

The Ontology of Quantum Field Theory: Structural Realism Vindicated?

David Glick

Faculty of Philosophy, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK

Abstract

In this paper I elicit a prediction from structural realism and compare it, not to a historical case, but to a contemporary scientific theory. If structural realism is correct, then we should expect physics to develop theories that fail to provide an ontology of the sort sought by traditional realists. If *structure alone* is responsible for instrumental success, we should expect surplus ontology to be eliminated. Quantum field theory (QFT) provides the framework for some of the best confirmed theories in science, but debates over its ontology are vexed. Rather than taking a stand on these matters, the structural realist can embrace QFT as an example of just the kind of theory SR should lead us to expect. Yet, it is not clear that QFT meets the structuralist's positive expectation by providing a structure for the world. In particular, the problem of unitarily inequivalent representations threatens to undermine the possibility of QFT providing a unique structure for the world. In response to this problem, I suggest that the structuralist should endorse pluralism about structure.

Keywords: Structural Realism, Structure, Quantum Field Theory, Pluralism, Unitary Inequivalence

1. Introduction

Broadly construed, structural realism (SR) is the thesis that our best scientific theories tell us only about the *structure* of the world. This is to be contrasted with traditional scientific realism, which attempts to find a complete ontology of individual objects and properties in scientific theories. Cast in this way, SR is a form of selective skepticism toward scientific theories; we should regard theories as approximately correct about the structure of the

world but have an attitude of skepticism toward everything they (apparently) attribute to the world beyond structure.¹

In this paper I will be concerned with two main questions. In section 2, I will ask what SR should lead us to expect. If SR is the correct view of scientific theories what follows? Here I will argue that SR should lead us to expect that future scientific theories (or at least theories of fundamental physics) will not provide an ontology of individual objects and properties as sought by traditional realists. If it is indeed *structure alone* which our best scientific theories are tracking, then we should expect ontological posits that go beyond structure to fall by the wayside.

In section 3, the question at issue will be whether SR is vindicated. Is there reason to think that science is developing theories that fail to provide a traditional ontology? I will argue that there is good reason to think quantum field theory (QFT) fulfills this expectation. Despite a concerted effort by philosophers of physics, an ontology of individual objects and properties—of the sort sought by traditional realists—has not been forthcoming in QFT. This should be troubling for the traditional realist, but is just what the structuralist expects.

SR makes a positive prediction as well: theories *should* provide us with a *structure* of the world. Here the case of QFT is not so easily embraced by the structural realist. In fact, QFT seems just as unable to provide us with an unequivocal structure as a traditional ontology. After outlining several responses to this problem, I suggest that SR may embrace a pluralism about models to resolve the issue.

2. What structural realism predicts

2.1. *Structure and success*

Following Worrall (1989), SR aims to be the “best of both worlds” by both accounting for theory change in science and providing an explanation for the instrumental success of individual scientific theories. Much work

¹Ontic structural realists such as James Ladyman and Steven French regard knowledge of the structure of the world as in principle complete, and hence, there is no room for skepticism on their view. Rather, the ontic structural realist takes the ontology of the world to be fundamentally structural. This means that the correct attitude toward the putative objects and properties recognized by traditional scientific realists is atheism, not skepticism.

has focused on the first component: does a commitment to only structure allow the realist to sidestep antirealist arguments such as the “pessimistic meta-induction”? But one may also wonder how SR is able to account for success in science.

Here the structural realist may adopt a strategy deployed by traditional realists.

