

Realism in the Desert and in the Jungle: Reply to French, Ghins, and Psillos

Anjan Chakravartty

Received: 15 November 2012 / Accepted: 15 November 2012 / Published online: 4 December 2012
© Springer Science+Business Media Dordrecht 2012

Abstract *A Metaphysics for Scientific Realism: Knowing the Unobservable* has two primary aims. The first is to extract the most promising refinements of the idea of scientific realism to emerge in recent decades and assemble them into a maximally defensible realist position, *semirealism*. The second is to demonstrate that, *contra* antirealist scepticism to the contrary, key concepts typically invoked by realists in expounding their views can be given a coherent and unified explication. These concepts include notions of causation, laws of nature, scientific kinds, and approximate truth, and consequently, the demonstration undertaken includes a metaphysical study of ideas more commonly employed unreflectively in epistemological assessments of the sciences. In this paper, I answer searching critiques of this project by Steven French, Michel Ghins, and Stathis Psillos.

1 Introduction: Semirealism

Scientific realism is an epistemic attitude towards the sciences, which takes the content of our best scientific theories to furnish knowledge of both observable and unobservable aspects of the natural world. This attitude has been the subject of a great deal of elaboration in recent philosophical history, punctuated by a surge of attention three decades ago following the demise of logical empiricism and the historical turn in the philosophy of science. This elaboration was and continues to be necessary, for the plausibility of the schematic characterization of realism given above is difficult to assess otherwise. One might wonder, for example, how one should determine which theories are our best, or what scientific knowledge amounts

A. Chakravartty (✉)
Department of Philosophy, University of Notre Dame, 100 Malloy Hall, Notre Dame,
IN 46556, USA
e-mail: chakravartty.1@nd.edu

to given the fact of theory change over time and the widespread use of idealizations in theories and models.

In a recent book (Chakravartty 2007, 'MFSR' henceforth), I sought to do two things, the first of which was to identify the most promising aspects of recent elaborations of scientific realism and fuse them together. I called the resulting position 'semirealism', in recognition of the fact that sophisticated, contemporary forms of realist commitment are characteristically and appropriately modest or measured, not least so as to survive pressing antirealist arguments that might otherwise skewer more naïve formulations of realism. In particular, I argued that attempts to refine realism by directing a positive epistemic attitude towards selective parts of our best theories—parts that most merit it—have taken important steps forward in rendering realism plausible. Positions such as entity realism and structural realism are selectively pessimistic about parts of theories they identify as least deserving of epistemic commitment, but *eo ipso* selectively optimistic about others. Semirealism is an attempt to marry the progressive insights of these views and leave behind their weaknesses.

My second goal was to give a coherent and unified account of the most important concepts generally invoked by realists in expounding their views, in part because the idea that there *can* be such an account *at all* is also a target of antirealist scepticism. For example, sophisticated versions of realism often appeal to ideas such as causal manipulation and intervention, or the discovery of laws of nature and natural kinds of entities and structures. The recurrence of these themes suggests their importance to realism, and it is for this reason that I gave significant attention to the metaphysical concepts underlying an otherwise epistemological position, for such is the centrality of these concepts and the weight they bear, that the position is not secure, I submit, so long as the coherence of them is in doubt. Thus, the realist conception of scientific knowledge naturally invites and arguably demands some ontological clarification, to make plain what a realist commitment commits one *to*, more precisely.

French (2013), Ghins (2013), and Psillos (2013) pose a number of serious challenges to both aspects of this project. On the metaphysical side of the balance, concerns are raised about my adoption of a dispositional account of causal properties, which I offer as an underpinning for semirealism. Both the ontology of dispositions and my contention that causal properties are generally intrinsic are problematized, as is my claim that many and perhaps most of what one calls laws of nature are relations between such properties. These and related concerns, about what view of particulars best suits realism, as well as my suggestion that laws can be regarded as "concrete structures", will be the focus of Sect. 2. I tackle these metaphysical issues first because in doing so, I am able to clarify certain aspects of the foundations of semirealism proposed in MFSR that may helpfully inform an understanding of the position itself. But those whose primary interest is in the epistemology of science and scientific realism more narrowly construed should feel free to skip straight to Sect. 3.

In Sect. 3, I turn to a number of epistemological challenges raised by French, Ghins, and Psillos, the first of which concerns the nature of explanations of the occurrence of kinds, especially in the realm of subatomic physics. A second worry

targets the extent to which the causal theory of reference may be applied in the service of realism, and a third the coherence of a distinction I introduce—between ‘detection properties’ and ‘auxiliary properties’—to help draw a line between the components of scientific theories about which realists should feel confident, and those about which they might reasonably be wary. A final challenge queries the very stability of the combination of epistemology and metaphysics to which I aspire, thereby raising important questions about philosophical methodology generally, and more specifically in the context of the philosophy of science.

Thus running the gamut from the minutiae of ontology through grand questions about the nature of philosophical inquiry, I will end with a few remarks on the relationship between metaphysics and the philosophy of science. For it is a consequence of my account of scientific realism—expected and I think unavoidable, but welcome—that some differences between realists and antirealists, as well as between realists of different stripes, may of necessity remain unresolved. Nowhere is this more likely than in connection with the metaphysical aspects of realism. It is nonetheless crucial to the plausibility of the position to see if it can be rendered in *any* way a conceptually coherent attitude towards the sciences. If in the end we find that there is more than one way to achieve this, more power to us all. But before becoming overly excited by this prospect, let us turn to the first item on the agenda: the metaphysical foundations of semirealism proposed in MFSR.

2 Metaphysical Challenges

2.1 Properties, Part I: Dispositions

One central feature of semirealism is that it is a realism about well-detected properties in the first instance. This, I believe, is the bedrock of any plausible realist epistemology. The greater our apparent abilities to interact with and otherwise manipulate such properties, especially in the case of the strictly unobservable, the greater our epistemic warrant for believing in them. The intimate relationship between epistemic warrant and causal connections to properties led me in MFSR to consider the nature of these properties in some detail. I argue there that a particular view pays rich dividends in this context: a dispositional view of causal properties naturally yields a unified elaboration of several concepts common to realist discourse, including laws and kinds, and a straightforward explication of the notion of *de re* necessity. The dispositional view holds that the identity of a causal property—that which makes it what it is—is exhausted by the dispositions for behaviour it confers on whatever has it.

