Synthese (2017) 194:2273–2289
DOI 10.1007/s11229-015-0913-z

S.I. : CAUSATION AND STRUCTURALISM

Particles, causation, and the metaphysics of structure

Anjan Chakravartty[1]

Received: 27 May 2014 / Accepted: 11 September 2015 / Published online: 29 September 2015

Abstract I consider the idea of a *structure* of fundamental physical particles (as described, for example, in quantum theory) being *causal*. Causation is traditionally thought of as involving relations between entities—objects or events—that cause and are affected. On structuralist interpretations, however, it is unclear whether or how precisely fundamental particles can be causally efficacious. On some interpretations, only relations (as opposed to entities) exist; on others, particles are ontologically dependent on their relations in ways that problematize the traditional picture. I argue that thinking about causal efficacy in this context generates an inevitable pattern of reasoning. To assess the cogency of a given structuralist proposal one must take a stand with respect to a significant metaphysical challenge. Two options then emerge: skepticism about the form of structuralism at issue; or a dissolution of the challenge by means of a contentious ontological primitive. I contend that the choice between these options cannot be forced on scientific or philosophical grounds alone.

Keywords Subatomic particles · Causation · Structuralism · Structural realism · Scientific realism

1 Causation and fundamental physics

The central question of this paper is posed at the confluence of two great debates in metaphysics and the philosophy of science. The first of these debates concerns causation: whether it is, in fact, part of the fabric of the world, as opposed to a concept

✉ Anjan Chakravartty
chakravartty.1@nd.edu

[1] Department of Philosophy, University of Notre Dame, 100 Malloy Hall, Notre Dame, IN 46556, USA

without a referent. The second debate concerns subatomic ontology: whether we have warrant for belief in the system of particles described by modern physics (the Standard Model) as opposed to agnosticism or some more severely skeptical attitude. The central question of concern here is addressed to the prospects of marrying realism in the former debate with realism in the latter. Regarding causation, I will proceed at a certain level of abstraction; by 'causation' I will intend any number of views that are realist as opposed to deflationary, in a sense to be specified (more on this momentarily). Regarding subatomic ontology I will have something more specific in mind; my interest here is in so-called "structuralist" approaches to understanding the nature of the world at the level of quantum theory. And so, the central question of this paper comes to the following: is structuralism compatible with causation?

Before delving into the question of whether it is possible to find a home for causation in the context of structuralist views of particles, it will help to have a better sense of what is meant here by 'causation' and 'structuralism'. Let us begin with causation, which is the easier of the two if only in the sense that I will have little to say about it here.[1] My present interest is not to defend any one specific view of causation, but rather to invoke and assume a highly generic realism concerning it for the sake of argument. For present purposes, the precise content of that realism may remain unspecified—it will not matter whether one thinks of causation as a *sui generis* relation between certain objects or events, or as a special sort of process in which objects transmit certain quantities along spatiotemporal trajectories and occasionally transfer them to others, or as a form of counterfactual dependence between objects or the states of objects or events involving them, or as the manifestations of dispositions. All of these accounts of causation have at least one thing in common, which will suffice for present purposes. Each of them—let us call them "traditional" views—describes or analyzes causation in terms of the relations of, or properties of, or relations between the states or properties of *objects* or *events* involving them. On the traditional views, objects and events are the locus of what, for lack of a better term, one might call the "oomph" of causation.

The terms 'object' and 'event' are helpfully vague. For the time being it will be sufficient to proceed, I hope, with the thought that the term 'object' is intended to label the referents of both count nouns as well as things that are often only awkwardly described as "objects" per se, such as the referents of mass nouns. Thus, a quantity of liquid will count here as an object in just the way an atom does. Both are putatively things that are, according to the traditional views I have gestured towards above, causally efficacious, in the more specific ways described on each view. The term 'event' is likewise a term of art, and there is some dispute as to whether it is best used only in the context of changes to objects, or whether its use should be extended to include temporal durations during which objects do not change as such. My use of the term here will not depend on this sort of detail and, to extend this rather tolerant linguistic practice further, let me refer to objects and events thus broadly conceived as *entities*. Entities are central to traditional accounts of causation as the repositories of the causal modality and, as we shall see, it is precisely their role in this regard that

[1] For a discussion and defense of some reasoned preferences regarding causation, see Chakravartty (2007, Chap. 4).

raises questions about the compatibility of causal notions with structuralist accounts of subatomic ontology.

Why begin with the thought that there *may be any* such compatibility? This question invites some clarification of the most significant assumptions of the discussion to follow. I will assume that there *is* such a thing as causation in the world—the thing or things that many have attempted to describe in terms of the traditional accounts mentioned above, or combinations thereof. This is by no means uncontroversial, but since the aim of the current exercise is to examine the compatibility of causal and structuralist notions, it seems a fair place to begin. Even supposing a generic commitment to causation, however, one might reasonably wonder whether it is appropriate to contemplate the compatibility of this concept with any sort of philosophical interpretation of *physics*. After all, there is no shortage of philosophical opinion to the effect that causal notions have no home in the context of modern physics *simpliciter*, let alone in considerations of fundamental physics, where it may be especially tempting to think of much of the content of theories as exhausted by mathematical equations, which are then interpreted as describing nothing beyond the temporal evolutions of certain systems, and certainly nothing that might be regarded as causation. The time-reversal invariance of such descriptions (in contrast to the putative asymmetry of causation) and the challenge of interpreting correlations between distant measurement events in cases of quantum entanglement causally are further grist to this mill.