If scientific realism is to be plausible and... in agreement with actual practice, then it must go for differentiated commitments to scientific theories, and what they entail about the world, in accordance with the evidence which supports them. (Psillos, 1999, 161)

Psillos goes on to argue that the ontological commitments of the realist are limited to those entities responsible for the novel predictive successes used to motivate realism with respect to some theory. This suggests a strategy—letting the drivers of empirical success determine ontological commitment—that may be borrowed by the structural realist with the addendum that *structure alone* is responsible for empirical success. Steven French describes this approach in his recent book:

The structuralist also focuses on the relevant success-inducing structures presented by the theory. However, instead of taking these to be the metaphysical outcome of properties and their interrelations [as the traditional realist does] ... she takes these structures themselves to be fundamental... (French, 2014, 44)

To illustrate, French uses the often discussed example of Fresnel’s wave theory of light. In this case, the structural realist argues that we should be committed to that which is responsible for the theory’s empirical successes, namely, Fresnel’s equations and the structure they encode. Of course, defending the structuralist’s claim that structures alone are responsible for empirical success may prove difficult. But my intention here is more limited: I wish to ask what follows from such a claim.

In particular, I want to highlight that for the structural realist adopting this strategy, structure is *doing all the work* in our scientific theories. Regardless of whether we follow Psillos’ emphasis on novel predictive success, the structural realist must claim that structure alone is capable of accounting

for empirical success. This is simply a result of the structuralist’s desire to make use of the “no miracles” argument for scientific realism. If structural realism is to provide a suitable explanation for success, then it had better be sufficient for that in which empirical success consists. Perhaps the claim that structure is sufficient for success is made more plausible if empirical success is limited to novel predictive success, but this needn’t be a commitment of structural realism. What *is* a commitment of structural realism is that structure alone—and not objects that may be posited over and above it—is all that is needed to account for the empirical success of those scientific theories worthy of belief.

Given this, it may appear as something of a mystery why some scientific theories seem to present us with *more* than structure. Why, for example, does classical mechanics seem to posit fundamental objects—in the form of Newtonian particles—in addition to the structures provided by the theory?²

One explanation open to the structural realist is that the traditional ontology associated with some scientific theories plays a merely heuristic role.³ Newtonian particles, for example, might not correspond to any individual objects in the world, but rather allow one to conceive of the structure of the classical world. Indeed, it might prove difficult to understand and apply a scientific theory without there being *some* basic ontology of objects and properties in the context of which its claims are cast. The structural realist can readily grant this point without being committed to the reality of such an ontology. As an analogy, consider a mathematical structuralist who recognizes that doing arithmetic is simplified by thinking in terms of numbers *qua* objects, but takes the natural-number *structure* as the proper subject matter of the field.

For my purposes, it doesn’t much matter how the structural realist accounts for apparent ontological posits that exceed structure in certain scientific theories. Regardless of the particular explanation offered, SR views the surplus ontology of a scientific theory as a *defect*, at least as far as the the-

²There are several candidate structures available in classical mechanics. One natural choice for the structural realist is to focus on state-space structure. North (2009), for example, argues that classical physics ascribes Hamiltonian (or symplectic) structure to the world while, Curiel (2014) argues instead for a Lagrangian state-space structure. SR claims that such structures—and not any objects thought to exist over and above them—are the proper target of the realist’s ontological commitment.

³French refers to this strategy as “Poincaré’s Manoeuvre” (2014, 67).

ory’s perspicuity is concerned. According to the structural realist, the reason why our best scientific theories are successful is that they (more or less) accurately reflect the structure of the world. Any ontological commitments that extend beyond structure must be regarded as surplus content.

2.2. *Surplus content in physics*

There is a reading of the history of physics as an elimination of surplus content. Consider, for example, the development of spacetime theories. Newton posited absolute space and time, but the dynamics encoded in his laws did not require a privileged inertial frame. The empirical success enjoyed by Newton’s theory is not the result of absolute space, because it is unaltered when one switches to a conception of spacetime that does away with a privileged inertial frame (Galilean spacetime). In the next major revolution in spacetime physics, the special theory of relativity did away with absolute positions and times altogether. Obviously there is far more to the story, but this rough sketch illustrates a trend toward the elimination of aspects of theories that are surplus in the relevant sense; idle wheels that do no work tend to be abandoned in future theories. Jenann Ismael and Bas van Fraassen describe this trend as follows:

The ontologies of our most fundamental theories are not guided by physical intuition; they are not shaped by philosophical prejudices, but led, at their best, by the ideal of a kind of formal simplicity. The history of modern physics has been... ‘a long, sustained effort to shed redundant concepts’ (Ismael and van Fraassen, 2003, 391)