One interesting consequence of this dispositional identity thesis is that the natures of causal properties admit of a form of holism. Dispositions are “powers” to stand in various sorts of relations; a specification of all such relations would thus constitute an exhaustive description of the natures of all causal properties (and thereby, any given causal property). This holism, Psillos argues, is pernicious. For on this view, he contends, ‘what a property is cannot possibly be identified unless what all other properties to which it is related are has already been specified; that is,

unless all other properties have already been identified. But since this tangle arises for *any* property whatever, it follows that no property can be identified unless some other properties have already been identified, and because of this, no property can be identified *simpliciter*'. Furthermore, if the identity of properties is in this way holistic, then surely it is laws (relations between properties) that determine the identities of properties, not dispositions. Thus, Psillos maintains that the dispositional view of property identity makes identity impossible to determine, and simultaneously undermines itself by inadvertently making the relations of properties, not dispositions, the determinant of property identity.

These would be serious consequences indeed, but neither follows from the dispositional identity thesis. First note that, as Psillos himself suggests in his further remarks, it is crucial here to distinguish between epistemic concerns—those regarding how properties are identified and individuated in actual scientific practice—and metaphysical concerns—regarding the ontological identities of properties quite independently of anything we might do. The dispositional identity thesis entails a form of holism with respect to the latter, but nothing with respect to the former. A scientist might detect a property by means of eliciting only one of the many relations it can manifest under different laboratory or experimental conditions, and that may be sufficient for identification in the epistemic sense. But this has little to do with the identity conditions of this same property in the metaphysical sense. In just the way that I might recognize my friend Stathis through a crowd by glimpsing his face and no other part (and though I suspect the face is merely one aspect of the man), one might recognize the presence of a mass by means of its relations to a balance quite independently of the many other possible relations into which that property can enter.

This resolution to the first worry suggests a resolution to the second, according to which the holism of property identity in the metaphysical sense makes laws of nature the determinants of property identity. Again, however, this is no consequence of the dispositional identity thesis. It is not relations between properties (laws) but dispositions for relations that determine the identity of a causal property on this view. Psillos expresses consternation regarding the idea that dispositions for relations could be “fixed” by anything other than the relations themselves; but what, one might ask, is “fixing”? There are two natural readings of this. Perhaps ‘fixing’ means something like ‘epistemically apprehending’, but if so, as we have seen, there is nothing here at odds with the dispositional identity thesis. Indeed, how else might one learn about a property except by standing in some relation to it, however direct or indirect, and via whatever means of detection (sensory modalities, scientific instruments, etc.) one brings to bear? On the other hand, perhaps ‘fixing’ means ‘determining the identity of’ in the metaphysical sense. But if so, this worry begs the question. For why should the identity of a causal property be determined by its relations? Antirealists about dispositions might think it so—it is a storied (and ultimately unsuccessful) empiricist project, for example, to dissolve disposition talk into talk of manifest relations—but this is simply to beg the question.

This last point reveals a common assumption about dispositional properties that is, I think, misleading in the present context. Many authors identify causal properties with *relational* properties, and given that dispositions are generally described in terms of relations that would be manifest if certain conditions were to

obtain, it is no doubt natural to do so. On the assumption that ‘relational’ is here opposed to ‘intrinsic’, however, this is not what I myself intend most generally. The dispositions of entities studied in the course of scientific work (*inter alia*) are often intrinsic properties. Unlike relational properties (here defined), they are not properties that something has in virtue of standing in a relation to something else, and this is the case independently of the fact that our epistemic access to them comes by means of their relations to other properties, including those of our instruments of detection, our visual apparatus, and so on. We often learn of the intrinsic natures of things by fashioning relationships to them; that is how we learn.

Does this threaten the idea of mind-independent knowledge so dear to scientific realists, as Psillos warns? It does not. For while it may be part of the nature of a property that it confers a disposition to stand in a particular relation to another property that may happen, contingently, to be exemplified by one’s sensory apparatus, the existence of such an observer is hardly a necessary condition for the existence of the original property (that, as it happens, may or may not be observed). Neither is the nexus of relations which things in the world are disposed to manifest a mystery, though Psillos worries it must be. What is the source of this putative mystery? Once one pays attention to the distinction between the epistemic and metaphysical dimensions of property identification, the former is revealed as no more mysterious than our everyday practices of detection, and the latter no more mysterious than the notion of a disposition. I am the first to admit, however, that dispositions are not for everyone (Chakravarty 2007, p. 118). I will return to this point in Sect. 3.4.

2.2 Properties, Part II: Intrinsicity

In clarifying the dispositional identity thesis for causal properties, I have just suggested that they are often intrinsic properties of the particulars that have them. French, however, wonders whether this can be so, and his worry is tendentious. If one could show that there are no such things as intrinsic properties, and demonstrate instead that all properties are relational, this might be taken to point the way towards ontic structural realist views of physics and science more generally, according to which at the fundamental level of ontology, there are only structural relations and nothing that can be said to be related *per se*—or at least nothing related that is not dependent in some deep sense (to be explained) on the relations in which they stand. I will not consider arguments for and against ontic structural realism here, but a few remarks on the idea of intrinsicity are called for in defence of my proposed basis for semirealism.

French begins with some scepticism regarding the ways in which the notion of an intrinsic property is commonly described. Thought experiments in which one designates all and only those properties a particular would have if alone in some possible world are unsatisfying, he suggests, insofar as they abstract particulars away from the physical context that one might think furnishes necessary conditions for their very existence. Is it an instance of a promising philosophical methodology to imagine the property of mass abstracted from the framework of general relativity? If in the actual world, elementary particles typically participate in certain

kinds of relations, what confidence should one have in the conclusions of thought experiments that describe them in the absence of those relations, inhabiting (for all we know) physically impossible worlds? Such scepticism is healthy and should, I think, give one pause. It does not, however, serve to undermine the idea that properties of scientific interest are often intrinsic. The care that one should rightly bring to considerations of putatively possible worlds concerns the epistemological project of classifying properties into types—intrinsic, relational, what have you. But this by itself yields no reason to think that there are only relational properties, or that commonly cited examples of intrinsic properties are not. Let us see why this is so.

