

Structuralism as a form of scientific realism

Anjan Chakravartty

Structural realism has recently re-entered mainstream discussions in the philosophy of science. The central notion of structure, however, is contested by both advocates and critics. This paper briefly reviews currently prominent structuralist accounts en route to proposing a metaphysics of structure that is capable of supporting the epistemic aspirations of realists, and that is immune to the charge most commonly levelled against structuralism. This account provides an alternative to the existing epistemic and ontic forms of the position, incorporating elements of both. Structures are here identified with relations between first order, causal properties: properties that confer specific dispositions for relations. This form of structuralism constitutes an explicit proposal for what seem implicit structuralist tendencies in sophisticated but more traditional characterizations of realism. An outline of the proposal's response to the anti-realist's pessimistic induction on the history of scientific theories is considered.

1. Two approaches to structuralism

Structural realism (SR) is the view that insofar as scientific theories offer true descriptions of reality, they do not tell us about the underlying nature of reality. Rather, they tell us about its structure. Contemporary proponents see in this a safe route between opposed forces. On the one hand, if theories are to some extent correct in mapping the structures of the natural world, we have an explanation for the tremendous empirical success they afford in allowing us to predict and manipulate natural phenomena. The realist's 'miracle argument' is thus accommodated. On the other hand, if we take only some of the structural parts of theories to represent the world, we have an explanation for what Henri Poincaré called the seeming 'bankruptcy of science'. We may sacrifice the remainder to the anti-realist's pessimistic induction on the history of past theories.

Correspondence: Institute for the History and Philosophy of Science and Technology, Victoria College, University of Toronto, 91 Charles Street West, Toronto, Ontario M5S 1K7, Canada.
E-mail: anjan.chakravartty@utoronto.ca

Advocates of SR take inspiration from precursors who are by no means unified in their philosophical motivations, such as Russell, Poincaré, Duhem, Cassirer, Schlick and Carnap.¹ As a result, discussions of structuralism are susceptible to conflation of different notions of structure; often there is ambiguity regarding these notions to begin with. By focusing clearly on the concept of structure, my goal is to outline a proposal for SR that learns lessons from previous (some argue) problematic accounts, and that achieves the promised compromise between realist and anti-realist arguments. I should stress that this does not by itself amount to a defence of scientific realism. I will not consider, for example, the comparative merits of realism and well-known forms of anti-realism, or defend realism against criticisms concerning the miracle argument or under-determination. The aim is rather to provide a workable account of structuralism that captures the most compelling intuitions of structural realists and those of realists more generally.

Informally, the idea of a structure has to do with relations between the elements of some system of elements. Structuralism focuses on the relations themselves, rather than on any putative relata. In the contemporary literature, SR comes in two flavours: epistemic (ESR) and ontic (OSR).² Epistemic versions place a restriction on scientific knowledge; proponents hold that we can know structural aspects of reality, but nothing about the natures of those things whose relations define structures in the first place. The natures of the objects are beyond the proper grasp of our quest for knowledge. Ontic versions, more radically, do away with objects altogether; proponents hold that at best we have knowledge of structural aspects of reality, because there is in fact nothing else to know. Here we have the denial of any traditional metaphysics of objects.³

I will contend that the most reasonable form of SR is *both* epistemic *and* ontic, but in different ways than those described by current proponents of ESR and OSR. The proposed structuralism will collapse the distinction between knowledge of structures and knowledge of natures. I will attempt to show how this position does justice to both realist and anti-realist arguments in a way that should be acceptable to sophisticated realists generally, but in a way that emphasizes a structuralist approach to scientific knowledge. In outline, I begin by briefly reviewing the epistemic tradition of SR. This account faces a potentially fatal difficulty, but teaches valuable lessons. Next, I consider the motivation for OSR and argue that it does not license the revision in ontology its exponents demand. I then describe the proposed SR, how it avoids prior difficulties and how it accommodates certain realist and anti-realist arguments. The proposal characterizes realism in terms of an explicit structuralism that is implicitly endorsed, I believe, by more traditional but sophisticated realist accounts of scientific knowledge.

2. The epistemic programme: Russell, Maxwell, Worrall

Poincaré (1952, ch. 10) champions a knowledge of relations at the expense of elusive, unknowable objects. Pierre Duhem (1954, p. 32) recommends the ‘representative’ parts of theories (mathematical laws expressing relations) over the ‘explanatory’ parts

(underlying entities and processes). In neither case, however, are we given a precise description of how best to understand the concept of structure. For clarity on this point, structuralists could later turn to Bertrand Russell, who offers the following definition of structural identity: '[w]e shall say that a class α ordered by the relation R has the same structure as a class β ordered by the relation S , if to every term in α some one term in β corresponds, and vice versa, and if when two terms in α have the relation R , then the corresponding terms in β have the relation S , and vice versa' (Russell, 1948, p. 271, see also 1927). Crucially, the members of α and β need not bear any qualitative similarity to one another. The only requirement is that R and S have a purely formal similarity. Structural identity on Russell's account does not depend on what sorts of relations define structures. So long as there is a relation of any sort between elements of a system, this is sufficient to determine a structure.

William Demopoulos and Michael Friedman (1985, pp. 628–629) argue that by suggesting only mathematical (formal, logical) properties of the world can be known, Russell invites a fatal objection. Structure is here a higher order property, a property of relations; no first order properties or relations of objects can be known. The authors cite M. H. A. Newman's (1928, p. 140) criticism that on Russell's definition, the claim that a system has a particular structure tells us nothing about the system other than its cardinality, for any collection of elements can be arranged so as to exemplify a given structure so long as there are enough of them. Given any set α and any arbitrary structure W , it is a consequence of set theory or second order logic that there exists some relation in α having structure W , so long as W is compatible with the number of elements in α . On Russell's approach, the claim that some aspect of the world has a particular structure is thus trivially satisfied.

Grover Maxwell marries Russellian structuralism to the Ramsey sentence to produce his own version of SR. An unobservable entity whose place is held by a predicate variable in a Ramsey sentence is 'whatever it is' that satisfies the relations specified by the sentence. This 'indirect' reference is achieved by 'purely logical terms (variables, quantifiers, etc.) plus terms whose *direct* referents are items of acquaintance [i.e. observables]' (Maxwell, 1970a, p. 16). We do not know what the natures of theoretical entities are, but we can assert that they exist and stand in certain relations. 'This ... may be taken as an explication of the claim of Russell and others that our knowledge of the theoretical is limited to its purely structural characteristics and that we are ignorant concerning its intrinsic nature' (Maxwell, 1970b, p. 188).⁴

Insofar as one interprets Maxwell's programme as an effective means of implementing Russell's, focusing on higher order properties ('purely structural') as opposed to first order properties of entities ('intrinsic natures'), it too is susceptible to the Newman objection. Some, however, dispute this charge. Redhead (2001a) claims that the objection is ineffective, for it correctly applies only to a position that denies the reality of first order properties and relations. Since Russell and Maxwell do not do this, and can admit that such properties and relations exist while restricting more detailed knowledge claims to higher order content, the Newman objection loses its force. Worrall and Zahar (2001) argues that the objection is effective only because of

an easily corrected error on Russell's part. The error is to speak as though our knowledge of reality is *purely* structural, but this applies only to our knowledge of unobservables. If a theoretical description includes observation terms, the corresponding Ramsey sentence is not trivially satisfied.