Of course, the elimination of surplus content from physical theories is not always without costs. Scientific theories that are more sparse in their posits may lack the resources necessary for their successful application. Returning to the case of spacetime theories, there are two ways of understanding the metaphysical change brought about by relativity. We may, as I suggested above, claim that relativity *eliminates* positions and times and all absolute kinematic quantities that depend on them. Tim Maudlin, for example, urges that “[t]o understand Relativity, we have to expunge all ideas of things having speeds, including light” (Maudlin, 2012, 68). Yet, the common “text-book” presentation of (special) relativity *does* make use of speeds and related quantities, but relativizes them to pragmatically chosen frames of reference.

If one is interested in applying relativity, then it is extremely helpful to invoke these “surplus” elements while recognizing that they are relativized to a given context or perspective.

We are now prepared to draw a prediction. Given that non-structural content—i.e., a traditional ontology of objects and properties—is surplus according to SR, we should expect that it will be eliminated from the appropriate ontological commitments of future physical theories. This does not mean, however, that such elements will not play a significant role in the *application* of the theory. What we should expect is that if non-structural content is significant, it depends on the *context of application* or *perspective of the user* in an important sense. In summary, if SR is true, we should expect physics to develop theories in which a traditional ontology of objects and properties is either eliminated outright or contextualized to suit particular applications.

Before considering whether this prediction is borne out, it is worth clarifying the claim being made in this section. Several familiar discussions of these issues are cast in terms of the elimination of surplus *structure* rather than surplus *ontology*.⁴ Indeed, the very case I mention—spacetime theories—seems to be a case of the elimination of spacetime structure in the form of a privileged inertial frame, rather than an elimination of “traditional” ontological elements. This may seem to count *against* SR insofar as the view requires that structure is preserved across changes of theory, not eliminated.

There are two important points to consider in reply to this concern. First, SR does not require *preservation* of structure across theory change, but only *continuity* at the level of structure. SR claims only that theories capture the structure of the world approximately and imperfectly, hence, it is to be expected that our understanding of the structure of the world will be altered by future scientific developments. The key here is that these modifications of structure are continuous in a way that changes in traditional ontology are not.

Second, it is not my present intension to defend SR as such, but rather to consider an interesting prediction that follows from its truth. It is a premise in my argument that that the history of physics reveals a trend toward the elimination of surplus content, non-structural or otherwise. It is only when

⁴See, for example, Redhead (2003); Ismael and van Fraassen (2003); Baker (2010); French (2011).

combined with the claim that traditional ontology is surplus that it yields the structuralist prediction. Thus, for my purposes it is enough that surplus content tends to be eliminated, and how this impacts the structuralist’s claim about structure preservation across theory change is an issue to be fully addressed elsewhere.

3. What we find

Quantum field theory (QFT) provides the framework for a great many instances of predictive success in the physics of elementary particles.⁵ Thus, QFT is just the sort of theory scientific realists claim provides a (roughly) faithful description of the world. Yet, it is not at all clear what this description is. The two most familiar ontological approaches are the particle interpretation and the field interpretation, but both face substantial challenges.

3.1. *The particle interpretation*

Given that QFT provides the framework for the Standard Model of particle physics, it is natural to take the theory to be fundamentally about particles. Moreover, a common way of representing systems in QFT—the Fock space representation—is strongly suggestive of a particle ontology. Let’s begin, then, by very briefly reviewing how quantum theory represents the states of systems.

In quantum theory, as in many contemporary scientific theories, the state of a system is represented by a mathematical object in an abstract state-space. To take the simplest case, a pure quantum state is represented by a vector $|\psi\rangle$ in a Hilbert space H .⁶ Such a state is bestowed with physical significance by considering the action of certain operators on this state vector. Physical quantities (“observables”) correspond to self-adjoint operators on that Hilbert space and the result of one of these operators acting on the state vector provides information about potential measurement results of that

⁵One notable example is quantum electrodynamics (QED), which predicts the magnetic moment of the electron with an accuracy measured in parts per billion (parts per *trillion* for free electrons) (Vogel, 2009).