There is much to say about causal skepticism in the interpretation of physics and, especially, in various specific contexts of fundamental physics, but I will not engage with these disputes here.[2] Instead, let me begin with the assumption that causation is something that can and should have a place in our best scientific conceptions of the world, including those derived from physics, including fundamental physics. One motivation for this assumption in the present context comes from philosophers who advocate structuralist understandings of fundamental particles themselves, for throughout the development of these approaches over the past couple of decades, several of these advocates have explicitly maintained that the relevant structures are properly construed as causal (sometimes they are described as 'modal').

For example, after describing the sorts of structures that he is interested in as those comprising certain 'webs of relations' (details to follow here in Sect. 3), French (2010) adds that 'this "web of relations" is inherently modal and, in particular, causal' (pp. 92–93); 'the quantum structure that we are concerned with…can be considered to be 'causally empowered' (p. 98; cf. French 2006 and his description of causation as 'functional dependence' in French 2014, Chap. 8.8). In a similar spirit, Esfeld (2009, p. 180) holds that 'fundamental physical structures are causal structures' and that his own variation of 'ontic structural realism, as a position in the metaphysics of science that is a form of scientific realism, is committed to causal structures'. The structuralist approach to fundamental physics is *supposed* to be compatible with causation, according to some of its prominent proponents. The issue before us then is to investigate whether and how this can be so.

[2] For some rebuttals of causal skepticism in this sphere, see Frisch (2012, 2014).

Prima facie worries regarding the compatibility of causation and structuralism stem from an apparent conflict between the traditional picture's association of causal oomph with entities, and a downgrading of the ontological status of entities on structuralist descriptions of the natures of subatomic particles. In Sect. 2, I foreshadow the argument to come by reflecting on the difficulties inherent in furnishing a metaphysics of particles; such difficulties constitute the primary motivation for structuralist approaches to quantum ontology. Section 3 considers these approaches in more detail, dividing them into two categories: 'eliminative' and 'non-eliminative' structuralism. In each case I attempt to answer the question of what might serve as the locus of causal efficacy, illuminating divergences from the traditional picture of causation. In Sect. 4, I take stock of certain metaphysical challenges that result, arguing that each form of structuralism presents the hope of causal compatibilism with a dilemma in responding to these challenges: in each case one is either driven to skepticism or one must accept a contentious ontological primitive. Grasping the second horn of the dilemma saves compatibilism, but at a cost, and neither physics nor philosophy can pay the freight. In conclusion, in Sect. 5, I consider the morals of this argument for structural realism and realism more broadly in the quantum domain.

2 The metaphysics and epistemology of particles

Let me begin with a note concerning what I have been loosely calling fundamental physics. My present interest is in thinking about structuralist accounts of the nature of subatomic particles, and it is reasonable to ask whether this focus is "fundamental" in quite the right way. What if, for example, what we loosely call particles are in fact modes of excitation of an underlying quantum field? The ontology of quantum field theory, itself contested, is beyond our scope here, but perhaps excitation talk will ultimately prove preferable to particle talk. That eventuality, however, would not obviously undermine the present investigation, for an excitation of a field is, arguably, an *event*, and events are perfectly good entities in the sense canvassed in Sect. 1. Indeed, while it is only natural for some to think of a particle as an object, with misleading analogues from everyday experience and classical physics lurking in the background, given the peculiar metaphysics and epistemology of subatomic particles as conceived in connection with our best descriptions in quantum theory, it is far from clear, absent some mere but nonetheless significant terminological stipulation, whether object or event or some combination of the two is the most applicable ontological category. Leaving that more precise taxonomic challenge to one side, all that is required for present purposes is that particles are entities.

This of course leaves open the possibility of even more fundamental theorizing at the level of research programs in quantum gravity, and such further theorizing will not be my concern here. The very existence of these research programs serves to raise two different sorts of worries, however, about the investment of philosophical energy into attempts to understand the ontological natures of subatomic particles. The first worry stems from the possibility of a reductive skepticism regarding the reality of composite or "higher level" entities very generally—members of classificatory categories at levels of description "above" our most fundamental theorizing—but I

will simply assume here, along with many and certainly most of those interested in the special sciences, that proteins and giraffes are no less real than whatever it is that our most fundamental theorizing may ultimately reveal. This leaves room, at least, for the reality of subatomic particles. A second worry targets the more specific philosophical preoccupation with the *particular* categories inhabiting the quantum zoo as we know it, owing to the familiar strains that our theoretical descriptions of them place on our ability to make them intelligible *qua* ontology, in any sense other than merely thinking of them as the extensions of certain mathematical representations.

There are many areas of the sciences where classificatory practices are fruitful subjects of philosophical attention focusing on the metaphysical natures of the members of the relevant categories. Here I have in mind theorizing about the precise ontological category or categories to which an entity may belong, where 'ontological category' refers not to the name of a *scientific* category (like 'protein' or 'giraffe'), but rather the finer-grained ontological *types* or *character* of its members. Fundamental physics is a pressing example: having determined the state-independent properties (mass, charge, spin) in terms of which particles are classified, one might reasonably feel pressed to respond in some way to finer-grained questions concerning the more precise ontological status of the things that ostensibly have these properties. The fact that the precise ontological natures of the particles described in the Standard Model are notoriously difficult to understand merely amplifies this pressure. In some contexts, including measurement, they appear to behave like familiar entities with well defined properties; in other contexts, also including measurement, they do not. There is significant controversy over the question of whether particles can be individuated or distinguished from one another on the basis of their qualitative properties, or their relations, or otherwise. What is a realist to say about this?

Elsewhere I have argued that one answer to the question of how to think about the reality of particles—suspending belief in them, pending further developments in fundamental physics—is entirely self-undermining for the realist.[3] For subatomic particles seem to satisfy almost all of the criteria that are typically held to indicate the viability of realist commitment elsewhere. On a realist interpretation of the data, the relevant entities are detectable; scientists have measured their properties with impressive accuracy and precision; they can be causally manipulated in intricate ways; they can be used to interfere with other entities; theorizing about them has generated novel predictions which have been confirmed in numerous experiments. On pain of renouncing the epistemic warrant that is supposed to be conferred by such indicators, which would then undermine realism very broadly, the case for realism here seems inescapable. And yet, the metaphysical perplexities remain, and the doubt they sow is understandably real.