It is not my aim here to engage with the controversy surrounding the Newman objection.⁵ I believe that the infatuation with structure as a higher order property is avoidable and best avoided. If structure can be understood in terms of relations between first order properties, as Newman (1928) himself makes clear, his objection does not arise. That is, the applicability of the Newman objection turns on the question of what sort of structure one thinks we can know on a tenable account of SR. Structuralists adduce cases of structural continuity across changes in theory to support their positions; this very same evidence, I will argue, supports a knowledge of structures that renders the Newman controversy immaterial. I detail this notion of structure in the next section. First, however, let us draw one further suggestion from the tradition of ESR.

Worrall (1989) brought SR back into the mainstream of philosophy of science after a significant hiatus, identifying Poincaré as its founding father. As mentioned earlier, Poincaré is quiet on the subject of what structures are, and Worrall inherits this ambiguity. His central case study is the transition in theories of light from the wave optics of Fresnel to the electromagnetic theory of Maxwell. Fresnel developed a set of equations relating intensities of incident, reflected and refracted light when a beam passes from one medium into another of different optical density. He believed that light was a disturbance in the ether, but Maxwell's theory was ultimately accepted, not as a description of ethereal disturbances, but as a description of oscillating electromagnetic field vectors. Nevertheless, Fresnel's equations survive intact in Maxwell's theory. The moral, says Worrall: despite changing views concerning the nature of light, here a description of the structure of light survives from one theory to the next. The structure of light is given by the mathematical equations common to Fresnel's and Maxwell's theories.

Here is where the ambiguity sets in. What does it mean to say that mathematical equations are indicative of structure? It is insufficient for a realist simply to point to the equations of theories and claim that they in some sense describe reality, for constructive empiricists, instrumentalists, logical positivists and idealists agree with this much.⁶ I have already hinted that the metaphysics of structure I will propose involves first order properties and relations. A second hint is provided by Maxwell. While the Ramsey sentence may be viewed as a means to the end of implementing Russell's programme, it is not clear that it must be employed in this way. Consider this (merely) suggestive remark from Maxwell (1970a, p. 17): 'Temporal succession, simultaneity, and causal connection must be counted among these structural properties, for it is by virtue of them that the unobservables interact with one another and with observables and, thus, that Ramsey sentences have observable consequences'. In suggesting causal connection as an important structural property, Maxwell was on to something of particular importance to a defensible, structuralist form of scientific realism.

3. Causal properties and structural realism

Let us take the satisfaction of Russell's definition of structural identity as a necessary, but not sufficient, condition of the notion of structure required. Structure, I suggest, correctly understood in this context is not a *higher* order property, a property of relations between first order properties of objects. Instead, let us identify structure *with* relations between first order properties.

What is the intuitive notion of a structure? To give the structure of something is to enumerate its parts and cite relations between those parts; consider, for example, the structures of things like tables, causal processes, societies. I suggest that we view structure as something tied to specific kinds of relations and their characteristic relations.⁷ It is important to distinguish this idea from the way the term 'structure' is often used in physics and mathematics. Redhead (2001a,b) furnishes some helpful terminology here, distinguishing between 'concrete' and 'abstract' structures. A concrete structure consists in a relation between first order properties of things in the natural world. An abstract structure is one of the sort Russell intended: a higher order property that may be instantiated by different concrete structures.⁸ The current proposal for SR concerns the concrete, not merely the abstract. When we say that blueprints display the structure of a house or that the structure of DNA is shown by Watson and Crick's demonstration model, these are instances of the identity of abstract structure. In order to achieve the desired representations, it is only necessary that a blueprint or model instantiate the same abstract structure as that which is represented (or something close, since representations are rarely perfect). As we shall see, however, SR permits epistemic access to more than just abstract structure.

Two concrete structures that have the same higher order, logical form but that are concerned with different sorts of elements and relations are *not* one and the same concrete structure. Realists have traditionally aspired to a knowledge of the concrete. An identity of concrete structure requires that the elements of the sets in question, α and β , as well as the relations R and S , are of the same kind, i.e. the respective elements instantiate the same relations between the same first order properties. This of course does not require that α and β be identical; more than one set of things can instantiate the same first order properties. Unless otherwise indicated, when I speak of structure henceforth it is this sense of concrete structure I intend. Advocates of SR have generally thought that it is only by retreating to abstract structures that realists can respond to worries such as the pessimistic induction. The remainder of this paper will attempt to show that structuralists can be more ambitious than this.

The notion of structure just outlined, being concrete, is immune to the Newman controversy, so let us now turn to Maxwell's hint that causal connection is a crucial aspect of structure. Here we may accommodate two birds with one stone. The first order properties whose relations comprise concrete structures are what I will call *causal* properties: those that confer dispositions for relations, and thus dispositions for behaviour, on the objects that have them. Let us consider this suggestion in some detail.

Why and how do objects causally interact? It is in virtue of the fact that objects have certain properties that they are causally efficacious. Let me attempt to clarify the nature of this ‘in virtue of’ relation. Properties such as charge, mass, acceleration, volume and temperature all confer on the objects that have them certain causal capacities. These capacities are dispositions to behave in certain ways when in the presence or absence of other objects and their properties. The property of mass confers *inter alia* the capacity of a body to be accelerated under applied forces. The property of volume on the part of a gas confers *inter alia* the capacity to become more highly pressurized under applied heat. It is the ways in which these dispositions are linked to one another, i.e. the ways in which objects with various properties are disposed to act in consort with others, that produce what we take to be causal activity. Ultimately, causation has to do with relations determined by dispositions, conferred by causal properties.⁹

Talk of dispositions raises questions I cannot consider here. There are well-rehearsed ways in which this talk might be elaborated. Many empiricists, worried by what they take to be the mysterious connotations of powers, opt for deflationary analyses in terms of conditionals. Of those who hold that dispositions are genuine, occurrent properties, some speak of causal properties in terms of categorical bases of dispositions and others speak of ungrounded dispositions. When I say that properties confer dispositions to enter into relations, the ambiguity of ‘confer’ here is deliberate. I will assume dispositional realism, but use ‘confer’ to signal neutrality on the precise details for present purposes. The present contention is simply that properties are responsible for the behaviours of things; objects behave as they do because they have causal properties.¹⁰

One important consequence of the proposed view is that the central thesis of ESR, that we can have knowledge of structures without knowledge of natures, cannot be maintained. ESR commits to the structural, but not the intrinsic. On the proposed view this prescription is rejected, for all knowledge of structure *contains* knowledge of the intrinsic natures of things. Structures are here identified with specific relations, which are determined by dispositions conferred by first order properties of objects. On this view, to say that two sets have the same structure is *ipso facto* to say something about the intrinsic natures of their members. Natures, comprised of the first order properties of things, are intimately connected to the relations into which objects enter. Speaking loosely, we might say that causal properties are intrinsic, but also have a ‘relational’ quality. They are ‘relational’ in that the dispositions things have as a result of possessing such properties determine the sorts of relations they can enter into. Knowledge of these relations thus gives us insight into the intrinsic natures of things.¹¹