⁶In addition to pure quantum states, there are *mixed* quantum states, which may be thought of as statistical combinations of pure states. Mixed states cannot be represented by a vector in Hilbert space, but can be represented by a density operator ρ on H .

physical quantity. A Fock space is a particular Hilbert space that is used to treat systems of “identical” particles (i.e., particles of the same kind).

Within the Fock space representation, certain states of the quantum field correspond to different values of a number operator (N_i). The result of a measurement of energy E or momentum p in such a state is just what one would expect if it contained N_i particles, each with energy E_i and momentum p_i . These states can be modified by the so-called “creation” and “annihilation” operators (a_i^\dagger and a_i , respectively). Unsurprisingly, the number operator is interpreted physically as the occupation number of the system (i.e., the number of these particles it contains) and the other operators are associated with the creation and annihilation of particles. This suggests that the realist about QFT should take it to describe a world of particles, or, in Paul Teller’s terminology, “quanta.”⁷ Teller advocates a view in which the state of a physical system is represented in QFT by a vector in Fock space that is the result of “adding particles” to the empty vacuum state by the action of the creation operator.

[O]ne can start from the idea that the world contains entities that can be aggregated, called quanta. One can then describe these with a Fock space and stipulate that... a collection of [creation] operators... will work to build up vectors describing exact momentum states from the vector $|0\rangle$. (Teller, 1997, 50)

But the particles (or “quanta”) of QFT must diverge from the ordinary conception of particles in several key respects. For example, we already knew from non-relativistic quantum mechanics that particles cannot be localized to a precise location, but things are worse in QFT. There are several no-go theorems that purport to show that particles in QFT cannot be localized to *any* finite region of space-time (Malament, 1996; Halvorson and Clifton, 2002). Entities that cannot be localized *at all* hardly seem to deserve to the title “particles.”

Another problem arises from the Unruh effect, which predicts that an observer accelerating uniformly when the field is in the vacuum state of the Fock space associated with unaccelerated motion would experience a thermal bath of particles known as Rindler quanta (Earman, 2011). Thus, from the

⁷Quanta are particles that can be counted or aggregated, but cannot be not labelled or numbered. See Teller (1997, chapter 2).

perspective of the accelerating reference frame, there would exist particles that do not exist according to an inertial frame. Yet, given that QFT is a relativistic theory, we should expect its fundamental ontology to be invariant under changes of reference frame. Surely what there is shouldn't depend on whether one is accelerating or not.

Finally, even if these problems can be overcome, the Fock space representation used to describe non-interacting quantum fields cannot easily be extended to deal with *interacting* quantum fields such as the quantized electromagnetic and electron fields whose interactions are described by quantum electrodynamics. In particular, Haag's theorem shows that there are no Fock space representations available for interacting quantum fields that can be regarded as equivalent to the Fock space that would be appropriate if they did not interact (Earman and Fraser, 2006).⁸ This suggests that in describing interacting quantum systems—and there are no truly “free” systems in nature—one cannot appeal to a notion of countable particles or quanta.

Obviously, there is more to be said about the role of particles in the interpretation of QFT, but clearly the particle interpretation faces serious challenges. This leads one prominent philosopher of physics to pronounce that “[t]he notion that quantum field theory (QFT) can be understood as describing systems of point particles has been all but refuted by recent work in the philosophy of physics” (Baker, 2009, 585).

3.2. *The field interpretation*

Perhaps, then, we are better off regarding QFT as a theory describing physical *fields*, that is, a collection of properties attributed to spacetime regions. But, while QFT can be understood as describing fields, they are *operator-valued* fields—an assignment of operators to regions of spacetime. As Teller (1997) emphasizes, operators in quantum theory are associated with *determinables* rather than *determinates* and a distribution of the former seems ineligible as a fundamental ontology for QFT. Consider, as a toy example, temperature. We might take the temperature in this room to be a field comprised of a property t_x associated with each location x in the room. Notice that in this case associated with each point in space is a determinate (t_x) of the same determinable (**temperature**). At most, what an operator-valued field seems to provide, however, is an attribution of determinables to

⁸More carefully, these Fock spaces are *unitarily inequivalent* in the sense defined in section 3.4.

locations, e.g., **temperature** associated with location x . It’s hard to understand what it would mean for the world to consist of a spatiotemporal distribution of determinables, but it certainly wouldn’t qualify as a physical field in the usual sense.