The intuitive notion of an intrinsic property is that of one possessed by something independently of externalities; an intrinsic property is possessed in virtue of the way something is in itself, not what other things are. Granted, this informal sketch is also imprecise, and attempts at precision here and elsewhere in this territory engender intricate challenges, but let us persist with an informal conception presently. A qualification must be introduced immediately if this conception is to avoid being trivially false. Let us say that my friend Steven is very good at math, and let us use the term ‘cleverness’ to label the property (or properties) of his brain that confers this ability. Intuitively, his cleverness is an intrinsic property, but this is not to say that it is exemplified in the absence of enabling conditions. These conditions can be described in terms of myriad relations to supportive parents in childhood, a fine education, environmental factors such as the presence of sufficient oxygen to sustain life, and so on, but Steven’s cleverness is nonetheless a feature of him in himself, and thus, intrinsic. Therefore, the idea that relations may be required in order to bring a property into being does not entail that the property is relational. Even if one is resistant to this intuition, the fact that some properties persist through changes in their relations furnishes some evidence of their intrinsicity. Let us see how this plays out in the scientific context.

French imagines a possible world inhabited by a lonely charged particle. But how can one say, he wonders, what properties this particle has? In the absence of other things, does it really have charge, obey Coulomb’s law, and so on? The only way to determine these things would be to bring in a test charge; that is, to put the subject charge into some relation with something else. Only then could one ascribe the relevant property. As an objection to the notion of intrinsic properties, however, this is to conflate epistemic and metaphysical questions in a manner similar to Psillos. On the dispositional view, a particle’s charge is something it possesses independently of its interactions with test charges—that is a metaphysical proposal. *How one comes to know* its charge is another matter, and may well require experiments (either real or in thought). To think that the relations manifested in such experiments somehow “make” charge the property that it is, however, is once again to beg the question. And even if it turns out that it is not physically possible for there to be a charged particle unless it exists in some background space, the relation of charge to this space is much like the relation of one’s cleverness to enabling conditions such as good genes and oxygen. Furthermore, the fact that a charged particle can survive changes in its relations involving charge (imagine a test charge at infinity, and then very close by) is some evidence of intrinsicity.

The hallmark of a relational property is that its existence depends on some relation obtaining in a constitutive (as opposed to a merely enabling or otherwise non-constitutive) way. If I were to move to another neighbourhood, the kindly old lady next door would no longer be my neighbour. But physics gives us no reason to think that the same is true of charge. If one were to take the test particle away (from the world, even), would the subject particle no longer have charge? On the dispositional view, the subject particle is still disposed to interact with any number of test particles in ways described by Coulomb's law, and there is nothing in the physics of such particles to suggest otherwise. Worse, the suggestion that it would not have charge in the absence of test conditions once again courts the empiricist's morass of attempting to dissolve dispositions into manifest relations. I submit that on an informal consideration of the notion of intrinsicity, there is no debilitating concern here regarding the intrinsicity of many properties of scientific interest.

2.3 Particulars

Having made the world safe for intrinsic dispositions, let us turn now to the particulars that have them. Recall that semirealism is a realism about well-detected properties in the first instance, and though in MFSR I suggest that the inference to the existence of particulars is inevitably weaker, where properties are systematically observed or detected as clustering in regular ways, talk of entities that have them reasonably enters the realm of realist ontological commitment. Ghins, however, presents two reasons for thinking that this is problematic. The first stems from his suspicion that the semirealist picture of particulars is tantamount to a bundle theory of objects, and that such theories are untenable. The second concern, connected to the first, stems from his contention that semirealism inevitably amounts to a form of Platonism, and that this too is untenable. Let us consider these charges, in turn.

Talk of particulars here is talk of concrete entities in the domain of scientific knowledge, but the very idea of "concreteness" is one that admits of ancient (and some more recent) dispute in metaphysics. Given that semirealism privileges an epistemic commitment to well-detected properties above all else, it is natural, no doubt, to think of it as meshing neatly with some version of the bundle theory of particulars. For if belief in certain properties identified in scientific investigation is the thing for which one has greatest epistemic warrant, it may appear that a conception of particulars on which they are merely bundles of such properties is conducive to ontological economy—it extends ontological commitment minimally, one might think, in comparison to other theories of particulars such as substratum theories, which posit yet further denizens of fundamental ontology (such as bare substrata). I have some sympathy for this line of thought, but it is not so strong as to amount to a commitment, and certainly not one on behalf of semirealism.

Once one engages in a metaphysical study of the core concepts implicated in scientific realism (or any epistemological position, for that matter), as I do in MFSR, one thereby opens the door to a potential regress of explanation. Having attempted to explicate the notions of causation, laws, and kinds in terms of an underlying ontology of dispositions, the stage is then set to ask further questions about the precise natures of the items featuring in the *explicans* in terms of yet more

fundamental ontology. I would not rule out of court these further metaphysical pursuits *ex cathedra*, but *qua* philosophy of science, it seems to me that one need go no further. On the assumption that there *are* internally consistent and coherent accounts of these more fine-grained ontological *explicanda*, from the point of view of scientific realism, any will do. Of course, merely assuming consistency and coherence does not make them so, but at this very fundamental level of ontology one reaches conceptual bedrock, and modulo the different primitive concepts adopted by the adherents of different approaches here, and after more than 2000 years of refinement, I have no doubt that these approaches are consistent and coherent by their own lights.

For example, consider Ghins' dissatisfaction with the bundle theory and his insistence instead on something more like the substratum theory. The bundle theory, he suggests, cannot deliver concreteness, whereas the notion of a substratum (or in Ghins' terms, 'an ingredient which is not a property') can. But given the entirely primitive nature of the concept of a substratum, what seems helpfully explanatory to the proponents of substrata seems hopelessly ad hoc and unilluminating to the proponents of bundles. This pattern repeats itself, for at this most fundamental level of ontology, conceptual bedrock must be fused together with primitives of one form or another, and the choice here extends beyond that which a realist *qua* scientific realist need make. It suffices that there exist *explicantia* that are not self-undermining, and this, to the neutral, must seem evident. Ghins worries that if a gas, for instance, is merely a group of properties, then its volume is either a second-order property (a property of some properties) and thus not a property of a particular, or it is a component property of the gas, in which case its conferral of dispositions to the set becomes a mystery. But no bundle theorist would agree: on this view, a group of properties arranged thus and so *is* a particular, and the participation of a causal property in that group affords certain dispositions *to* that particular. If one is troubled by the putative abstractness of these properties, think of them as tropes. And so on.