In the following section I will codify the various attempts that have emerged to cope with some of these perplexities by means of some form of structuralism and, for each form, attempt to locate the possible ontological home or homes of causal efficacy. In the section after that I will show how each of these approaches falls inevitably

[3] See Chakravartty (2011), which argues for a combination of realism and pragmatism regarding answers to questions about the finer-grained ontological status of certain entities. Sections 4 and 5 of the current paper can be viewed as an attempt to exemplify this combination of views in a case study.

into further perplexity in locating causal efficacy where it does—puzzlement that can be relieved only by departing from traditional understandings of causation in the direction of new and arguably strange territory. I will then go on to suggest that one plausible response to this situation for the realist is to suspend belief, not in the reality of subatomic particles, but rather with respect to their precise metaphysical natures, taking a pragmatic attitude towards the use of metaphysical characterizations in contexts where they are useful, in much the same way that realists are happy to use incompatible models of one and the same entity or phenomenon for predictive and other purposes in other domains, while maintaining a realist commitment to the relevant entity or phenomenon itself.

This immediately raises a deep (I think) question about whether believing in the existence of an entity while suspending belief in, and adopting a pragmatic attitude towards, its precise metaphysical character is a coherent combination of beliefs for the realist. I believe that this is indeed a coherent combination, and furthermore that it is one that many and perhaps most philosophers of science already accept, if only implicitly. For while most philosophers of science believe at least some parts of our best theories, their tolerance for metaphysics varies enormously. As one delves deeper and deeper into the fundaments of ontology, one finds that most people have a breaking point, beyond which they simply find the questions being asked and the answers being given too far removed from actual empirical results to be interesting or perhaps even worth considering at all. This is the point at which one may cross over, I suggest, from realism into pragmatism. Let us see now how this epistemology of particles plays out in the context of structuralism.

3 Structuralist interpretations: relations and relata

The past decade and a half has witnessed a surge of interest in structuralist accounts of scientific knowledge generally, but most intensively in connection with fundamental physics. Yet more specifically, as a generic thesis regarding the nature of subatomic particles, structuralism can be understood as an attempt to reorient what I will call the "standard" metaphysical picture of the relative ontological status of entities and their relations. On the standard picture, physical entities have what one might call a non-derivative ontological status vis-à-vis the physical (i.e. non-logical or mathematical) relations in which they stand; their existence is independent of any such particular relation. Conversely, physical relations have a derivative status; their existence depends on some things *of which* they are relations. To illustrate the picture with an everyday example, an arrangement of chairs around a table—a given set of concrete spatial relations between these entities—depends on the existence of the entities themselves. Without some such entities there could be no arrangement of table and chairs at all, let alone any particular arrangement. Any given arrangement, or structure, is derivative of the entities. Conversely, any particular set of items of furniture can exist in any number of arrangements and can survive the creation or destruction of any one such arrangement (as any serious home decorating show will attest).

Conversely, the structuralist proposal seeks to reverse the traditional direction of ontological dependence between relations and relata at the level of particles. Thus, to

illustrate the alternative picture in its intended domain, consider the structure of the Poincaré group. One may identify quantities that are invariant under the associated symmetry group transformations with properties of particles—mass and spin—but on a structuralist interpretation the properties and, concomitantly, the entities that have them depend on the group structure itself for their very existence. The structure of the group, in a way to be specified, provides some form of ontological grounding for the properties and particles. In this way, the structuralist interpretation turns the standard picture of ontological dependence between physical relations and relata on its head.

This example of structuralism as applied to quantum theory requires further unpacking, but the details vary significantly according to the form of structuralism at issue. Further complicating matters, it is fair to say that while a significant number of structuralist proposals have emerged, the precise ontologies and forms of dependence intended by their proponents are often somewhat opaque. Rather than attempt to unpack and clarify every proposal, my approach here will be conceptual. It is possible to codify the conceptual landscape of structuralism by noting that all such positions fall ultimately into one of two camps, which I will refer to as 'eliminative' structuralism and 'non-eliminative' structuralism. The non-eliminative camp would itself appear to admit of two variants: one which takes the existence of entities to be derivative of their relations and not vice versa, thus asserting a kind of ontological *priority* of the latter over the former; and another variant which takes the existence of entities and their relations in this domain to be mutually derivative, thus asserting a kind of ontological *parity* of relations and relata. Let us consider each of these conceptual camps and variants (qualifications to follow shortly) in turn and, in each case, consider what ontological resources they furnish for a possible explication of causation.[4]

Eliminative structuralism, to be perfectly accurate, does not reverse the direction of ontological dependence between physical relations and relata suggested by the standard metaphysical picture so much as rule out the direction suggested by the standard view *tout court*. That is, while it does suggest a dependence of relata on relations, it is not an ontological dependence *per se*, for on the eliminativist view, to be perfectly accurate, there are no relata. Eliminativism reverses the direction of dependence only in a sense: it *replaces* the notion of ontological dependence with one of conceptual dependence, for on this view the very notion of an entity—that is, conceived as part of the fabric of reality—is illusory. Nevertheless, our notions of particles are entirely acceptable *concepts*; their use is only misleading if one takes them to have physical counterparts in the world. The usefulness of these concepts can be demonstrated on several grounds, analogous, I would suggest, to the various ways in which, historically, fictionalists, instrumentalists, and logical empiricists were quick to point out the usefulness of terms for unobservable entities even while problematizing an ontological commitment to their referents on anything resembling a literal construal. The eliminative structuralist adopts a non-literal construal of particle talk.