What’s more, Baker (2009) argues that the problems that confront the particle interpretation mentioned above carry over to the field interpretation as well. The best developed version of the field interpretation represents the state of a system by a wave-functional $\psi[\phi(x)]$, a function that maps classical field configurations $\phi(x)$ to probability amplitudes (Huggett, 2000). Importantly, the Hilbert spaces in which these wave-functionals “live” are formally equivalent to Fock space, which lends the field interpretation much of the plausibility of the particle interpretation. But unfortunately, this equivalence also means that the wave-functional approach inherits the problems of the particle interpretation outlined in the previous section. In particular, the alternative representations of the Unruh effect and the difficulties in representing interactions carry over to the field interpretation.

3.3. *The “scandal” of QFT*

These remarks merely scratch the surface of the debate over the ontology of QFT, but they suffice to make the following observation. The difficulties associated with the ontology of QFT stand in stark contrast to the success of physicists in applying it. What results is nothing less than a “scandal” for scientific realists (Ruetsche, 2011, §15.1): the very theories for which we have best reason to let guide our understanding of reality are seemingly unable to provide such guidance.

Of course, one way to respond to the scandal is to abandon the aims of realism in the context of QFT. If QFT is not in the business of describing the world, then it is no mystery why commentators have struggled to find in it a suitable ontology. Richard Healey, for example, argues that QFT does not make “descriptive claims about physical reality” but instead serves to “advise an agent on the scope and limitations of descriptive claims it may make in a given situation” (2012, 731–732) and a broadly similar position is taken by Quantum Bayesians (Fuchs, 2010; Caves et al., 2007).⁹

⁹A key difference between these two approaches is that Healey’s pragmatism, unlike Quantum Bayesianism, claims that QFT enables us to make objectively true claims about quantum states and probabilities, with the crucial qualification that neither states nor probabilities attempt to represent elements of physical reality. Nevertheless, to the extent

The structural realist has a different take on the scandal. The reason the traditional realist approach—interpreting the mathematical formalism of QFT in terms of a fundamental ontology of objects and properties—fails is not a problem with realism, but a problem with the traditional approach to it. If we abandon the search for a traditional ontology appropriate to QFT, the scandal is avoided. Indeed, for the structural realist, the failure of QFT to provide such an ontology might be seen as a vindication of the prediction made above. Physics seems to have given the structuralist just what she was looking for: a successful scientific theory free of ontological baggage.

Laura Ruetsche (2011) draws a different moral from the issues surrounding the ontological interpretation of QFT.¹⁰ Ruetsche’s *coalescence approach* argues that “the uses to which theories are put have a hand in determining the content of those theories” (Ruetsche, 2011, 288). In terms of the debate above, the upshot of this approach is that there are contexts in which a particle interpretation applies and others to which a field interpretation is better suited. The coalescence approach obviously does not sit well with the traditional understanding of scientific realism, which seeks to derive a single ontological picture from our best theories.

From the perspective of SR, however, the situation may be seen as roughly analogous to the spacetime theories mentioned above. Just as relativity theory may be seen as *relativizing* certain kinematic quantities to a pragmatically chosen reference frame, so too is our ontological interpretation of QFT relativized to the context in which it is applied. In both cases there is an underdetermination of representations; there are multiple ways to model the same situation. This underdetermination is resolved by adopting a particular representation on the basis of practical considerations. In the case of relativity, this means choosing a reference frame that is convenient for some particular application of the theory. The structural realist can claim that different representations in QFT (e.g., a given Fock space or wave-functional space representation)—and their ontological interpretations in terms of particles or fields—are similarly underdetermined. Thus, the same moral applies to these representations, namely, that the relativized elements they contain are artifacts of our representation rather than genuine features of the world.

that the realist requires the success of QFT to be grounded in a description of the world provided by the theory, neither approach is acceptable.