Ghins sees a form of Platonism naturally emerging from a commitment to the bundle theory, but even if this is so, having dispensed with any necessary commitment to bundles, there is no pressure here for a semirealist to be a Platonist. And importantly, the idea that dispositions are occurrent properties independently of whether they are manifested—a view I do indeed develop—does not by itself connote Platonism. For manifestations are simply the coming to pass of relations for which these properties are dispositions, and whether or not these relations do come to pass, dispositions might exist in an Aristotelian or trope-theoretic manner all the same.

2.4 Laws of Nature and Concrete Structures

In MFSR, I characterize laws as relations between properties. This integrates neatly with the account of causal properties and processes I describe, thus serving the aim of producing a unified account of the conceptual apparatus of semirealism. If these relations are laws of nature, then statements describing them are statements of law, or law statements. It should be clear I hope that these uses of the terms 'law' and 'law statement' are simply one sort of regimentation of terminology. There are, of

course, other sorts of things that are often called law statements in the sciences—things that do not appear to describe relations of properties directly *per se*, but rather describe the putative behaviours of kinds (‘mammals give birth to live young’) or the putatively characteristic properties of kinds (‘mammals are warm-blooded’). I call these ‘behavioural generalizations’ and ‘definitional generalizations’ respectively, and together with law statements in my sense, take them to exhaust the space of things commonly called statements of law in the sciences.

Ghins takes issue with my identification of laws with relations between properties. Rather, he thinks that the surface form of law statements that appear to quantify over properties is in fact elliptical for genuine law statements that quantify over particulars. For example, the Boyle-Mariotte law is often written in a form that suggests a description of properties, or more specifically, the constant product of values of pressure and volume at a constant temperature ($\forall p \forall V (pV = K)$), but this is merely elliptical, he contends, for a description of the behaviour of particular gases ($\forall x (Gx \rightarrow \forall p \forall V (px \times Vx = K))$). It is unclear to me, however, what this latter interpretation of law statements achieves, other than conformity with an extra-scientific commitment to the idea that laws must concern particulars rather than properties. Acknowledging as I do that ‘law’ is a term of art, I do not suppose that anything philosophically momentous hangs on how one chooses to regiment it, so long as one can account for the various uses of it in scientific discourse. And this, as I suggest above, is something I take myself to have done.

French has a different axe to grind with respect to laws, targeting the idea that one can or should invoke properties (and *a fortiori*, dispositions) at all in giving an account of them. Recall that the ontic structural realist aims to give a description of fundamental physics in terms of relations that have some form of ontological priority over their putative relata, and thus, the prospect of laws in the absence of dispositional properties they might otherwise relate is for her a desideratum. French offers two considerations in support of the idea that one might happily keep the laws and dispense with the properties. The first appears to be that in the context of basic entities like elementary particles, their relations are always manifesting. This is presumably offered as disanalogous to the case of other entities, in which one might reasonably view properties as having an “anchoring” role: the relations of which other entities are capable are not all always manifesting; arguably, some intrinsic properties serve to anchor the existence of such entities in the face of changeable relations. But if the relevant relations are always present in the case of elementary particles, what need is there of an anchor here? French’s second consideration is a worry to the effect that if dispositions and the laws relating them are ontologically distinct, there is a “metaphysical gap” between them, in which case one might reasonably ask how the former determines the identity of or otherwise “governs” the latter. In the absence of a response, why not again simply dispense with dispositions altogether?

Neither of these considerations, I believe, undermines an ontology of dispositions, or of properties *simpliciter*. Even if it were true that the relations of fundamental physical entities are always manifesting, one might nevertheless favour the economy of an account of properties and laws that applies across the sciences, not merely to basic physics. The account of properties and laws I elaborate has this virtue. More

importantly, the premise that all the relations of which the entities of basic physics are capable are always manifesting seems false. There is a level of abstraction at which one might reasonably think it true: for example, assuming there are such things, one might think that an elementary particle is always manifesting gravitational attraction, conceived as a determinable relation. More concretely, however, the determinate magnitudes of this and other such parameters are functions of changeable relations to other particles; what is to serve as an anchor then? Only if all the currently manifesting relations of basic physical entities were frozen in time could these relations anchor in the way properties do, and therefore, since our universe is not the frozen kind, it appears they do not. French hopes to replace dispositions for different relations manifested in different circumstances with ‘the modal abstraction of certain aspects of structure’. One might worry that this is to replace an admittedly abstruse notion with something significantly more so.

What of the worry that given the assumption that laws and the dispositional properties related by them are distinct things, it is a mystery how laws can govern things with dispositions? The concern here is that, since the notion of a disposition already includes the idea of associated behaviours in appropriate circumstances, laws are somehow shorn of their proper function—is not governing the phenomena, after all, part of the very conception of lawhood? It is no part of my conception. Laws are simply relations between properties, and these, I take it, amount to a large proportion of the things called ‘laws’ in scientific discourse. They need contribute no modal force, however that is analyzed, above and beyond that which is already supplied by dispositions. They are nonetheless distinct things, for a disposition for a relation is not the same thing as a relation, and so long as the latter is of interest in scientific contexts, I see no reason to dispense with it as a useful concept in describing scientific knowledge.

To conclude this section with one last concern on the topic of laws, Psillos takes issue with my suggestion that laws conceived this way can be thought of as “concrete structures”. The idea here is to distinguish knowledge of structures that a realist might reasonably aspire toward—qualitative knowledge of the relations of first-order, causal properties of things—from knowledge of higher-order mathematical or logical properties that some structuralists contend is the best anyone can hope for. A concrete structure is one that relates particular *kinds* of relata. Knowing a concrete structure thus involves knowing something about qualitative properties and relations (for example, opposite charges and electrostatic attractions, or gaseous pressures and volumes and their inverse proportionality at constant temperature), as opposed to merely formal properties and relations (for example, the relation of total ordering, which is shared by the qualitative relations taller-than and shorter-than). But surely, Psillos wonders, ‘isn’t the very idea of a *concrete* structure an oxymoron?’ Structures are by their nature abstract: they are things that different concrete systems can exemplify. To imagine such a thing as concrete might well be ‘neo-Aristotelianism gone wild’.