Stathis Psillos (1995, 1999) argues that there is no principled distinction between the structure and nature of an entity or process and that this is a reason to reject ESR. I agree with Psillos that the form of structuralism he criticizes should be rejected, but *contra* Psillos, there is a real distinction between the natures of entities and structural relations. An entity’s first order properties are part of its nature, and natures are possessed whether or not entities are, at any given moment, manifesting all of the particular relations of which they are capable. Structures are identified with relations

and unless these relations obtain, there are no structures to speak of. Thus, although entities always have natures (first order properties), whether or how they can be described structurally will depend on whether they are, at any given moment, manifesting particular relations, and this is generally a contingent matter. However, given that causal properties are understood in terms of dispositions for relations, structural knowledge contains a knowledge of properties, and thus natures. Structures are, metaphorically speaking, 'encoded' in the natures of entities, because first order properties confer dispositions for specific relations, those we recognize as structures.

A knowledge of structure thus entails some knowledge of the first order properties conferring dispositions to bring about the relations detected. It is thus no miracle, the realist claims, that good scientific theories are empirically successful, for they describe the structures of reality: relations that obtain between objects with causal properties. This is far from a vindication of the entire contents of theories, however. In order to cope with the pessimistic induction, the realist requires a principled means of differentiating parts of theories that are likely to be retained as sciences move on from those that are apt for replacement. In sections 5 and 6 I will argue that the proposed account of structure provides such a means.

First, however, let us consider two possible criticisms (one now, the second in section 4). One might complain that the *point* of SR is to do away with intrinsic natures, to speak of relations between elements of reality, not the elements themselves. By denying that we can separate a knowledge of relations and relata in this way, am I no longer describing a form of structuralism? If the proposed view permits knowledge of not only structures but natures and causal relations, one might think this a *reductio* of the very idea of structuralism. On such a view, suggests Papineau (1996, p. 12), 'restriction of belief to structural claims is in fact no restriction at all'. Defining structures in terms of relations between first order properties weakens SR to the extent that it collapses into standard scientific realism.

This charge has things back to front. 'Standard' realism is not what it once was; it has been refined in response to anti-realist criticism. The extent to which sophisticated realists are able to respond to such arguments, I believe, is proportional to the extent to which they have moved, implicitly, towards a form of structuralism. Richard Boyd (1981, pp. 613–614) characterizes realism by saying that 'typically, and over time, the operation of the scientific method results in the adoption of theories which provide increasingly accurate accounts of the causal structure of the world'. Psillos describes scientific knowledge this way:

When scientists talk about the nature of an entity, what they normally do – apart from positing a causal agent – is to ascribe to this entity a grouping of basic properties and relations. They then describe its law-like behaviour by means of a set of equations. In other words, they endow this causal agent with a certain causal structure, and they talk about the way in which this entity is structured. (Psillos, 1999, p. 155)

Many sophisticated accounts of realism are closet versions of structuralism. The problem with these accounts is that they generally give little detailed consideration to what is meant by 'causal structure', and to the implications this concept has for where realists should draw the line between aspects of theories we have good reason to

endorse and those we do not. The current account is an attempt to address this problem. Before elaborating further, however, we must consider a potentially fatal challenge posed by OSR.

4. The ontic programme: French and Ladyman

Advocates of OSR hold that realists should believe only in structures described by theories, because structure is all there is to reality. The more traditional realists' talk of objects is misguided, engendering fatal metaphysical difficulties. Proponents of this position are happy to speak of objects, but only as a *façon de parler* informed by a more enlightened metaphysical picture. An integral part of this picture is the idea that objects, conceived of as bearers of properties that stand in relations, are metaphysically otiose. I will argue, however, that this charge is too quick. Though the metaphysical status of objects is an open question, it is not clear that the idea that objects comprise a genuine ontological category is a moribund stance for the realist.

The case for OSR proceeds from the interpretation of quantum mechanics (QM). QM appears to underdetermine the nature of quantum particles as regards their identity or individuality. This underdetermination, it is claimed, is fatal to any realism other than OSR.¹² Macroscopically and in classical physics we think of objects as having identities which distinguish them from other things. This notion of individuality is reflected in classical, Maxwell–Boltzmann statistics, which recognize different permutations of objects as distinct arrangements. It is unclear, however, whether quantum statistics respect individuality. Consider the possible arrangements of two particles of the same kind (i.e. having the same state-independent properties such as mass, charge, spin) distributed across two energy states. Neither Bose–Einstein nor Fermi–Dirac statistics count particle permutations as constituting different arrangements. Interchanging the particles has no physical significance according to QM.

This suggests that QM objects are not objects in the classical or everyday macroscopic sense and generates a dilemma concerning their putative objecthood. One might hold that quantum particles are peculiar individuals, ones that seem to violate Leibniz's principle of the identity of indiscernibles (PII). Unlike classical objects, quantum objects appear to violate even weak versions of PII, for QM state functions describing assemblies of particles attribute to them all of the same intrinsic and relational (e.g. spatio-temporal) properties. On this understanding, then, one would commit to individuals that are indistinguishable in principle, and individuality must be understood in terms of individual essences (haecceity, primitive this-ness), as opposed to determinate properties. On the other hand, one might hold that particles are not individuals, but non-individuals of some sort. This interpretation is favoured by some who view 'particles' as excitations in a quantum field.

It is argued that these two different understandings of the nature of objecthood are underdetermined by QM. Granting this, what is the significance of this underdetermination? Steven French and James Ladyman contend that the mere fact of underdetermination scuppers any form of realism involving objects, for '[i]t is an

ersatz form of realism that recommends belief in the existence of entities that have such ambiguous metaphysical status' (Ladyman, 1998, p. 420).¹³ OSR is immune to this charge, because it makes no recommendation on behalf of objects. It advocates a conception of reality according to which objects are relinquished in favour of structures taken as 'primitive and ontologically subsistent'.

The construction of an object-free ontology is an intriguing metaphysical programme. The argument from underdetermination, however, is suggestive at best. Mere underdetermination tells us little. In order to demonstrate that an ontology lacking objects is required, one must show that no account of objects is tenable. However, some prefer one of the above metaphysical packages to the other, and even assuming that both are problematic, it is not clear that these options are exhaustive.¹⁴ It is further unclear why, if there are problems of interpretation regarding the metaphysical status of quantum particles, this should infect the status of objects at non-quantum levels of description, thus rendering realism about non-quantum objects otiose. Perhaps some form of reductionist argument could be made to serve this end, but none is forthcoming, and the presumption of such an argument is highly contentious.