¹⁰Ruetsche’s target is not QFT as such, but rather quantum theories dealing with systems that have infinite degrees of freedom (QM_∞) of which QFT is an instance.

This means that even on Ruetsche’s coalescence approach, it remains the case that QFT fails to provide a single, context-independent (traditional) ontological picture, just as SR would lead us to expect.

3.4. *The problem of inequivalent representations*

Unfortunately for the structural realist, the analogy with spacetime theories can only be taken so far. Crucial to our understanding of relativity theory is the idea that content expressed in different reference frames is physically *equivalent*; different frames offer different representations of the same situation. If this analogy is to be extended to QFT, it seems the structural realist must regard the particle and field interpretations as also equivalent. That way the structural realist can be as sanguine about the adoption of a particular ontology as she is about adopting a given reference frame. But, by the usual standard of equivalence in quantum theories, the different Hilbert space representations deployed in QFT *fail* to be equivalent.

In order to understand the problem it is necessary to understand two senses in which Hilbert spaces may be properly called “representations.” Thus far, I have been referring to Hilbert spaces as *representations* in so far as they are used to model a particular physical system, for example, a free (non-interacting) electron. In order for a Hilbert space to serve as a suitable state-space for quantum theory, though, it is a further requirement that the operators defined on it must conform to certain rules. These rules take the form of relations between operators, such as the following relation between position (x) and momentum (p): $[x, p] = i\hbar$. These rules are known as the Canonical Commutation Relations (CCRs), or, in more general form, the Weyl relations. When the operators on a Hilbert space conform to these relations, we say that it is a *representation* of the CCRs. Thus, all Hilbert spaces used to represent systems in quantum theory are also representations in this latter sense.

The question now before us is when two representations of the CCRs are physically equivalent. Intuitively, a promising answer is that we must be able to transform one representation to the other via an operation analogous to a change of reference frame in relativity. The standard candidate for such an operation in quantum theory is a unitary transformation, and hence, the standard criterion for physical equivalence is *unitary equivalence*. Two representations π_1, π_2 of an algebra of operators \mathcal{A} in Hilbert spaces H_1, H_2 (respectively) are unitarily equivalent iff: $U\pi_1(A) = \pi_2(A)U$ for a

certain unitary map (a one-to-one, linear, norm-preserving transformation) $U : H_1 \rightarrow H_2$ and all $A \in \mathcal{A}$.

This criterion has the welcome result that all Hilbert space representations of the CCRs in non-relativistic quantum theory are equivalent (a result established by the Stone-von Neumann theorem), but this is no longer the case when we consider systems with infinite degrees of freedom as in QFT. This gives rise to the problem of unitarily inequivalent representations: QFT provides multiple irreducible and inequivalent Hilbert space representations.

The presence of unitarily inequivalent representations gives rise to two related problems. First, what seems to be the same physical situation can be represented in a number of unitarily inequivalent ways. For example, in the Unruh effect mentioned above, the accelerated and inertial frames correspond to unitarily inequivalent representations of the CCRs. This leads to uncertainty about which of them captures the actual state of the world. Second, inequivalent representations are deployed in different contexts. As noted above, free and interacting quantum fields require inequivalent representations, and hence, it becomes difficult to apply a single interpretation of QFT across different contexts. To summarize, unitarily inequivalent representations give rise to two sorts of worries: an underdetermination worry and a lack of common representation worry.

The problem for the structural realist is that these unitarily inequivalent representations each seem to attribute a different *structure* to the world. A unitary transformation is an isomorphism between Hilbert spaces, so it's hard to see how there could be any notion of "shared structure" without it. This means that many of the difficulties raised above for the particle and field interpretations seem to face the structural realist interpretation as well. As Ruetsche puts the problem, "the collection of significant successes enjoyed by theories of QM_∞ defy attribution to a single type of structure"¹¹ (Ruetsche, 2011, 348). Thus, we might conclude that QFT is exactly the *wrong* sort of theory from the perspective of SR; it is equivocal not only about traditional ontology, but also about the structure it ascribes to the world.