[4] See Ladyman (2014/2007) for a comprehensive accounting of the positions on offer. I will assume for the sake of argument in what follows that eliminative and non-eliminative structuralism are, in fact, distinct options. For an argument to the effect that insofar as non-eliminative positions ascribe only extrinsic properties to particles, they inevitably collapse into eliminativism, see Chakravartty (2012).

If eliminativism entails a non-literal employment of terms like 'electron' and 'neutrino', just what are these terms elliptical for? That is a fair question but not easily answered, and the proponent of the position can only begin by pointing out that the subject-predicate form of standard theorizing in ontology makes it difficult for her to express what *she* means by these terms. A clue here can be gleaned from the original motivation for the position in the phenomenon of permutation invariance: quantum theory assigns no physical significance to permutations of particles, which is reflected in the ways that states of affairs involving them are enumerated using quantum statistics; permutations do not constitute distinct physical arrangements. Eliminative structuralism extends the idea of "no physical significance" to the particles themselves. As a consequence, any description of causation here, whose referent is not merely a concept but also some actual, physical modality, must locate that modality in a system of relations that lacks physical relata. The locus of causation will have to be the structure itself, on this view, absent the sorts of entities that the traditional picture would otherwise describe as central to causation. We will return to this idea shortly.

Turning to the idea of non-eliminative structuralism, one might again inquire into the location or source of causal efficacy. One possibility, it seems, is once again to ascribe causation wholly to the relations of particles, which is simply to say that something very much like or identical to the understanding of causation possessed by the eliminative structuralist is likewise available to the non-eliminativist. In addition to this possibility, however, non-eliminativism has another option, for unlike eliminativism, non-eliminative versions of structuralism propose ontologies of 'particles' literally construed, and this opens the door to the possibility that entities might play a role in causation after all, just as traditional views of causation suggest. The two variants of non-eliminativism differ only with respect to whether they take particles to be ontologically derivative of or dependent on their relations in an *asymmetrical* manner—that is, such that the relations are not likewise dependent on their relata—or whether they conceive of particles and their relations as *symmetrically* dependent.[5] While this difference is itself worthy of attention, the crucial observation presently is simply that non-eliminative structuralism, on both the priority view and the parity view, includes an ontology of entities that might serve in the analysis of causation.

The idea that entities could play causal roles given the ontology of non-eliminative structuralism must be understood in a very specific way, however. These versions of structuralism, while allowing that entities are genuinely part of the fabric of reality, nonetheless regard their existence in a way that makes them derivative of or dependent on their relations, and this is typically cashed out by saying that the very natures of particles—what it is that makes them what they are—is relational. Put another way, this is to say that the identity of a particle is determined by something extrinsic, such as the property of occupying a particular place in a system of relations, as opposed to any intrinsic property or properties. But this immediately courts ambiguity: is it the individual identity or the kind identity of a particle that is so determined? That is to ask, for example, whether it is the identity of an electron or a neutrino *qua* individual

[5] Esfeld and Lam (2011) suggest a twist on the second variant: that the distinction between particles and their relations should be understood as merely conceptual rather than in terms of ontological dependence. It is unclear to me that this describes a form of structuralism as such.

entity that is determined by the relevant structure, or whether it is the identity of an electron or a neutrino *qua* member of a category of particle that is so determined.

Insofar as one hopes to find a home for causation here, I submit that the answer to this question must be the latter, and that the issue of individuality, though interesting in its own right, is immaterial presently: only the kind identity of an entity is relevant to questions regarding its causal efficacy. It is solely in virtue of having the properties attributed to a kind—such as a category of entity described by subatomic physics, associated with a specific mass, charge, and spin—that particles have causal profiles, and thus play a role in causal phenomena. It is unclear whether one can even make sense of what a particle might be apart from the properties that make it the kind of particle that it is, but even if we were to imagine, for the sake of argument, that a particle *qua* individual can survive alterations to its properties (in virtue of some other, stable, intrinsic or extrinsic feature or features) that would result in its transformation into some other kind of thing, it is only the identity of the entity as regards its kind characterization that is relevant to its causal profile.[6] Individuals engage in causal relations in ways determined by the properties of the kinds to which they belong, and it is precisely these kind-determining properties that non-eliminative structuralism views as extrinsic. So, any description of causal efficacy here (which does not attribute it to systems of relations alone, a possibility shared with eliminativism) must attribute it to the extrinsic properties of entities.

4 Evaluating the options: a pattern of reasoning

Having canvassed the possibilities for thinking about the locus of causation on both eliminative and non-eliminative versions of structuralism, let us now consider the tenability of these options. I will suggest in this section that each of them generates metaphysical challenges, all of which have been discussed in some form or other in the literature.[7] These discussions, however, have been inconclusive, and I will argue that there is good reason for inconclusiveness here: it arises naturally from the fact that arguments for causal compatibilism in this context inevitably exemplify a pattern of reasoning which terminates in a dilemma, and neither horn of the dilemma can be forced merely on the basis of physical or metaphysical considerations. The result is a form of principled disagreement regarding how best to render the ontology of subatomic particles. As a prelude to elaborating this contention, let me first outline in general terms the pattern of reasoning and dilemma at issue.

[6] Thus, the lively debate surrounding the question of whether particles are relationally discernible as individuals (for an overview, see Caulton 2013) is largely irrelevant to the present discussion. Cf. Wolff (2012, p. 614) on the singlet state of two electrons: 'The question is not so much whether electrons are individuals, but whether they are reducible to the role they play.' More generally, most theorizing about causation is unconcerned with the question of whether the relata of causal relations are individuals or not.

