the correct structure of the world (assuming of course that they 'surface' in the same way and they don't conflict with observations.) Realists may want to take sides on this issue. But I don't think they have to. For the only possibly substantive issue that remains is to name the individuals of the domain.

But how can this be a substantive issue? Since by hypothesis, the individuals in the two domains instantiate the very same interpreted (natural) structure, for each individual in domain M there is another individual in domain M' such that the two individuals share exactly the same causal role. Whether or not these two individuals will be taken to be the same will depend on how one thinks about individuation. If one thought that properties require a substratum to which they inhere, then one is also open to the thought that there may well be two distinct individuals with all their properties in common. If so, the whole issue between MR and RSR hinges upon the metaphysics of properties. I don't intend to resolve this issue here. Suffice it to note that under (A), RSR is, in a principled way, weaker than MR only if MR is taken to accept (the questionable thesis) that two individuals can share all their properties and yet be different. It may be thought here that this might be a case of multiple realisation. But in this context, this case is uninteresting. For we are not talking about two systems which can be described differently on physical terms but instantiate the same higher-level structure. Rather, we are talking about two (natural) structures whose only difference is that their individuals can be named differently.

Suppose RSR means (B). Then RSR implies that MR should endorse the following thesis: if there were two semi-interpreted structures which are exactly alike in all respects except their domains of discourse and the first-order properties attributed to individuals, then there would still be a fact of the matter as to which of those is the correct structure of the world. This is a view that some realists would accept. But has RSR come up with a principled restriction? (B) amounts to Carnapian "relation descriptions" (1928, 20). Relation descriptions offer less information than "property descriptions" (associated with (A)). They describe an object as that which stands in certain relations to other objects, e.g., 'a is the father of b', without specifying its properties. What this object is is not thereby fully specified. Only what it is in relation to other objects is being given. Although "relation descriptions" don't entail unique "property descriptions", they do offer some information about an object, because, generally, they entail some of its properties. For instance, from the relation description 'a is the father of b' we can conclude that a is male, that a is a parent, etc. More interestingly, from relational descriptions about, e.g., electrons, we can legitimately infer the existence of some first-order properties, namely, negative charge or mass. There is simply no natural epistemic cut between the relational and the first-order properties. Hence, given a rich enough relation-description, one may infer enough properties of the object to identify this object to a great extent. Hence, (B) doesn't imply inherent unknowability. That some properties may remain

unspecified doesn't justify the claim that only relations are knowable. Although this version of RSR would make it weaker than a realist position which says that all properties and relations can be known, the difference is only one of degree. For this version of RSR doesn't preclude any specific property of being knowable. Rather, it says that if we go for relation descriptions, some properties may be unknowable. MR should be happy with this.

Having shown that structural restrictions of types (A) and (B) above don't pose any severe restriction to MR, we should consider (C). So, take RSR to mean that only the higher-order properties of properties and relations are knowable. Here we have reached pure structure, as it were. We can therefore talk of Carnapian "purely structural definite descriptions" (1928, 25): not only are we ignorant of the objects and their properties, but also we don't know what their relations are. All we can know is the totality of the formal properties and relations of the structure, where these are those which can be formulated without reference to their meanings. Should the domain of a relation be finite, the structure of this relation is given by its logico-mathematical graph. This is a complete structure description, whose full verbal equivalent is a list of ordered tuples. This, I think, is the only characterisation of RSR which can impose a principled limitation on what is knowable. To return to the original graph metaphor, the claim is that all that can be known is the structure of the graph. Note that RSR(C) doesn't deny that there are definite relations which structure the world. It just claims that we cannot know what they are.

Is RSR, in the form (C), a defensible position? Despite the 'downward path' we followed, we have just reached the highest point of the empiricist 'upward path'. So, all the problems met there apply here too. In particular, there are two problems. The problem of motivation and the problem of content. Why should we consider accepting the restriction RSR places on what can be known of the world? And, what is the formal structure of the world a structure of? If we could just infer the formal structure of the world from the phenomena, then the motivational question would be answered. But RSR fares no better on this score than Russellian empiricism. For the reasons expressed in section 2, there is no path (at least no stable path) from the structure of the phenomena to the structure of the unobservable world. (Note that the question of motivation has already been implicitly raised for versions (A) and (B) above.) What about the problem of content? Two issues are relevant here. First, in empirical science we should at least seek more than formal structure. Knowing that the world has a certain formal structure (as opposed to definite properties and

relations—a natural structure) allows no explanation and no prediction of the phenomena. Second, the claim that the world has a certain formal structure is, from a realist point of view, unexciting. That the world has a formal structure follows trivially from set–theory, if we take the world to be a set–theoretic entity. In fact (this is Newman’s objection), it follows that the world has any formal structure we please (consistent with its cardinality). That it has the formal structure that corresponds to a definite natural structure (that is, that the formal structure has a certain natural structure as its content) is a much more exciting claim, which however cannot be established by purely structural means. The ‘downward path’ to SR should at least retain the realist outlook. But this cannot be retained unless there is a prior commitment to the not–purely–structural thesis that the world has a definite natural structure. That this thesis isn’t purely structural follows from the fact that its defence (and possibility) requires a commitment to objective (and knowable, at least in principle) objective similarities and differences between natural kinds in the world.