Without a compelling reply to the problem of unitarily inequivalent representation, SR loses out to antirealist and non-representationalist approaches to QFT. After all, those approaches may also predict that contemporary

¹¹ QM_∞ is the class of quantum theories dealing with systems that have infinite degrees of freedom of which QFT is an instance.

physical theories will fail to provide a traditional ontology. SR is distinguished from such views by its positive expectation that our best theories *are* capable of capturing the *structure* of the world, and hence, this claim must be defended from the challenge of unitarily inequivalent representations. In the next section I will briefly consider several general strategies the structural realist may deploy toward this end. I hope to motivate a reserved optimism about the prospects for defending SR in the context of QFT.

Before turning to that task, let's quickly take stock. We began by drawing a prediction from SR: If structure is doing all the work and traditional ontology is surplus, and if physics tends to eliminate surplus elements over time, we should expect to physics to develop successful theories which lack a traditional ontology of objects and properties. QFT, in light of the deep problems facing its fundamental ontology, may be taken to be a candidate for such a theory. In this sense, QFT may be taken to vindicate the *purely negative* prediction just mentioned. The structural realist cannot rest content with this result, however, for at least two reasons. First, other accounts share this negative prediction. The same argument could be deployed by antirealists of various stripes, who could claim that only the empirical content of scientific theories is significant, and hence, we should expect successful theories to emerge for which no single ontological interpretation is tenable. Second, to the extent that SR is a version of realism, it must be committed to a positive claim concerning what theories *can* tell us about the world. This positive expectation is that empirically successful theories will tell us about the structure of the world. Inequivalent representations are a problem because they seem to undermine the positive expectation of SR.

3.5. Replies

3.5.1. Three extant strategies

The problem of unitarily inequivalent representations for SR is the result of three key claims:

1. The structure attributed to the world by QFT is given by a Hilbert space representation.
2. There are unitarily inequivalent representations.¹²
3. Unitary equivalence is necessary for physical equivalence.

¹²It is a further requirement that each representation is *irreducible*: it does not contain a subset capable of representing the CCRs on its own.

It follows from these claims that QFT provides us with physically inequivalent structures, a conclusion which seems at odds with SR. Three general strategies of response, naturally, correspond to the denial of each of these three claims.

The first strategy rejects that Hilbert space representations are the proper sources for the physical content of QFT. Recall that each Hilbert space is a representation of the CCRs (or Weyl relations) in that the operators on it (which represent physical quantities) must conform to these relations. This means that, despite their unitary inequivalence, distinct representations do share some features. In particular, if we consider only those operators fixed by the relations, which form the *Weyl algebra*, these will be common to every representation. There is a large contingent of philosophers working on QFT who argue that such algebras are of primary importance—those who favor Algebraic QFT (see Halvorson and Müger, 2007)—which suggests that perhaps the realist should limit her commitments to the shared Weyl algebra.

If our aim is to find a structure that is shared among unitarily inequivalent representations, and we are inclined toward taking algebras of operators to be of particular importance, a natural suggestion emerges: the real structure QFT attributes to the world is given by the Weyl algebra of operators, which may be represented in different ways. Inequivalent representations, on this view, are nothing more than different depictions of the same underlying structure. Thus, the “Algebraic Imperialist” seeks to avoid the problem by limiting her ontological commitment to that which is common among representations; anything beyond the Weyl algebra of operators is regarded as devoid of physical content.

While such an approach can respond to the problem of unitarily inequivalent representations, it is seen by some as throwing the baby out with the bathwater. Laura Reutsche (2002; 2011) argues that by only recognizing the algebra of operators common to all representations, the Algebraic Imperialist lacks the resources to account for all of the empirical successes of QFT. In particular, cases such as spontaneous symmetry breaking and phase transitions in quantum statistical mechanics seem to require taking as physically significant operators associated with inequivalent representations. This worry is familiar from spacetime case mentioned above: if we eliminate all “surplus” content we may be left with too few resources for the theory to capture all of the phenomena.