[7] For a provenance of earlier forms taken by these arguments (offered by many including myself and Psillos 2006), Ladyman (2014/2007) is a helpful resource. The intended novelty of what follows is not so much in these arguments as in the meta-philosophical diagnosis of them presented. It is also worth noting that variations of these arguments can be formulated independently of the subject of causation, though the assignment of this issue of the journal—to consider whether structuralism is compatible with causation—does yield an especially pressing formulation of them.

Let us assume that compatibility with the idea that there is causation in the world involving particles is a conceptual constraint on any acceptable form of structuralism. Then, regarding any particular version of the position, the pattern of reasoning I have in mind takes the following form:

i. Given a version of structuralism, one determines a possible locus of causal efficacy.
ii. Evaluating the tenability of this proposal generates a metaphysical puzzle or challenge.
iii. In order to resolve the puzzle one must posit a contentious ontological primitive.
iv. Dilemma: accept the contentious ontological posit, or reject this version of structuralism.

At the end of this pattern of reasoning one is faced with a choice. Solutions to the metaphysical challenges that versions of structuralism invite are by no means obvious and, as a consequence, the choices to be made are themselves far from clear. To go with the first horn of the dilemma is to orient one's thinking in such a way that what might seem *prima facie* contentious or, some might argue, fatally undermining, is accepted as reasonable all things considered. To go with the second horn is to view the metaphysical commitment represented by the contentious posit as a *reductio* of the relevant form of structuralism. In what follows, I will suggest that this kind of choice applies to both eliminative and non-eliminative structuralism, and should thus inform our understanding of structuralism *simpliciter*. The upshot is that when it comes to causal compatibilism in this context, there are two kinds of people: those who are sanguine in the face of metaphysical puzzlement, and those who are repelled.

Let us begin with eliminative structuralism. Recall that the eliminativist happily subscribes to the use of particle locutions and concepts in practice, but only subject to a strict interpretation according to which they have no referents in the world, literally construed. Recall further that as a result, if causation is to be located anywhere on this picture, it must be found in a system of relations lacking the sorts of entities—the relata—that are traditionally held to play a central role in causal phenomena. Again, in the words of French (2010, p. 98), 'the quantum structure that we are concerned with does not exist independently of any exemplifying concrete system, it *is* the concrete system and as such it can be considered to be "causally empowered"'.[8] But then a puzzle arises immediately: how can a system of relations be concrete in the absence of the sorts of entities whose *being related* is what is standardly thought to constitute the concreteness of a system of relations in the first place? Standardly, relations in the absence of relata are regarded as abstract, not concrete; they are physically realized and thereby causally efficacious only when they take entities as their relata.

It is important to note that the metaphysical challenge of making sense of the possibility of causal compatibilism here does not depend on any elaborate or technical definition of the distinction between 'concrete' and 'abstract'. All that is required is the ordinary distinction between things that are spatiotemporal (whether described as having determinate locations or not, and without assuming a classical picture of

[8] For an approach to eliminativism focusing on category theoretic as opposed to group theoretic descriptions, see Bain (2013), and for skepticism regarding this use of category theory, see Lam and Wüthrich (2013) and Lal and Teh (2014).

things being "in" spacetime per se) and thus capable of acting or being affected—that is, participating in causation—and things that are not so capable, in virtue of not being spatiotemporal. On this commonplace understanding of concreteness and abstractness, the former is a prerequisite for causal efficaciousness. The question then facing the eliminative structuralist is that of whether a relation can be a spatiotemporal thing without itself being a relation of anything in particular. The standard answer to this question is 'no': relations *in themselves* are abstract, multiply realizable sorts of things; they are "concretized" when entities have them, otherwise not.

Conversely, the answer the eliminative structuralist must give to the question of whether relations can be spatiotemporal in the absence of relata is 'yes', and it is immediately clear that the only way to generate this positive answer is by reifying the relevant relations—let us call them "relations-in-themselves", to distinguish them from relations of entities. In making this move, the eliminativist is proposing a new ontological category, one heretofore uncountenanced on the standard metaphysical picture. Furthermore, this new ontological category must be taken as primitive. For on the standard picture, a state of affairs in which relations in the absence of relata are concrete is inexplicable, and eliminativism does not appear to furnish any further explanatory resources in terms of which to elaborate the assertion that such a state of affairs is possible, let alone how. Thus, any way you slice it, the metaphysical status of this new ontological category of spatiotemporally located relations-in-themselves would seem to be that of a primitive.

Here we come to the crux of the pattern of reasoning outlined earlier. In locating causal efficacy, eliminative structuralism invites the challenge of understanding how relations all by themselves can be causally efficacious, and in order to dissolve this challenge it adopts a new primitive: spatiotemporal relations-in-themselves. As is common when a primitive notion is invoked to promote or save a hypothesis, responses quite reasonably vary. One response is simply to accept that there is something here about which nothing further can be said; on the other hand, one might regard this as failing to provide insight into something of which it is fair to expect some significant understanding. The choice between these options splits the world into two kinds of people. Embracing the new posit saves causal compatibilism, but those who cannot bring themselves to embrace relations-in-themselves must then reject eliminative structuralism for failing the test of making causal compatibilism intelligible. Since the cost-benefit analysis of adopting or rejecting a new primitive in these circumstances may vary between otherwise rational agents, the decision does not appear to be forced on metaphysical grounds alone. Given that eliminativism is but one interpretation of the mathematical formalism of the physics of particles, the choice does not appear to be forced on scientific grounds either.