Scientific realists are often guilty of ignoring the metaphysical details required to support their epistemological positions. The claim that quantum underdetermination renders certain kinds of realism *ersatz*, however, is premature. Consider macroscopic objects; here too there is underdetermination. Physics underdetermines the choice, for example, between bare substrata instantiating properties and bundle theory, and both accounts are associated with well-known problems (the mystery of somethings-I-know-not-what, spatio-temporal persistence for bundles, etc.). This has not led us to renounce the idea of macroscopic objects, however. In response, French and Ladyman (2003, pp. 50–51, n. 14) argue that there is an important disanalogy between the quantum and macroscopic cases, for macroscopic objects can be experienced 'directly' and identified ostensively. Access to unobservable objects is unavoidably theoretical, so if our best theories underdetermine their precise metaphysical natures, this is reason enough to reject them.

The assertion of this disanalogy is not intended to suggest that realists should believe in the existence of some objects but not others; OSR maintains that there are *no* objects. Rather, the idea is that underdetermination at the quantum level is more severe than underdetermination at the macroscopic level; in the former case only, the very applicability of the concept of individuality is underdetermined. The disanalogy suggested, however, does not bear the weight of this argument. Observation and ostension *also* underdetermine the precise metaphysical natures of objects, including questions of individuality. While it is true that the causal chains that connect unobservable entities to our senses are generally (but not always) longer and more complex than those required to detect observables, this is a matter of degree, not kind. Realists generally think that such differences in degree are not by themselves epistemically significant. Whether an entity is observable or unobservable thus by itself has no bearing on the question of whether it is properly regarded as an individual. In fact, for the proponent of OSR this must be the case, since this person denies the

existence of entities *simpliciter*. For her, macroscopic objects are no more real than quantum objects.

To think that the disanalogy suggested is important to an assessment of whether the concepts of individuality and objecthood are applicable is to bestow an extraordinary privilege on merely putative objects of sensory experience. Certainly we can and do believe in the existence of observed, individual entities prior to philosophical argument and stories about causal chains. But we are concerned here not with pre-reflective beliefs; we are examining considered, philosophical hypotheses. Presumably we should not believe in the existence or properties of ‘objects’ experienced in hallucinations or optical illusions. These pass the test of ostensive presentation, but not that of the justificatory context of scientific realism. In order to distinguish between veridical and non-veridical perception, we must invoke theoretical beliefs, and the force of the disanalogy between our access to observable and unobservable entities is lost. Our considered application of the concept of individuality is theoretical in either case.

I take OSR to be especially valuable in focusing attention on a different question: one may define structure in terms of relations, but what are the relata? The proposal of section 3 is that the relata are causal properties, understood in terms of dispositions for relations. This form of structuralism is foremost a realism about relations of properties. How we then go on to think about objects, the bearers of properties, is an important but secondary question. One reason for not wanting to address the secondary question here is that there is no *a priori* reason to think that just one account should apply across the board. Quantum objects, if there are any, may not be the same *sorts* of objects as macroscopic objects. Some of what qualify as ‘objects’ are countable (proteins, cells), but others are merely quantifiable (plasma, light). Some appear to persist in time, others may exist only in the context of specific events during which their causal properties are instantiated. It is likely that the question of how we get from causal properties to objects is best answered in different ways, depending on the objects in question. Scientific advances shed light on the natures of particular objects, but objects in general no doubt comprise a heterogeneous kind. They may yet, however, serve as a useful category in terms of which to understand reality.¹⁵

As mentioned earlier, even ontic structuralists are happy to talk about quantum objects so long as they are understood structurally. Consider Ladyman’s view:

Objects are picked out by individuating invariants with respect to the transformations relevant to the context. Thus, on this view, elementary particles are just sets of quantities that are invariant under the symmetry groups of particle physics. (Ladyman, 1998, p. 42)

This seems compatible with the idea that the natures of particles can be understood in terms of collections of causal properties (‘invariant quantities’). To the extent that OSR countenances objects conceived structurally, one might wonder whether there is any substantive disagreement between it and the account of structuralism proposed in section 3. There is at least one very important difference. French (2003) holds that entities ‘cannot be regarded as prior to or ontologically separate from the structure which yields them, analytically’. As I suggested previously in connection with Psillos,

however, there is a genuine distinction between the natures of entities and structural relations. On the proposed view, the causal properties we associate with an entity may be present in circumstances where various relations (manifestations these properties confer dispositions *for*) do not obtain. No doubt *some* relation(s) involving it obtains at any time during which an entity can be said to exist, but dispositions and the causal properties that confer them often exist independently of whether the dispositions happen to be manifesting.

5. Optimistic and pessimistic inductions on past science

Realists generally do not believe the entire contents of scientific theories. The form of structuralism outlined in section 3 suggests a different strategy for demarcating parts of theories for which we have reasonable warrant than either ESR or OSR. Only some aspects of theories, I will argue, offer descriptions of structures that are likely to be retained as the sciences move on; other aspects are apt for replacement. The structures to which realists should commit, echoing the most persuasive insight of entity realism, are those involving properties and relations that are reasonably believed on the basis of our causal connections to the world.¹⁶ Let us consider this as a response to the pessimistic induction on past theories.

Many reactions to the pessimistic induction focus on Laudan's (1981, pp. 121–122) influential version of the argument and the list of theories he cites as evidence. Some maintain that if we factor in restrictive criteria (the notion of the maturity of a scientific theory or discipline, the importance of novel predictions, scepticism about *ad hoc* measures, etc.) the data for pessimism are greatly reduced.¹⁷ I will focus here on one argumentative strategy in particular. Even if for the sake of argument we grant the pessimistic induction, realism thrives on a structuralist interpretation. Earlier I suggested that various realists make implicit use of a structuralist strategy in responding to anti-realist pessimism. On this approach, it is granted there is a significant amount of discontinuity in scientific theorizing over time. What there is not, however, is much in the way of radical discontinuity in what is properly believed. The trick is to separate aspects of theories most worthy of belief from those for which we have less warrant. Once this distinction is made, the realist can admit a pessimistic induction on the history of past science *simpliciter*, while simultaneously asserting an optimistic induction on the parts of theories to which she commits.

Several authors have proposed variations on the theme that the important parts of theories for realists are those that 'do the work', that are responsible or indispensable for the predictions, retrodictions and explanations characteristic of mature sciences. Philip Kitcher (1993, pp. 140–149), for example, distinguishes between 'presuppositional posits', the 'idle' parts of theories, and 'working posits', the parts that are required to generate predictions and explanations. Psillos claims that:

it is enough to show that the success of past theories did not depend on what we now believe to be fundamentally flawed theoretical claims...it is enough to show that the theoretical laws and mechanisms which generated the successes of past theories have been retained in our current scientific image. (Psillos, 1999, p. 108).

He considers various incarnations of the caloric theory of heat and 19th century optical ether theories in arguing that both the conception of caloric as a fluid and the ether as an elastic solid, now discarded, were heuristic aids as opposed to essential parts of past theories. The essential bits, he maintains, are retained in current theory.