The term ‘concrete structure’ is a technical term in MFSR, variably interpretable on the basis of one’s preferred ontology of properties and relations. The central idea is that whatever more fine-grained ontology one prefers, what I call ‘concrete structures’ are less abstract—that is, closer to the realm of the concrete—than the

highly abstract, formal structures invoked by some structuralists. I am myself uncertain whether talk of “degrees” of abstraction should be understood figuratively or literally here, but in any case, how precisely what I call concrete structures are described in terms of concreteness and abstractness will vary, according to commitments more subtle than semirealism requires.

For instance, on a theory of transcendent (Platonic) universals, the taller-than relation is abstract, for it can be multiply exemplified by different particulars; the total-ordering relation is a second-order abstraction, for it is exemplified by (abstract) things such as the taller-than relation. The taller-than relation is thus less abstract, and thereby closer to the realm of the concrete, than the total-ordering relation. This is all I mean to convey with the term ‘concrete structure’. On a theory of immanent (Aristotelian) universals, what I call concrete structures exist only when exemplified by concrete particulars, but they are no less abstract for that. (What Psillos calls neo-Aristotelianism gone wild, I suspect, is simply Aristotle on a regular day.) The trope theorist will regard what I call concrete structures as relational tropes (particulars), which may be less or more (up to exactly) similar to one another. The unremitting nominalist will view them as classes of concrete relational systems.

These different elaborations of the notion of concrete structure are all on the table, and I express no preference on behalf of semirealism. They each face challenges, and their advocates resolve them in their own ways. Quite independently, however, I hope the idea of a structure conceived as a relation between first-order causal properties, such that a knowledge of it entails a knowledge of some qualitative properties and relations of things in the world, is clear enough for realism.

3 Epistemological Challenges

3.1 Scientific kinds, Sociability, and Symmetry

Let us now turn to the various epistemological issues raised by my interlocutors, beginning with the putative explanatory value of the notion of natural kinds in discussions of scientific realism. The appeal to kinds in this context functions primarily as a buttress for inductive inference. If there are naturally occurring categories of things in nature—even better: whose members share a common essence—and if one thinks that through scientific inquiry we discover these categories (and essences), what better ground could there be for successful scientific generalization and prediction? One of the core themes of MFSR, however, is that traditional conceptions of kinds in terms of essences, and the independent but also traditional view that there is one correct kind taxonomy of the natural world, are hopelessly outmoded as accounts of modern scientific classification. In place of essences I suggested a new metaphor: the idea of “sociability”. It is an empirical discovery that in some cases, certain groups of properties appear to cohere spatio-temporally with invariable regularity (these are cases in which talk of essences is most at home), but in other cases, groupings are regular enough to support degrees of

inductive success without invariability of association. Thus sociability admits of degrees, and I argue that this picture better accords with modern scientific taxonomy.

French wonders whether the idea of sociability is explanatory of the clustering of properties associated with kinds, and more specifically, suggests that in the case of basic physics, it is not. In this domain, he argues, it is symmetry that is explanatory of kindhood. Sociability is at best a labelling of phenomena that must be cashed out in terms of symmetry principles if one is to have a genuine explanation of the associations of properties described by the Standard Model of elementary particles. Certain quantities are invariant under specified symmetry group transformations—as represented by the permutation group, for example, in the case of bosons and fermions—and this is what ultimately explains the fact of the sociability of properties associated with the relevant particles. Sociability by itself does not seem to explain much in the absence of this more detailed group-theoretic understanding of elementary particles and their properties.

Is sociability an empty explanatory device? Though I am sympathetic to French's invocation of symmetry considerations in the taxonomy of subatomic physics, I do not think it diminishes the concept of sociability. In order to assess the merit of a given *explanans*, one must of course consider it in relation to an appropriate *explanandum*, and there are at least two distinct *explananda* here. Recall that the idea of natural kinds is primarily employed in the context of scientific realism to help underwrite a degree of optimism regarding candidate knowledge claims produced by scientific generalization and prediction. In explaining the success of these sorts of inductive practices, sociability is an effective component of what seems a compelling *explanans*: entities behave in certain ways in certain circumstances as a function of the (causal) properties they possess; therefore, the greater the extent to which the members of a class of entities share (causal) properties, the greater the success one should expect of inductive generalizations and projections over their members. As a measure of the degree to which properties are shared, sociability is thus directly correlated with this success, and serves well as part of its explanation. The remaining part of the *explanans* comprises whatever account one gives of how the sharing of properties results in shared natures and behaviours—in terms of dispositions, laws of nature, or what have you.

Another *explanandum* of interest is that corresponding to the question of why any particular measure of sociability applicable to a scientifically interesting class of entities is exhibited in the first place, and here the concept of sociability does no work. Indeed, there is no expectation here that the concept *should* explain its own instances. Rather, it seems likely that explanations of instances of sociability will require significantly different types of *explanans*, depending on the phenomena at issue. Some explanations of sociability are given in terms of underlying causal processes, hence fuelling causal explanation in these cases. Other explanations of sociability may exemplify different types, and so far as our current best physics is concerned, the sociability of certain properties of elementary particles is a case in point. Now, the question of whether in this case the relevant symmetry principles are explanatory will turn on the question of whether they are part of an *explanans* exemplifying a defensible form of explanation, and given the highly unificatory employment of symmetries in the Standard Model, one might argue that they

facilitate a unificationist explanation of these particular instances of sociability. But whatever one's views regarding whether or not this would constitute a genuinely explanatory unification – a fascinating subject that I will not examine here—the modest point remains that the explanatory aim is something other than that for which the concept of sociability is designed.