To sum up, the ‘downward path’ to RSR either fails to create a sustainable restriction on MR or, insofar as it focuses on the knowability of purely formal structure, it fails to be realist enough. It can be realist by talking of ‘natural structures’, but then again it gives up on pure structuralism.

What about ESR? Remember that ESR is eliminative: only structure can be known, simply because there is nothing else to know. A proposal in this direction has been recently defended by Ladyman (1998) and French (1999). Ladyman urges us to take SR as a metaphysical thesis. He notes: "This means taking structure to be primitive and ontologically subsistent" (1998, 420). And he adds that the empirical success of theoretical structures shouldn’t be taken to "supervene on the successful reference of theoretical terms to individual entities, or the truth of sentences involving them" (422). French (1989, 203) stresses that on their view "there are no unknowable objects lurking in the shadows (...)". The details of this view are not specified yet, so it is difficult to evaluate it properly. Nonetheless, I fail to see how the eliminativism it suggests is possible. If structures are independent of an ontology of individuals and properties, then we cannot even speak of any structural relation (be it isomorphism, or embedding or what have you) between structures. I doubt that we have any special insight into structural relations—we establish them by pairing off individuals and mapping properties and relations onto one another. Even in group theory—Ladyman and French’s paradigm case for their thesis—the group–structure is detached from any particular domain of individuals, but not from the very notion of a domain of

individuals. To hypostatise structures is one thing (may be legitimate). But to say that they don't supervene on their elements is quite another. It implies the wrong ontological thesis that they require no individuals in order to exist and the wrong epistemic thesis that they can be known independently of (some, but not any in particular, set of) individuals which instantiate them. Note that if the structures "carry the ontological weight" (French, 204), we can only take the identity of structures as something ontologically primitive (since the notion of isomorphism requires different domains of individuals which are paired-off). But I am not sure whether we can even make sense of this primitive structural identity. And if we introduce individuals as "heuristic" devices (ibid.) whose sole role is "the introduction of structures" (only to be 'kicked away' later on), we need to justify why they are just "heuristic devices" if the only road to structure is via them. I conclude that it is hard to see how the ontological revision ESR suggests is possible. Besides, if "theories tell us not about the objects and properties of which the world is made, but directly about structure and relations" (Ladyman, 1998, 422), it is a major miracle that they can be used to represent the world we live in. Unless we buy into some problematic metaphysical thesis which somehow 'constructs' the individuals out of relations, the world we live in (and science cares about) is made of individuals, properties and their relations.

4. Conclusion. Let me close with a positive note. One way to read SR is as a modest epistemic thesis which emerges from looking into the history of scientific growth. There is no heavy metaphysical machinery behind it, nor absolute claims about what can or cannot be known. It is just a sober report of the fact that there has been a lot of structural continuity in theory-change: we have learned a lot about theoretical and empirical laws, although our views about what entities in the world are related thus have seen some major discontinuities. In a certain sense, this is the insight behind Worrall's motivation for SR. All this can reasonably be accepted without abandoning realism. What isn't acceptable is any form of strong thesis which draws a principled division between the (knowable) structure of the world and some forever elusive (or, worse, non-existent) X.

References

- Carnap, R. (1928) The Logical Structure of the World, University of California Press.
- Demopoulos, W. & Friedman, M. (1985) 'Critical Notice: Bertrand Russell's The Analysis of Matter', Philosophy of Science, 52: 621-639.
- French, S. (1999) 'Models and Mathematics in Physics', in J. Butterfield & C.

Pagonis

(eds) From Physics to Philosophy, Cambridge U.P.

Hochberg, H. (1994) 'Causal Connections, Universals and Russell's Hypothetico-Scientific Realism', Monist, 77: 71–92.

Ladyman, J. (1998) 'What is Structural Realism?' Studies in History and Philosophy of Science, 29: 409–424.

Maxwell, G. (1970) 'Theories, Perception and Structural Realism', in R. Colodny (ed.)

The Nature and Function of Scientific Theories, University of Pittsburgh Press.

Maxwell, G. (1970a) 'Structural Realism and the Meaning of Theoretical Terms' in Minnesota Studies in the Philosophy of Science, 4, University of Minnesota Press.

Newman, M.H.A. (1928) 'Mr. Russell's "Causal Theory of Perception"', Mind, 37: 137–148.

Psillos, S. (1999) Scientific Realism: How Science Tracks Truth, Routledge.

Psillos, S. (forthcoming) 'The Present State of the Scientific Realism Debate', British Journal for the Philosophy of Science.

Russell, B. (1919) Introduction to Mathematical Philosophy, George Allen & Unwin.

Russell, B. (1927) The Analysis of Matter, London: RKP.

Weyl, H. (1963) Philosophy of Mathematics and Natural Science, Atheneum.

Worrall, J. (1989) 'Structural Realism: The Best of Both Worlds?', Dialectica, 43: 99–124.