Case studies are valuable here, but from the realist perspective it is undesirable to hold the tenability of the position hostage to an exhaustive series of studies. And as it stands, the advice to believe the working parts of theories, since these are likely to be retained, is rather vague. An explicitly structuralist approach offers a much more specific proposal for distinguishing between what works and what does not, and furnishes an *a priori* reason for thinking that certain structures will be retained. Descriptions of causal properties do the work. These properties confer dispositions for behaviours that we detect and represent in the form of relations (concrete structures). The capacity to distinguish between parts of theories that do and do not warrant belief thus boils down to an ability to distinguish between parts that we have good reason to think concern genuine causal properties, and those about which we are less sure.

To facilitate a discussion of this, let me introduce a distinction between the attribution of *detection* properties and *auxiliary* properties. Detection properties are causal properties that we have managed to detect; they are causally linked to the regular behaviours of our detectors. Auxiliary properties are any other putative properties attributed to entities by theories. This is an epistemic distinction. Detection properties are the causal properties we know, i.e. properties in whose existence we most reasonably believe on the basis of our causal contact with the world. The ontological status of auxiliary properties is unknown; they may be causal properties or fictions. An auxiliary property is one attributed by a theory but regarding which we have insufficient grounds, on the basis of our detections, to determine its status. Whether the attribution of a property qualifies as a detection or auxiliary attribution depends on the state of scientific inquiry at the time. As the sciences move on, some auxiliary properties are retained as auxiliary, some are converted into detection properties and others are rejected.¹⁸

Figure 1 maps these distinctions among properties. All detection properties are causal properties. To detect is to establish a causal link with the entity under investigation. The attribution of auxiliary properties is non-committal with respect to their ontological status. Further investigation may allow us to detect them, thus converting them into detection properties, or may rule them out altogether. Causal properties themselves are not exhausted, of course, by those attributed to entities by theories. If it turns out that detection content is generally retained as theories change over time and that what is left behind is generally auxiliary, the realist would have a systematic basis for an account of theoretical knowledge, past and present. (I will come to some reasons for expecting this sort of retention momentarily.) A realist could commit to relations of detection properties and remain agnostic or sceptical about auxiliary properties.

Let us probe the property attribution distinction further. The realist requires a practical means of demarcating detection properties from auxiliary properties. Here is a suggestion. Detection properties are connected via causal processes to our means of

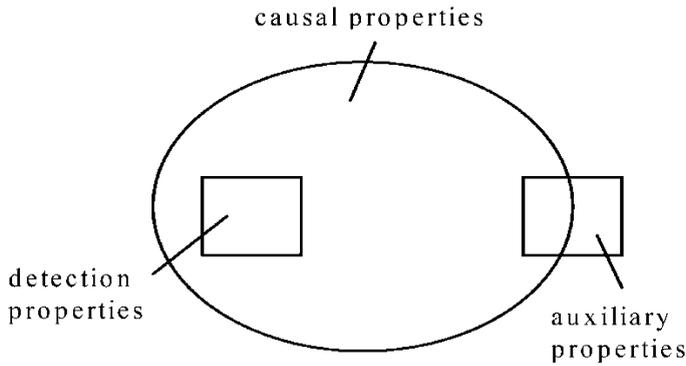


Figure 1. Property distinctions underlying structuralism

detection. We generally describe these processes in terms of mathematical equations that can be taken to represent the relations of properties. As I will attempt to show, we can thus identify detection properties as those that are required to give a minimal interpretation of these sorts of equations. Anything that exceeds such an interpretation, including interpretations of equations that are unconnected to practices of detection, go beyond what is minimally required to do the work of science: to make predictions, retrodictions and so on. The excess is auxiliary.

Recall Worrall's and Poincaré's example of the transition in theories of light from Fresnel to Maxwell. The existence of certain properties is minimally required to give a realist interpretation of the equations shared by these two theories, namely intensities and directions of propagation.¹⁹ These are properties of light. But what about the ether, or the electromagnetic field? In the very limited context of these specific equations, ethers and fields are auxiliary posits. Our theories incorporate such entities as important heuristic devices; they help to fill out our conceptual pictures of the phenomena. The structuralist's advice for the realist, however, is clear: believe in the relations of detection properties and treat anything that exceeds these structures with caution. On this basis, the realist has principled grounds for believing in some, but not all, aspects of theories. It is of course disputable whether, in practice, detection and auxiliary properties can be distinguished in the manner suggested. I will return to this important matter in section 6.

One might wonder why descriptions of detection properties identified by giving minimal interpretations of mathematical equations connected to detection are likely to be retained in some form in later theory. The answer is that these structures are *required*, for the realist at any rate, to interpret the mathematical formalism we use to describe the regular behaviours of our detectors; it is these structures to which we have the closest epistemic access. Only these structures are causally connected in an appropriate way to our means of detection. (We might regard auxiliary properties as 'causally' connected, but in this context inappropriately, in the sense that scientific investigation 'causes' us to posit them as explanatory or heuristic devices.)

Descriptions of the causal properties to which we have the best epistemic access should remain relatively stable as theories are modified and improved.

In fact, the realist is in a position to assert something even stronger than this. In some cases we need to retain specific structures involving detection properties (in some form) if we are to retain the ability to make decent predictions. Consider the theory of electromagnetism. No matter what form descendants of this theory take, we would presumably lose the result that the speed of electromagnetic radiation is c if we did not retain something like Maxwell's equations as a component of the theory.²⁰ Of course, current theories do not retain all of the structures described by their predecessors. But then not all structures are causally connected in appropriate ways to our practices of detection. We should expect to retain only those structures required to give a minimal interpretation of mathematical equations used to describe well-established empirical phenomena. The structuralist thus accepts both optimistic and pessimistic inductions on the history of science. Many contemporary theories will not survive in their current form. Given the important heuristic role played by auxiliary properties, however, we should nevertheless expect and commend their presence in theories.

In a classic paper on inter-theory relations, Heinz Post (1971) argues that heuristic principles typical of scientific investigation promote theories that 'conserve' the successful parts of their predecessors. He thus defines his 'General Correspondence Principle' (p. 228): a new theory generally accounts for the success of the one it surpasses by 'degenerating' into the older theory under conditions in which the older theory is well confirmed.²¹ In one of many similar examples, Laudan (1981, pp. 127–131) claims that there can be no continuity of this sort between classical and relativistic mechanics, for although some laws of the former are limiting cases of laws of the latter, others cannot be, because they invoke an entity that is simply not countenanced by the later theory: the ether. But *of course* there are parts of theories that later theories do not retain. That (in addition to incorporating new content) is the whole point of scientific change. This is hardly surprising or injurious to the structuralist. The ether is part of the auxiliary content of earlier theories. It is not required to give a minimal interpretation of equations describing relations between detection properties. The structuralist has nothing invested in the ether.