Let us turn now to the realm of non-eliminative structuralism. Recall that the non-eliminativist is at liberty to emulate the eliminativist in associating causation entirely with what I have dubbed relations-in-themselves, and we have seen where this leads. What is distinctive about non-eliminativism is that it has the further option, in virtue of its (not merely conceptual) commitment to entities, of analyzing causation in a manner that appears, *prima facie*, to fit seamlessly with traditional accounts of causation, which associate the oomph of causation in various ways with entities and their properties. The important qualification here is that in contrast to the standard metaphys-

ical picture according to which physical relations are (asymmetrically) ontologically dependent on their relata, the non-eliminativist takes entities and their properties to be dependent on their relations (whether asymmetrically or symmetrically): their kind identities are relational; they are constituted by an extrinsic system of relations. But then a worry arises immediately. Can the identity of a particle depend solely on the relations in which it stands? Typically, on the standard metaphysical picture, intrinsic properties of entities play an important role in constituting identities, and one might wonder whether purely extrinsic characterizations can suffice in making such identities intelligible.

Indeed, the distinction between the intrinsic and the extrinsic here generates a puzzle for relational identity, as follows. Much has been said about the contrast between intrinsic and extrinsic properties, but it will suffice here simply to proceed with the common understanding according to which an intrinsic property is one that is possessed by an entity independently of its relations to other things. They are thus, to use an awkward expression non-literally, "contained within" the entity in question. This is why, on the standard picture, they so often play a role in thinking about identity. Conversely, an extrinsic property is one that is possessed in virtue of an entity's relation or relations to some other thing or things. But then, on a purely extrinsic characterization, how is the entity that stands in these relations to be apprehended? On the ontology of non-eliminative structuralism, an entity can be understood only in terms of extrinsic features—that is, in terms of relations to *other* things. This generates a circularity or regress. In order to understand what the entity is, one must appeal to something else. But then one might reasonably wonder what it is that stands in these relations to some other thing or things, and the only answer available must again be given in terms of something purely extrinsic. Put another way, one might naturally wonder what the *it* is that is related to something else, in virtue of which one is able to speak of extrinsic features in the first place. To answer this question by citing extrinsic features yet again merely takes one in a circle.

Thus it appears that non-eliminative structuralism is susceptible to a puzzle concerning the idea of extrinsic identity. In order to describe an entity extrinsically, one seems to require a conception of something *with prior reference to which* something else might count as external (e.g. something intrinsic), but that is precisely what is lacking on the non-eliminativist view.[9] Given that extrinsic properties do not themselves yield an understanding of what entities are above and beyond relations-in-themselves—a worry that eliminativists may ignore, but that non-eliminativists who associate causation with entities cannot—one is faced here with a significant challenge. In answer, if the non-eliminativist is to treat entities as a locus of causation for purposes of causal compatibilism, she must then simply posit the existence of entities whose kind natures are exhausted by purely extrinsic properties. Let us call them "internally-extrinsic

[9] There is a long tradition of concerns about circularity that have a family resemblance, from worries in early analytic philosophy about the coherence of internal relations, to worries about forms of functionalism in the philosophy of mind, to worries about dispositional accounts of properties. Lam 2014 (p. 1163) suggests that some worries about circularity in the present context are 'mere prejudice', but I take the one I describe here to be substantive.

entities", to distinguish them from entities that have qualitative intrinsic properties.[10] In making this move she is proposing a new ontological category, one which departs significantly from those of the standard metaphysical picture. Since nothing can be said about the natures of these entities that does not appear to invite a goose chase of circularity or regress, they must be taken as primitive. The category is inexplicable on the standard picture and non-eliminativism provides no further resources in terms of which to elaborate the posit that such things are ontologically tenable.

Once again we see exemplified here the pattern of reasoning outlined earlier. The initial promise of non-eliminativism of associating causation with entities and their properties, as in traditional views, gives way to a puzzle concerning what these entities are, precisely. In order to dissolve the challenge inherent in this puzzle, one must adopt a new posit: internally-extrinsic entities. Just as in the case of eliminative structuralism, a primitive notion is thus invoked to promote or save causal compatibilism. One possible response to this maneuver is simply to accept that there is nothing more that can be said about such entities; another is to balk at the invocation of entities whose natures, one might contend, should be amenable to some better understanding. Once again, this choice splits the world into two kinds of people. Accepting the posit of internally-extrinsic entities may allow for causal compatibilism, but those who cannot accept the new primitive (and who are likewise suspicious of any strategy for causal compatibilism involving relations-in-themselves) must thereby reject non-eliminative structuralism. Again, it seems that otherwise rational agents may assess the virtues and vices of adopting the posit in different ways, thus ruling out a definitive metaphysical demonstration of the superiority of one option over the other. And again, given that non-eliminativism is just one among other interpretations of the relevant physics, the science itself is likewise inconclusive.

5 Escaping the dilemma: realist pragmatism

This paper began with the aim of examining the hypothesis that structuralist accounts of the ontology of subatomic particles are compatible with the idea that causation is a genuine feature of the world at this level of description. The subsequent consideration of the conceptual spaces in which different forms of structuralism might realize this hope terminates, in each case, in a dilemma, the form of which will be familiar to anyone who takes an interest in analytic metaphysics. At the end of the day it seems that a given hypothesis—in this case, causal compatibilism—is defensible, but only by way of primitive notions whose acceptability cannot be demonstrated beyond an illustration of the sorts of metaphysical work the posits might do. At this point it seems that different choices regarding trade-offs between accepting or not accepting arguably contentious posits, and rescuing or not rescuing the hypothesis at issue, are possible.[11] Certainly, it would seem to involve no *logical* error to go one way or the

[10] By 'qualitative' here I mean to distinguish these properties from primitive principles of individual identity such as haecceities or primitive this-ness. The former commonly feature in descriptions of the kind identities of entities, whereas the latter (cannot and thus) do not.