3.2 Theory Change, Part I: The Causal Theory of Reference

However well the sociability of properties helps to explain the success of inductive practices in the sciences, it is reference to the entities conceived as having these properties—objects, events, and processes—that is more often targeted by antirealists. Since scientific knowledge is often described in terms of such entities, referential discontinuity across theory change in the history of a given domain, so the argument goes, presents a *prima facie* reason for scepticism about such knowledge at any given time. It is for this reason that many realists are drawn to the causal theory of meaning and reference, for this approach yields continuity of reference across radical discontinuities in theoretical descriptions of entities, so long as they continue to be viewed theoretically as causally responsible for the same observable phenomena. Ghins has a great deal invested in this view, contending that ‘a particular is more than a coherent grouping of causal properties at some location [as suggested by semirealism]... It is also something that we identify as a “this” or a “that” in actual perception’; ‘These things are the concrete entities that are first given to us as “this” in actual perceptual awareness’. In other words, his realism is constitutively linked to the idea of reference not merely to properties but to particular entities that have them, which are connected in some way to perception.

As a consequence of this commitment, Ghins is critical of the semirealist idea that one should admit referential continuity only in cases where fairly specific dispositions for well-detected relations are preserved across theory change, not merely in cases of the retention of more general or more vaguely described causal roles in processes ultimately linked to observation. Is it not the case, for example, that certain perceptions were generated in scientists who worked with cathode ray tubes, and that these same perceptions are generated in connection with electrical phenomena in the laboratory today? Surely the term ‘electron’ thus simply refers to whatever it is that has the causal properties necessary to produce these observations. As Ghins puts it, ‘the causal properties of electrons are sufficiently minimally identified as the ones which are responsible for the occurrence of the perfectly identified luminous *phenomena* in precise laboratory conditions’.

Despite the unmistakable appeal of this strategy as a means by which to respond to Kuhnian pronouncements of referential discontinuity across theory change, I submit that it comes at a cost so high as to render it a poison chalice for realism. While there is no doubt something correct in saying, for instance, that one may regard the term ‘electron’ as referring to an entity detected from the late nineteenth century to the present, there is also a sense in which this is seriously misleading. One clue that there is something not quite right here is revealed by the extent to which theoretical characterizations of what we now call ‘the electron’ have changed over time. If realism is reduced to *mere* continuity of reference, it is very much

denuded, amounting only to the disappointing claim that there exists something, about which the realist can say nothing—other than the fact that it is linked (somehow!) to perception. One might wonder whether this is a realism worth fighting for, and there is worse news. If all realism need amount to is the claim that reference is preserved—independently of some significant knowledge of properties and relations—then realism is virtually guaranteed to be true so long as one specifies the relevant causal role or roles vaguely enough, for the claim that there is *something* (or some things) causally responsible for a given observation is almost undeniable. What began as an epistemically substantive position is now almost certainly true a priori, and thus, utterly trivial.

It is with this in mind that I insist that realists should not be so quick to identify the referents of ‘oxygen’ and ‘dephlogisticated air’, for example. Perhaps one should, but that is a judgment call, to be made on the basis of whether the putative referents of these terms are described by their respective theories as having sufficiently similar sets of well-detected properties. In MFSR, I argue that in this particular case, the sets are so different as to raise serious doubts about continuity of reference. While that is my view, it is by no means obvious that it need be shared universally. Realists must walk a line here between two dangers: on one side, lapsing into triviality; and on the other, making continuity of reference unattainable by too strict a requirement for shared properties, thus empowering antirealist scepticism. It is implausible to think that there could be any one, general formula that will produce univocal agreement in such cases. Judgements will depend, and rightly so, on the facts of the case, and one’s risk assessments in the face of these two dangers.

Ghins suggests that by identifying particulars with sets of properties, semirealism is bound to identify entities that are described as having different properties (by different theories) as different entities, and as a consequence, terms for such entities are on this view incapable of referential continuity across theory change. This is not quite so; indeed, I believe it is once again to conflate metaphysical and epistemological senses of ‘identity’ and ‘identification’. The identity of an entity is determined, in the metaphysical sense, by whatever properties it has (or some subset thereof), but in MFSR, I introduce an epistemological distinction between what I call the ‘detection properties’ and the ‘auxiliary properties’ of entities, and argue that only the former are relevant to considerations of reference. Scientific terms may retain reference across theory change even after significant changes in theoretical description, so long as these changes are all or primarily alterations to the auxiliary properties attributed by theories. Ghins’ realism is fastened so tightly to entities that it precludes the possibility of this sort of distinction—a distinction *between properties*—serving as the lynchpin of realist commitment. I submit, however, that precisely this sort of distinction is required in order to walk the line between triviality and scepticism.

3.3 Theory Change, Part II: Detection Properties and Auxiliary Properties

Clearly, the distinction between detection properties and auxiliary properties does important work for semirealism. It is thus crucial that the distinction be tenable, but

Psillos has doubts. The first concerns whether the distinction is in fact epistemic, as I contend, or rather something more pragmatic or perhaps even ad hoc, in either of which cases it would not serve the cause of realism. A second doubt focuses on the question of whether it is even possible to distinguish properties in this way, given that detectability admits of degrees. A final doubt targets the specific recipe I give in MFSR for determining which properties are detection properties and which are auxiliary. Let us consider these worries in turn.

The first concern is about the status of the distinction as I have presented it: does it mark an epistemic difference? Detection properties are those with which, on the basis of our current knowledge, we have managed to forge some significant causal contact—we have been able to detect and ideally to manipulate them, so that our warrant for believing in them is thereby enhanced. Auxiliary properties are all other properties attributed by theories to the entities in their domain. Psillos worries that because the line dividing detection properties from auxiliary properties changes over time as science develops, the distinction seems pragmatic, for ‘there is no epistemic mark of being auxiliary apart from the fact that there has not as yet been a causal detection of the property’. But surely the fact that a property has not been detected is an important epistemic fact. Regarding such a property one rightly says that our current knowledge is insufficient to allow detection let alone manipulation, and this should, or so I suggest, significantly weaken one’s warrant for belief. Thus, the distinction is doubly epistemic: it is fashioned on the basis of what current knowledge allows one to do; and based on what one does, warrant for belief is directly affected.

It should be immediately apparent as a consequence of this conception of the distinction between detection properties and auxiliary properties that the division is not an ad hoc device, identifying the former simply with descriptions of those properties noted in retrospect to have survived theory change, and the latter with descriptions of those properties which have not been retained. Indeed, this sort of just-so, post hoc story telling is of no help to realism. What is required is a forward-looking criterion with which to identify parts of theories that are likely to survive, and this is precisely what semirealism offers. Realists should expect—for reasons of epistemic warrant—that detection properties will survive, and that auxiliaries will survive only until such time as they are converted into detection properties or scientific theory has no further use for them.