It would be disingenuous, however, to suggest that this advice is easily followed. In particular, the notion of a minimal interpretation of equations, on which the identification of detection properties depends, bears great epistemic weight. In many cases a determination of the minimal interpretation may prove a challenge. Furthermore, if structures are retained in some form from one theory to the next, it is desirable that the meanings of terms used to describe the relevant properties and relations likewise extend across theories. This would seem to require a theory of meaning with a referential component, construed at least partly in terms of a causal theory of reference. This invites two concerns. First, one might argue that the task of giving a minimal interpretation is impossible in anything other than retrospect, where hindsight makes the identification of detection properties trivial. Secondly, one might hold that the appeal to some version of a causal theory of reference makes it impossible

that the unobservable terms of most theories could fail to refer. These worries are linked and threaten to turn the proposed structuralism into an empty position, uninformative and trivially satisfied. Let us consider these charges in more detail.

6. The minimal interpretation of structure

Hindsight is 20/20. It is easy to claim that we can identify aspects of past theories that did the real work of science after the fact. We then 'find' that these aspects are retained in later theory. The accusation here is of rationalization *post hoc*. Looking back from the perspective of the present, we are bound to identify aspects of past theories that have been retained as those that do 'the real work'. If this is correct, the realists' commitment to these aspects is suspect; it would seem that she believes in retained elements because they are retained, not because they are more likely to be true or approximately true. And if this were the case, it would be impossible to know at any given time what new structures posited by current theories should be believed. Such a determination could only be made in retrospect.

To defend against these concerns, realists must do more than merely identify belief-worthy parts of theories with parts that are preserved. This merely suggests a correlation, but offers no explanation for why the correlation obtains. Conversely, the proposed account offers a formula for identifying the belief-worthy parts of theories, and this formula is recommended on the basis of its epistemic value. When the realist determines which relations and detection properties are minimally required to interpret the mathematical formalism of a theory, her belief stems from the fact that these are the structures that cannot be denied if one accepts that the theory is reasonably successful in describing parts of the world and their relations to our detectors. These are the aspects of theory for which we have the greatest epistemic warrant, for these are the elements that cannot be done without in making predictions, forming generalizations and so on. (Anti-realists will of course have different views regarding what constitutes an appropriate minimal interpretation and, consequently, what is warranted.) Thus, when we identify detection properties, whether in new theories or old, the structuralist has *a priori* reason to believe that descriptions of them will be retained in some form.

Clearly, however, this would-be recipe for success, the minimal interpretation, will prove difficult in some cases. Theories often describe quite general causal frameworks and the portion of a theory to which the realist should commit may be embedded in a larger framework of this sort. Consider once again the Fresnel–Maxwell case study. Though Fresnel had a particular view of a general causal framework involving the behaviour of light, not all of this understanding is required to give a minimal interpretation of his equations. Fresnel thought that light is a disturbance in the ether, an elastic solid medium. Structuralists are wont to say: 'wait, not so fast, *look at the equations*'. Lacking a clear notion of structure, this is a vague recommendation, but given the concept of concrete structure described in section 3, the realist can make something of this advice.

So let us examine the equations and consider the properties we find described there.

The variables represent amplitudes (intensities) and angles (directions of propagation). But are these not intensities and directions of propagation in the ether? The structuralist is unmoved by this appeal to the greater causal framework. To suppose that a direction of propagation is furthermore a direction in the ether is to go beyond what is minimally required to give an interpretation of this particular set of equations. The structuralist thus commits to relations of intensities and directions and remains agnostic or sceptical about further embellishments.

The reason this prescription may be difficult to follow is that it asks realists to refrain from commitments to parts of theories that play explanatory roles. For Fresnel and other ether theorists the causal story told by their equations is *ipso facto* part of a causal story involving the ether. The structuralist asks for a suspension of belief: in this case we are asked to separate various aspects of the overarching causal story of Fresnel's theory and to believe only those we cannot do without in giving a minimal interpretation of his equations. This involves separating aspects of theories that, for psychological, professional, theological or other reasons, the scientists whose theories they are may find difficult to disentangle. However, it is not impossible. In fact, we know this to be true, for Fresnel's equations were ultimately accepted as part of Maxwell's theory in the context of a non-ethereal physics.

On this account of realism, however, reference to the relevant structures cannot be fixed purely by description. Theoretical descriptions are refined as theories improve, yet we continue to refer to the same properties and relations. It seems that we require an account of reference that has at least a causal component (for example, a causal-descriptive theory). This raises many questions I cannot consider here, but the issue demands attention, given the common complaint that causal theories trivialize reference. If we can ensure that we refer by assigning the referents of terms vague enough causal roles, this would seem to make error impossible, and no realist can claim that her commitment to various structures described by theories is infallible. Let me make a few, admittedly cursory, remarks here.

So long as the structuralist is wary of the charge of triviality, she need not succumb to it. Responding to the claim that the central terms of many past theories are non-referring, Hardin and Rosenberg (1982) suggest that many such terms can be construed as referring on the basis of the causal roles attributed to their referents by the theories in which they occur. This is essentially a variation on Putnam's principle of the 'benefit of the doubt'.²² We can construe Mendel's use of the term 'gene' as referring, even though nothing called a gene in contemporary genetic theory answers to the description Mendel gave, because in current theory configurations of DNA and their polypeptide products perform the same causal role as that attributed to genes by Mendel.

As a general strategy, however, this is too permissive. A realist must know where to draw the line between reasonable applications of the benefit of the doubt and applications that go too far, attributing approximate truth to arcane theories merely because their ontologies were posited to causally account for some of the same phenomena that interest us today.²³ Structuralism provides the beginnings of an answer here. It seems reasonable to give the benefit of the doubt in cases not just where

general causal roles are retained, but where quite specific dispositions for relations conferred by particular causal properties are preserved. On this view it would be unreasonable to apply the principle in such a way as to identify (with one another) the putative referents of significantly different systems of properties.

Consider an infamous case of theory change for the realist: the 18th century transition from Georg Stahl's theory of phlogiston, defended by Joseph Priestley, to Antoine Lavoisier's theory of oxygen. Phlogiston theory accounts for combustion, calcination (rusting) and respiration in terms of the removal of phlogiston. Oxygen theory accounts for the same phenomena in terms of the absorption of oxygen. A distinction between truth and reference is useful here; successful reference does not seem to require that the referring expression is true of its referent (cf. Leplin, 1997, p. 147). Present-day chemistry denies the existence of phlogiston, yet 'dephlogisticated air' might refer, not to air which is lacking in phlogiston conceived as a real substance, but rather to oxygen. The greater the extent to which air is 'low in phlogiston', the greater the extent to which it is rich in oxygen, and vice versa.

This is precisely the sort of conclusion that is dismissed by those critical of the realist's appeal to causal theories of reference. It trivializes reference, they claim, to say that Priestley was talking about oxygen all along. There is a sense in which Priestley *was* talking about oxygen (assuming that the relevant experiments involved *inter alia* the presence and absence of oxygen as opposed to phlogiston), but realism appears ridiculous if we say that that phlogiston theory and oxygen theory are mere linguistic variants. From a structuralist perspective, however, there is no question of identifying Priestley's descriptions with Lavoisier's. For although some of the causal roles described by their theories for dephlogisticated air and oxygen are the same, it is plainly not the case that the putative causal properties of dephlogisticated air are co-extensive with those of oxygen.