[11] This raises concern about Dorato and Morganti's (2013, p. 592) view that fine-grained ontology can be 'straightforwardly extracted from the relevant scientific description'. I share their pluralistic inclinations

other. In this sense, my description of the choice in terms of a dilemma may seem overly dramatic, for if there is nothing ultimately damning *qua* inference in the choice to grasp one horn of a dilemma as opposed to the other, the horns themselves may not seem especially pointy.

This dilemma is a serious one, however, in at least two ways. The first way is obvious in that, ideally, if one's aim is to realize the hope of causal compatibilism, one might reasonably further hope to achieve this aim without having to appeal to contentious metaphysical posits, and neither horn of the dilemma allows for this. The second way in which the dilemma is serious has much broader significance, and it is to this larger import that I turn now. Debates about the different forms of contemporary structuralism and the notion of causal compatibilism can be viewed as case studies in a much larger debate about the tenability of scientific realism, specifically as regards the relation of this epistemological worldview to metaphysical theorizing. The issue here is not merely one concerning arguments between proponents and critics of versions of structuralism *inter se*. It also concerns the larger question of whether resolving such arguments (for example, deciding definitively which of the two horns of the dilemma presented here is the right one to grasp) is a requirement for scientific realism. To put the question another way, one might wonder just how much metaphysics a scientific realist must do, if any, in order to defend the plausibility of her realism.

Earlier I mentioned that the sorts of criteria generally cited in arguments for the sorts of beliefs typically associated with scientific realism are amply satisfied by subatomic particles. If a commitment to the reality of particles and their properties all by itself qualifies as metaphysical (as it would, for instance, on the conception of metaphysics suggested by logical empiricism), then the answer to the question just posed is that at least *some* metaphysics is implicit in defending realism. But this much "metaphysics" is clearly acceptable to any scientific realist, whether or not one thinks the moniker appropriate. The metaphysical theorizing at issue here is a deeper variety: the kind that is engaged in responding to questions that go well beyond a commitment to the mere reality of certain scientific entities and a knowledge of their properties, to matters such as the kinds of ontological categories these entities inhabit and, as in the cases considered here, how they participate in causation. I submit that if one takes definitive answers to questions such as these (including the question of which way one should go in the dilemmas explored above) as a constraint on the tenability of scientific realism, one is asking for rather a lot. One is likely asking for too much. In conclusion, let me sketch an alternative path—one which allows realism to escape dilemmas of the sort explored here; a "third way".

Scientific realism is often thought to imply the view that our best scientific theories describe well, that is, with some significant success, the ontology of the world. This view is taken to follow from the primary description of realism as the position that our best theories are true or approximately true, describing things that actually exist even if undetectable by the unaided senses. There is a temptation to carry this implication too far, however. Consider the following. Molecular cell biology describes processes

Footnote 11 continued
towards such ontology, as may be appropriate across different domains, but agree with Arenhart and Krause's (2014) diagnosis that "extractions" inevitably make recourse to extra-scientific, metaphysical assumptions.

in which certain chemical reactions are catalyzed by enzymes. Presented with the evidence for these happenings, a certain kind of realist might commit to the reality of enzymes, reactions, and processes in which enzymes catalyze reactions. The further and more deeply metaphysical questions as to what a process is, precisely, and what distinguishes causal processes such as these from other possible kinds of processes, go well beyond the sorts of questions that can be answered definitively on the basis of the scientific evidence for the reality of enzymes and chemical reactions. A knowledge of the latter is detachable from a knowledge of the former. In just the same way, I suggest, a knowledge of the ultimate nature of particles is detachable from the knowledge that they exist, along with some significant knowledge of what they are like.

Some antirealists and perhaps even some realists will question the detachability of realism at "coarser levels" of description from realism concerning specific proposals for finer-grained metaphysical analysis. What are the grounds of such questioning? Surely it is not the conviction that every last question concerning the metaphysical natures of entities must be settled in order for realism about them to be viable. Arguments for scientific realism are simply not premised on settled answers to most fine-grained, metaphysical questions. They are premised on our seeming abilities to manipulate, in highly systematic ways and with amazing degrees of accuracy, the putative entities and structures in question. They are premised on the putative indispensability of specific items of knowledge in generating the most impressive predictive and explanatory successes of our best theories and models. In the course of these arguments, *ex hypothesi*, significant knowledge of the entities and structures in question is invoked precisely so as to manipulate, predict, explain, and so on, and all of this amounts to a considerable understanding of the natures of these things. These arguments are contested, but note: nowhere here is realism held hostage to settling every last metaphysical dispute that one might enter into regarding the relevant natures. Realism at one level of description is compatible with uncertainty at another.

This suggests a third way to proceed in the face of dilemmas of the sort I have articulated above. *Scientific* realism is compatible with the idea that it may be very difficult and perhaps even impossible for us, on the basis of the kinds of evidence we have, to argue definitively for the one true account of the fine-grained ontology of the world. This suggests that realism about certain kinds of things should be compatible with a more pragmatic attitude towards the answers one might give to questions concerning their finer-grained ontological status. The example of structuralism is an excellent case in point. It should hardly come as a surprise that the number of variants of structuralism has, in recent discussion, multiplied impressively. As a thesis concerning how best to interpret fundamental physics, structuralism has splintered from eliminativism into non-eliminativism and then again into what I earlier labeled the priority and parity variants of non-eliminativism, in part because of the very different intuitions that different agents bring to bear in evaluating which of the metaphysical puzzles confronting these views are most puzzling and, subsequently, which are the most innocuous or problematic.