The senses in which the distinction between detection properties and auxiliary properties is epistemic also helps to dissolve, I think, other potential confusions here. For example, Psillos wonders whether auxiliary properties are acausal (causally isolated or inert), since I claim that their ontological status cannot be determined on the basis of our causal contact with the world. But again, ‘our causal contact’ must be read epistemically in this context: on the basis of what we know, we have not managed to make causal contact with these properties, and as a consequence, their status as either existing or as fictional is not something we are in a position to determine. As scientific techniques, instrumentation, and experiments develop over time, our knowledge is extended in ways that allow us, in some cases, to detect what were previously auxiliary properties, at which point they are auxiliary no more. It is precisely this I have in mind when I say that auxiliary properties often

have a heuristic function, acting as “methodological catalysts”; the very presence of them in theories and models is a spur to scientific attempts to extend the range of our detections, thus ultimately extending the range of scientific knowledge. Consider, for instance, the failure to detect properties of the aether at the end of the nineteenth century. It was no doubt this failure that led, *inter alia*, to the eventual consensus that there is no such thing. Thankfully, sometimes results go the other way (consider the properties of neutrinos, and DNA, and so on).

None of this clarification will serve realism, however, if Psillos is right that since detectability is a matter of degree, one is never in a position to determine whether a property is a detection property or auxiliary. By “degrees of detection”, Psillos has in mind the fact that some detections are more “direct” than others, in the sense that they employ shorter causal chains of interaction leading to our senses than other, less direct detections. At the latter end of the scale, he contends, there is no clear or sharp distinction between being detected and undetected. I suspect that the “length” of so-called causal chains is a red herring in this context (no doubt some lengthy causal processes are epistemically sound and some short ones dubious), but I will not argue for this here. Let us assume more generally, as seems plausible, that some detections are more epistemically sound than others, and that at the more dubious end of the spectrum, there may be reasonable doubt as to whether a successful detection has taken place.

Would this render the distinction between detection properties and auxiliary properties untenable? It would not. It is certainly the case that we have better evidence for some detections than others. Some properties (such as, for example, certain properties of gene sequences) can be manipulated in extraordinarily precise ways so as to allow experiments in which highly novel predictions are borne out in observation. Others (such as, for example, certain properties of elementary particles detected in collider experiments) cannot be manipulated so intricately, but can be detected nonetheless. At the end of the spectrum where evidence of detection is especially weak, the realist should, like any reasonable epistemic agent, adopt reasonably apportioned degrees of belief. Credence should follow the evidence, and no doubt there will be cases in which one may wonder whether one’s degree of belief is sufficiently high to merit realism about a given property. In such cases, unlike others in which realists will naturally feel more epistemically secure, our abilities to detect and manipulate the relevant properties will be attenuated in various ways. But that is simply life in the real world, and not an objection to realism (or semirealism) more generally, or in principle. Hard cases, as they say, make bad law.

In MFSR, I offer a suggestion for identifying detection properties which I call the strategy of minimal interpretation. The idea, in a nutshell, is to interpret the variables of mathematical equations describing well-detected relations as naming detection properties, and regard any further elaborations of the natures of these properties, as given by the more general causal narratives in which such descriptions are often embedded, as auxiliary. Thus, take the term representing the intensity of a beam of light in Fresnel’s equations as naming a property—intensity—and consign any further embellishment—that it is an intensity of vibration in the aether, for example—to the auxiliary. Psillos worries that minimal interpretations are often

insufficient to pick out a given property and its relations, which may require appealing to the larger causal story in which they are embedded. Insufficient in what sense? I have no doubt that these larger causal narratives are psychologically important, for example. Perhaps Fresnel could not imagine an intensity being anything *other* than an intensity of aetherial vibration. But this by itself does not suggest that the property could not be conceived otherwise. Indeed, we know it can, because it *was* conceived otherwise after the demise of the aether. A minimal interpretation may require psychologically demanding restraint, but so be it; it is this for which we have greatest epistemic warrant.

Psillos ends his assessment of minimal interpretations with a *coup de grâce* comprising two presumptively fatal consequences. The first is that minimal interpretations of properties and relations would at best amount to knowledge of purely phenomenological laws, but it is difficult to see, however, why this should be so. The intensity of an electromagnetic disturbance, for instance, may be minimally construed quite independently of whether it is detected or even detectable by the unaided senses (and thereby phenomenological). Also puzzling is the second presumptively worrisome consequence, which takes the form of a dilemma. On the first horn, if descriptions of detection properties are interpreted independently of the larger theoretical framework in which they are embedded, what need is there to interpret theories so as to identify them? Conversely, on the second horn, if descriptions of detection properties are interpreted in a theory-dependent manner, presumably the theory has an interpretation prior to the determination of detection properties, in which case this prior interpretation will be required in order to distinguish between detection properties and auxiliary properties.

Neither horn of this dilemma is damaging, however. Regarding the first horn, perhaps detection properties *could* be attributed independently of knowing the larger theoretical frameworks inhabited by descriptions of well-detected properties and relations. This would require knowing, somehow (perhaps one is told by a reliable colleague), that the relevant mathematical descriptions apply to well-detected properties and relations without actually knowing much if anything about the theory itself. But this is not what happens in practice. In practice, one begins with some knowledge of the relevant theory, the evidence for it, its successes and failures, and so on; this brings us to the second horn of the dilemma. No doubt one's understanding of the semantics of a theory (and pertinent accompanying information, just mentioned) is a starting point. Having interpreted a theory in the semantic sense, one is now at liberty to identify the aspects of it that are most sound from an epistemic perspective. This is precisely what the minimal interpretation seeks to achieve. It aims to identify properties described by theories for which one has greatest epistemic warrant—an aim which is in no way undermined by the fact that one usually begins with an understanding of the theory *in toto*.