For example, oxygen has a fixed chemical composition, but dephlogisticated air *ex hypothesi* does not. Different combinations of gases may lack phlogiston and different combinations have radically different dispositions. The same cannot be said of oxygen. Clearly, we are dealing here with different putative sets of detection properties. As suggested above, the structuralist should thus not accept the claim of identity of reference. The clarity of hindsight may assist, of course, in determining that 'dephlogisticated air' and 'oxygen' should not be thought of as co-extensive. In such cases, however, retrospection is no cause for concern, since the point of the exercise is to consider cases of theoretical change, something that can be done only in retrospect. Structuralism thus provides an antidote to the threat of triviality.

I have outlined a form of structuralism and a version of scientific realism. The central idea is that if we wish to be both structuralists *and* realists we should focus on concrete structures. A knowledge of merely abstract structure, assuming it can be defended against the Newman objection, gives too weak a purchase on reality to constitute much of a realism. The evidence of mathematical continuity across theories commonly cited by structuralists, however supports more than a knowledge of abstract structures. We can know first order properties and relations by minimally interpreting the mathematical equations we use to describe the natural world. On this

basis it is arguable that we can also know conservation laws and other principles (e.g. Pauli exclusion) which are not concrete structures *per se* but which may comprise ways of summarizing aspects of structural relations. These descriptions do not underdetermine the detection properties of entities, but rather the auxiliary contents of theories, and this demarcates plenty of theory about which we may be cautious, agnostic or sceptical, and gives us reason to expect both optimistic and pessimistic inductions on the history of the sciences.

Acknowledgements

Parts of this paper were presented and benefited greatly from discussions at CUNY (New York), the Moral Sciences Club (Cambridge), the annual PSSA Conference (Durban, 2001), the Sigma Club (LSE) and the University of Western Ontario. I am especially grateful to Steven French, James Ladyman, Peter Lipton and Michael Redhead for penetrating, constructive comments. This work was supported by the SSHRC and the Darwin Trust of Edinburgh.

Note on contributor

Anjan Chakravartty (Ph.D. Cantab., 2001) is an Assistant Professor in the Institute for the History and Philosophy of Science and Technology and the Department of Philosophy at the University of Toronto (<http://individual.utoronto.ca/anjan>). His research focuses on central issues in the philosophy of science (realism, antirealism and modeling) and metaphysics (causation, laws of nature and natural kinds).

- [1] See Gower (2000) for an historical account of early structuralism.
- [2] The distinction is due to Ladyman (1998). The same distinction is found in Psillos (2001a) under the labels 'restrictive' and 'eliminative' SR, respectively.
- [3] For a defence of the epistemic view, see Worrall (1989, 1994), Zahar (1996) and Worrall and Zahar (2001). The ontic view is favoured by Ladyman (1998), French (1998, 1999) and French and Ladyman (2003).
- [4] Maxwell (1970b, p. 181), like Russell, believes that 'the observable is instantiated only in inner events of observers'. Contemporary scientific realists generally take observables to be external objects and processes detectable by the unaided senses. However, as Maxwell himself notes, one may adopt the Ramsey sentence approach whatever one's view of where to draw the line separating observables and unobservables.
- [5] I am, however, sceptical of these rebuttals. Redhead's argument is open to the criticism that Newman's objection cannot be dissolved merely by invoking objects with first order properties and relations as a postulate. It is only through some detailed knowledge of these properties and relations that we have substantive (as opposed to trivial) knowledge of Russellian structures (cf. Psillos, 2001b, p. 369). Worrall and Zahar's argument depends crucially on the contestable distinction between observable and unobservable terms. They also accept that no empirically adequate Ramsey sentence can be false, which renders Maxwell's approach unattractive to most realists. For a critique of this approach see Ketland (2004).

- [6] Compare Redhead (2001a, p. 345): ‘The vague reference to mathematical equations is not sufficient to get a proper handle on this notion [i.e. structure]’.
- [7] As will become clear in sections 4–6, the relata I have in mind are determinate properties in the first instance. How kinds of objects are then constituted from kinds of properties is a further question.
- [8] This suggests a commitment to universals, but the language of instantiation is merely expedient. One might refer instead to the class of concrete structures. I will continue to speak of properties *simpliciter*, but everything here might be understood in terms of transcendent universals, immanent universals, tropes or resemblance nominalism.
- [9] For detailed views of this sort see Shoemaker (1980) and Swoyer (1982). In these accounts the identities of properties were determined by the dispositions they confer, but this further claim is not crucial to the current proposal.
- [10] See Mumford (1998) for a comprehensive survey of accounts of dispositions and a defence of realism, and Cartwright (1989) for a defence of dispositional realism (‘capacities’) in the interpretation of scientific theories.
- [11] Maxwell (1970a, n. 19, pp. 33–34) himself mentions dispositions *en passant*. He states that if we were to redefine higher order property terms by means of ‘a viable causal redefinition’, we could predicate them of physical entities. For example, we might redefine ‘red’ so that it refers not to a property of visual experience, but rather to a disposition on the part of the object that appears red to us. But he concludes that such a disposition, though a structural property, would not be a first order property of the object in question. This conclusion is precisely what I deny.
- [12] For details of the physics see French and Redhead (1988), French (1989), van Fraassen (1991) and Huggett (1997). My primary concern is with the possible import of underdetermination, discussed in French (1998), Ladyman (1998) and French and Ladyman (2003).
- [13] See also Bueno and French (1999) and French and Ladyman (2003).
- [14] See Lavine (1991), who argues for a putatively different conception of quantum particles. Redhead and Teller (1992) and Teller (1995, 2001) support the non-individuals view, while van Fraassen (1991) appears to dispute it. For work on related issues see Castellani (1998) and Huggett (1997).
- [15] French (2003) notes Eddington’s view that in the relativistic framework the relata are events, another potentially useful category for collecting groups of causal properties.
- [16] Morrison (1990) argues that entity realists cannot restrict knowledge to ‘low level’ generalizations concerning manipulated entities, since in order to generate the causal processes by which they are known we require further theoretical knowledge. I agree; the proposal here will take a different approach.
- [17] See Worrall (1989, pp. 153–154) for a discussion of maturity and *ad hoc*ness. Psillos (1999, pp. 105–108) discusses *ad hoc* theories and the importance of novel predictions. See also Leplin (1997).
- [18] The distinction between detection and auxiliary properties is central to Chakravartty (1998).
- [19] $R/I = \tan(i - r)/\tan(i + r)$; $R'/I' = \sin(i - r)/\sin(i + r)$; $X/I = (2\sin r \cdot \cos i)/[\sin(i + r) \cos(i - r)]$; $X'/I' = 2\sin r \cdot \cos i/\sin(i + r)$. I^2 , R^2 and X^2 are the intensities of the incident, reflected, and refracted components polarized in the plane of incidence; I'^2 , R'^2 and X'^2 are the components polarized at right angles to the incident plane; i and r are the angles made by the incident and refracted beams with a normal to the plane of reflection.
- [20] This was suggested to me by Steven French (private communication).
- [21] Post’s examples come primarily from the history of physics. See French and Kamminga (1993) for studies which exemplify and discuss Post’s retentionist views, including examples from chemistry and biology.
- [22] See Putnam (1978, pp. 22–25). This is one of two strategies suggested by Hardin and Rosenberg (1982).