Why not, then, take a pragmatic attitude towards these finer-grained metaphysical proposals? In coming to grips with the idea of causal efficacy, different versions of structuralism may be interpreted as jointly furnishing a plurality of fine-grained ontologies, all of which are compatible with scientific realism. One may have good

evidence to support realism concerning an aspect of the world while simultaneously lacking sufficiently telling evidence with respect to its finer-grained ontological profile. This opens up the possibility of realist pragmatism in at least two dimensions. Theorizing about finer-grained ontology allows us to fill in our conceptual pictures of the ontology of science in different ways, which then has the potential to lead to deeper understandings (or at least possible understandings) of our commitments pending further scientific investigation. The pragmatic commitment here is to the potential heuristic utility of attempts to understand the natures of things beyond the reach of our current science, as a primer for scientific conceptions of things to come. As the history of natural philosophy and even the modern sciences has illustrated on a number of occasions, sometimes metaphysics is proto-science. (The same can be and has been said about mathematics.)

There is also pragmatic virtue in the application of different conceptual pictures to different contexts of scientific work in the present. Different accounts of fine-grained ontology may well best suit different *explananda*. There is no obvious reason simply to assume, for example, that all the subject matters of fundamental physics let alone the natural sciences more generally should be described correctly in terms of any one account of fine-grained ontology. There are different *kinds* of things whose natures may turn out to be rather different, in which case it would be a mistake to analyze them all in the same way. There is no obvious reason simply to assume, for example, that if there are such things as spacetime points, they should have precisely the same fine-grained ontology as subatomic particles. There are open questions here regarding which ontological descriptions best correspond to the various categories of scientific taxonomy, and where there are open questions like this there is surely something pragmatic in investigating the options for how one might best conceive these different things.

Thus, the investigation into how structuralist views of particles might be conceived in such a way as to accommodate the notion of causation in the world is a microcosm of a much larger problematic. The project of elaborating these views inevitably runs into metaphysical challenges that divide the world into two kinds of people: those who are willing to adopt arguably surprising ontological primitives; and those who take such maneuverings as signs of degenerating research programs. With a dose of pragmatism regarding descriptions at such fine-grained levels of ontology, the realist may both preserve her scientific commitments and make something of deeper metaphysical theorizing. This conviction stands opposed to attempts by some antirealists to dismiss versions of scientific realism on the grounds that it is all so much metaphysics, and to the view shared by some realists and antirealists that unless we can resolve certain ambiguities concerning the metaphysical natures of scientific entities, we have no good reason to believe in them at all.

As suggested earlier, the notion of realist pragmatism also has the virtue, I think, of being an idea that most philosophers of science already accept, if only unconsciously or in an unarticulated sort of way. As they sink into the depths of ontology, most people have a breaking point—a point at which they find themselves unable to manifest the conviction that the exercise is capable of generating knowledge of the world. This breaking point, the location of which I admit and expect will vary from one person to the next, is the point at which I recommend trading realism for pragmatism.

Acknowledgements For comments and discussions relating to aspects of this material I am grateful to Michael Esfeld, Steven French, Johanna Wolff, and terrific audiences at Case Western, Northern Illinois, Oxford, Pretoria, California–San Diego, and Western Ontario, as well as at meetings of the Eastern APA (Atlanta) and the International Ontology Congress (San Sebastián).

References

Arenhart, J. R. B., & Krause, D. (2014). Why non-individuality? A discussion on individuality, identity, and cardinality in the quantum context. *Erkenntnis*, *79*, 1–18.

Bain, J. (2013). Category-theoretic structure and radical ontic structural realism. *Synthese*, *190*, 1621–1635.

Caulton, A. (2013). Discerning indistinguishable quantum systems. *Philosophy of Science*, *80*, 49–72.

Chakravartty, A. (2007). *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge: Cambridge University Press.

Chakravartty, A. (2011). Scientific realism and ontological relativity. *The Monist*, *94*, 157–180.

Chakravartty, A. (2012). Ontological priority: The conceptual basis of non-eliminative, ontic structural realism. In E. M. Landry & D. P. Rickles (Eds.), *Structural realism: Structure, object, and causality, western Ontario series in philosophy of science* (pp. 187–206). New York: Springer.

Dorato, M., & Morganti, M. (2013). Grades of individuality. A pluralistic view of identity in quantum mechanics and in the sciences. *Philosophical Studies*, *163*, 591–610.

Esfeld, M. (2009). The modal nature of structures in ontic structural realism. *International Studies in the Philosophy of Science*, *23*, 179–194.

Esfeld, M., & Lam, V. (2011). Ontic structural realism as a metaphysics of objects. In A. Bokulich & P. Bokulich (Eds.), *Scientific structuralism* (pp. 143–159). Dordrecht: Springer.

French, S. (2006). Structure as a weapon of the realist. *Proceedings of the Aristotelian Society*, *106*, 1–19.

French, S. (2010). The interdependence of structure, objects and dependence. *Synthese*, *175*, 89–109.

French, S. (2014). *The structure of the world: Metaphysics and representation*. Oxford: Oxford University Press.

Frisch, M. (2012). No place for causes? Causal skepticism in physics. *European Journal for Philosophy of Science*, *2*, 313–336.

Frisch, M. (2014). Physics and the human face of causation. *Topoi*, *33*, 407–419.

Ladyman, J. (2014/2007). Structural realism. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy*. Stanford: Stanford University.

Lal, R., & Teh, N. J. (2014). Categorical generalization and physical structuralism. arXiv:1404.3049.

Lam, V. (2014). Entities without intrinsic physical identity. *Erkenntnis*, *79*, 1157–1171.

Lam, V., & Wüthrich, C. (2013). No categorial support for radical ontic structural realism. arXiv:1306.2726.

Psillos, S. (2006). The structure, the whole structure and nothing but the structure? *Philosophy of Science*, *73*, 560–570.

Wolff, J. (2012). Do objects depend on structures? *British Journal for the Philosophy of Science*, *63*, 607–625.