3.4 A Tale of Two Methods

Let me conclude with some reflections on the very nature of the attempt to explore the conceptual foundations of scientific realism. French, Ghins, and Psillos are all scientific realists, but they each attach metaphysical foundations to their conceptions

of realism that are at odds in one way or another with the position I explore in MFSR. This I applaud, not least because of the excellent criticisms which have pressed me so thoroughly in productive ways here. I am struck, however, by the conviction with which they offer alternative conceptions of the metaphysics of scientific realism as fitting (in various ways) “more naturally” with science. Our disagreements, I believe, have important meta-philosophical dimensions, and by exposing some of them in closing, I hope to offer a final defence of the project of MFSR.

This project began with the assumption that *some* metaphysical theorizing is required in defence of realism, in order to answer antirealist scepticism to the effect that the core concepts most commonly invoked in its defence—things like causation, laws, and natural kinds—are not well articulated in this context, or are incapable of coherent articulation. Having accepted the job of articulation, however, it should be clear immediately that the sciences *do not tell us how* to articulate these concepts. Indeed, the sciences underdetermine the accounts one might give. Our best contemporary scientific theories contain no explicit pronouncements in favour of any of the rival conceptual foundations one might consistently adduce in articulating realism. And thus, when French recommends that one ‘simply “read off” the [ontic structuralist] metaphysics from the theoretical context’, or when Ghins declares that scientific laws are really about particulars (not properties) in the first instance, or when Psillos identifies Humean regularities as the detection properties of science...let us see these claims for what they are. The sciences do not tell us these things. They are metaphysical proposals.

How should one view metaphysical proposals of this kind? One should view them, not as following somehow transparently from scientific investigation, but as attempts to render our conceptualizations of the fruits of scientific labour in ways that are maximally consist, coherent, and unified. The assessment of these parameters is inevitably susceptible to different outcomes based on the sorts of consistency, coherence, and unification one values, and this is not something that can be fixed by some imagined, absolute, epistemological principle. If one hopes for deeper understanding, one may value a broader range of ontological tools with which to provide it. If one despairs of such understanding, the ontology one is willing to countenance may be sparser for it. These are the sorts of predilections that marshal philosophers into Aristotelian, Humean, and other camps, and they do not stem from the sciences. Thus, when French suggests that talk of dispositions is unmotivated by the sciences, I say: unmotivated for whom? Not obviously for one who aims to give a unified account of causal necessity and laws of nature in interpreting scientific knowledge. The sciences do not tell us whether to live in the desert or in the jungle.

This, I believe, helps to expose the error in maintaining, as Psillos does, that there is a tension between the epistemic prescription I call semirealism, and the articulation of its core concepts in what he calls neo-Aristotelian terms. Is the former not constrained by the austere epistemic criterion of empirical detection while the latter is licensed to be profligate, constrained only by the criterion of explanatory power? By the austere criterion, surely the proposals of the latter should be rejected as potential candidates for belief, and by the profligate criterion,

semirealism is surely too restrictive. This is a double standard, so the argument goes: the semirealist should either be austere throughout and adopt a Humean metaphysic (as the “detection properties of science”), or profligate throughout and thus more liberal in describing the commitments of scientific realism.

There are, I think, several confusions here. The first is the idea that a Humean metaphysic is that which is given to us by science—a default setting, as it were. But as I have already suggested, this can only be wishful thinking, because scientists do not detect Humean regularities. They detect instances of what are inferred to be regularities, which are then interpretable in a number of ways, metaphysically speaking. Another confusion, to pick up the analogy to detection properties and auxiliary properties, is to think of the Humean picture as one that can be embedded within a metaphysically richer, neo-Aristotelian one, and thereby recommend (as semirealism does) that one withhold belief from what is auxiliary—the non-Humean excess. The Humean picture, however, is not embeddable within the neo-Aristotelian picture. These are, in fact, fundamentally opposed proposals for conceptualizing the world. Another confusion pertains to the dichotomization of detection and explanation as criteria for epistemic assessment. It is certainly true that the denizens of metaphysics (regularities, powers, necessities) are generally undetectable (though when metaphysics is done well, it takes our best empirical evidence as a starting point for inference). But there are many explanatory considerations that enter into a judgement that a successful detection has occurred, not least in the case of unobservable entities. Every such judgement may be regarded as an inference to the best explanation of the outputs of scientific instruments and experiments.

One must use the right tool for the right job. Where empirical evidence is available, it rightly bears strongly on what one should believe, and semirealism takes this instruction seriously. Indeed, as I argue in *MFSR*, by not taking it seriously enough, more forgiving, less austere versions of realism thereby fall prey to antirealist arguments. Where detection is impossible in principle, as in more general metaphysical theorizing, explanatory considerations are all one has. But one should not regard the different epistemic criteria one may adopt in these cases as constituting a double standard, for the contexts are epistemically dissimilar. Particularly in the latter case, I see no reason to fear the consequence that those with different explanatory values and inclinations may favour different approaches.

I expect that some may recoil from this as relativism, but when it comes to elaborating the core concepts underpinning scientific realism, I believe that one must and should admit a degree of voluntarism. The challenge for the realist, in the face of no small amount of scepticism, is to show that the relevant concepts can be given *any* sort of consistent, coherent, and unified explication. I believe I have done so, and invite others to do the same in accordance with other possible explanatory priorities, and thus make semirealism their own.

Acknowledgments I am much obliged to Curtis Forbes and Greg Lusk for discussing the content of this paper with me. It is an honour (and not a little bit daunting) to have the perceptive scrutiny of Steven French, Michel Ghins, and Stathis Psillos focused on one’s ideas with such rigour and generosity. I could not be more grateful to them, and to Gabriele Contessa whose idea this was and who skilfully managed each stage of the process. For their continuing inspiration and encouragement, my sincere thanks.

References

- Chakravartty, A. (2007). *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge: Cambridge University Press.
- French, S. (2013). Semi-realism, sociability and structure. *Erkenntnis*. doi:[10.1007/s10670-012-9417-0](https://doi.org/10.1007/s10670-012-9417-0).
- Ghins, M. (2013). Semirealism, concrete structures and theory change. *Erkenntnis*. doi:[10.1007/s10670-012-9416-1](https://doi.org/10.1007/s10670-012-9416-1).
- Psillos, S. (2013). Semirealism or Neo-Aristotelianism? *Erkenntnis*. doi:[10.1007/s10670-012-9418-z](https://doi.org/10.1007/s10670-012-9418-z).