- [23] See Laudan (1983) for a discussion of this and related concerns. Niiniluoto (1997) argues that the similarity theory of approximate truth can be used to make the benefit of the doubt precise, but one still requires a criterion with which to establish when terms have the same reference.

References

- Boyd, R. 1981. Scientific realism and naturalistic epistemology. In: *PSA 1980*, edited by P.D. Asquith & R.N. Giere. East Lansing, MI: Philosophy of Science Association 2: 613–662.
- Bueno, O. & S. French 1999. Infestation or pest control: the introduction of group theory into quantum mechanics. *Manuscripto* 22: 37–68.
- Cartwright, N. 1989. *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.
- Castellani, E. (Ed.) 1998. *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*. Princeton, NJ: Princeton University Press.
- Chakravartty, A. 1998. Semirealism. *Studies in History and Philosophy of Science* 29: 391–408.
- Demopoulos, W. & M. Friedman 1985. Critical notice: Bertrand Russell's *The Analysis of Matter*: its historical context and contemporary interest. *Philosophy of Science* 52: 621–639.
- Duhem, P. 1954. *The Aim and Structure of Physical Theory*. Princeton, NJ: Princeton University Press.
- French, S. 1989. Identity and individuality in classical and quantum physics. *Australasian Journal of Philosophy* 67: 432–446.
- French, S. 1998. On the withering away of physical objects, In: *Interpreting Bodies: Classical and Quantum Objects in Modern Physics*, edited by E. Castellani. Princeton, NJ: Princeton University Press 93–113.
- French, S. 1999. Models and mathematics in physics: the role of group theory, In: *From Physics to Philosophy*, edited by J. Butterfield & C. Pagonis. Cambridge: Cambridge University Press 187–207.
- 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. & H. Kamminga (Eds) 1993. *Correspondence, Invariance and Heuristics*. Dordrecht: Kluwer.
- French, S. & J. Ladyman 2003. Remodeling structural realism: quantum mechanics and the metaphysics of structure. *Synthese* 136: 31–56.
- French, S. & M. Redhead 1988. Quantum physics and the identity of indiscernibles. *British Journal for the Philosophy of Science* 39: 233–246.
- Gower, B. 2000. Cassirer, Schlick and 'structural' realism: the philosophy of the exact sciences in the background to early logical positivism. *British Journal for the History of Science* 8: 71–106.
- Hardin, C.L. & A. Rosenberg 1982. In defence of convergent realism. *Philosophy of Science* 49: 604–615.
- Huggett, N. 1997. Identity, quantum mechanics and common sense. *Monist* 80: 118–130.
- Ketland, J. 2004. Empirical adequacy and Ramsification. *British Journal for the Philosophy of Science* 55: 287–300.
- Kitcher, P. 1993. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford: Oxford University Press.
- Ladyman, J. 1998. What is structural realism? *Studies in History and Philosophy of Science* 29: 409–424.
- Laudan, L. 1981. A confutation of convergent realism. *Philosophy of Science* 48: 19–48.
- Laudan, L. 1983. Realism without the real. *Philosophy of Science* 51: 156–162.
- Lavine, S. 1991. Is quantum mechanics an atomistic theory? *Synthese* 89: 253–271.
- Leplin, J. 1997. *A Novel Defence of Scientific Realism*. Oxford: Oxford University Press.
- Maxwell, G. 1970a. Theories, perception, and structural realism. In: *The Nature and Function of*

- Scientific Theories*, edited by R.G. Colodny. Pittsburgh, PA: University of Pittsburgh Press 3–34.
- Maxwell, G. 1970b. Structural realism and the meaning of theoretical terms. In: *Minnesota Studies in the Philosophy of Science*, edited by M. Radner & S. Winokur. Minneapolis: University of Minnesota Press.
- Morrison, M. 1990. Theory, intervention and realism. *Synthese* 82: 1–22.
- Mumford, S. 1998. *Dispositions*. Oxford: Clarendon Press.
- Newman, M.H.A. 1928. Mr. Russell's 'causal theory of perception'. *Mind* 37: 137–148.
- Niiniluoto, I. 1997. Reference invariance and truthlikeness. *Philosophy of Science* 64: 546–554.
- Papineau, D. (Ed. 1996. *The Philosophy of Science*. Oxford: Oxford University Press.
- Poincaré, H. 1952. *Science and Hypothesis*. New York: Dover.
- Post, H.R. 1971. Correspondence, invariance and heuristics: in praise of conservative induction. *Studies in History and Philosophy of Science* 2: 213–255.
- Psillos, S. 1995. Is structural realism the best of both worlds? *Dialectica* 49: 15–46.
- Psillos, S. 1999. *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Psillos, S. 2001a. Is structural realism possible? *Philosophy of Science* 68: S13–S24.
- Psillos, S. 2001b. Quests of a realist. *Metascience* 10: 366–370.
- Putnam, H. 1978. *Meaning and the Moral Sciences*. London: Routledge.
- Redhead, M. 2001a. Quests of a realist, review of S. Psillos. *Scientific Realism: How Science Tracks Truth, Metascience* 10: 341–347.
- Redhead, M. 2001b. The intelligibility of the universe. *Philosophy* 48 (suppl.): 73–90.
- Redhead, M. & P. Teller 1992. Particle labels and the theory of indistinguishable particles in quantum mechanics. *British Journal for the Philosophy of Science* 43: 201–218.
- Russell, B. 1927. *The Analysis of Matter*. London: Kegan Paul, Trench, Trubner.
- Russell, B. 1948. *Human Knowledge: Its Scope and Limits*. London: George Allen & Unwin.
- Shoemaker, S. 1980. Causality and properties. In: *Time and Cause*, edited by P. van Inwagen. Dordrecht: D. Reidel.
- Swoyer, C. 1982. The nature of natural laws. *Australasian Journal of Philosophy* 60: 203–223.
- Teller, P. 1995. *An Interpretative Introduction to Quantum Field Theory*. Princeton, NJ: Princeton University Press.
- Teller, P. 2001. The ins and outs of counterfactual switching. *Noûs* 35: 365–393.
- van Fraassen, B.C. 1991. *Quantum Mechanics: An Empiricist View*. Oxford: Clarendon Press.
- 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*, edited by D. Hull, M. Forbes & M. Burian. East Lansing, MI: Philosophy of Science Association 1: 334–342.
- Worrall, J. & Zahar, E.G. 2001. Ramseyfication and structural realism, In: *Poincaré's Philosophy: From Conventionalism to Phenomenology*, E.G. Zahar. Chicago, IL: Open Court 236–251.
- Zahar, E.G. 1996. Poincaré's structural realism and his logic of discovery, In: *Henri Poincaré: Science and Philosophy*, edited by J.-L. Greffe, G. Heinzmann & K. Lorenz. Berlin: Akademie Verlag & A. Blanchard 45–68.