David Wallace (2011) has urged that the algebraic approach to QFT, despite its popularity among philosophers, is the wrong way to think about

QFT. Instead, he argues that we should focus on the “conventional” Lagrangian version of QFT deployed by many physicists. Importantly, Wallace takes this theory to include “cutoffs” that limit its application to systems with only *finite* degrees of freedom. Recall that in non-relativistic quantum theory the problem of unitarily inequivalent representations does not arise because the Stone-von Neumann theorem ensures the equivalence of all of the relevant Hilbert space representations. This is usually taken to fail in QFT because the degrees of freedom associated with a field in spacetime are infinite, but if we limit the theory to systems of finite degrees of freedom, the theorem can be resorted. Thus, for Wallace, the problem of inequivalent representations does not arise because all representations of the CCRs *are* unitarily equivalent.

The greatest challenge facing Wallace’s approach is that Lagrangian QFT, while perhaps fine for working physicists, has features that make it unattractive as a philosopher’s guide to the world. Imposing cutoffs only seems justified if we are really dealing with finite systems, but for a field in spacetime to have a finite number of degrees of freedom seems to require that spacetime itself is both finite and discrete (Fraser, 2009). To avoid these consequences, Lagrangian QFT must be viewed as an “effective field theory” that is only applicable within a given energy scale and, thus, it can at best tell us what the world is like at that energy scale. In other words, this understanding of QFT prevents the realist from being able to say what the world is really like at the most fundamental level. Whether this is a suitable position for a realist to take toward QFT remains to be seen.

A final strategy denies that unitary equivalence is required for physical equivalence. Above, the requirement was motivated by the intuition that unitary transformations play an analogous role to reference frame changes in relativity; both implement “changes” that make no physical difference. The general term for transformations that, like these, map solutions into physically equivalent ones are *symmetry transformations*. But, as David Baker argues, “since symmetries can fail to be unitarily implementable... unitary equivalence is not necessary for physical equivalence” (2011, 146). In essence, what Baker argues is that the presence of a symmetry mapping and unitary inequivalence come apart, and hence, it is the former rather than the latter that should be taken to establish physical equivalence. The (structural) realist, then, may claim that unitarily inequivalent representations needn’t be a problem so long as there is a symmetry mapping between them. Indeed, French suggests that the structural realist can “appeal to Baker’s condition

to assert the intertranslatability of representations” in just this way (2012, 133).

Unfortunately, it’s not clear that Baker’s condition—the existence of symmetry mapping—can be put to the use French suggests. The formal definition of symmetry¹³ that Baker deploys only relates the operators in the Weyl algebra (those shared by all representations); once we include other operators (which differ between representations) all bets are off. Of course, one could limit commitment to only the Weyl algebra, but this is to endorse Algebraic Imperialism, a position whose costs have already been reviewed. If the physically significant structure provided by a representation exceeds the shared algebra of operators, then there is no reason to think that the existence of a symmetry mapping between representations establishes their physical equivalence (Baker, 2011, 147).

While each of these strategies remains a live option for the structural realist to pursue, each faces substantial challenges as well. Thus, it is worth considering whether there are other options available. In the next section I will consider a reply that attempts to *embrace* the three claims that generate the problem while maintaining SR.

3.5.2. *Pluralist SR*

This final strategy of reply takes inspiration from a remark in Don Howard’s discussion of the problem of unitarily inequivalent representations for SR.

If the structuralist is committed to the view that theory affords a unique structural representation of nature, and if the algebraic point of view is the right point of view in QFT, the structuralist has a problem. My own view is that structuralism should never have committed itself to a uniqueness claim in the first place. But then, I am not a structuralist. So I have no right to an opinion. (Howard et al., 2011, 231)

Howard may not be a structuralist, but it is worth considering whether the structural realist can drop the uniqueness claim he alludes to. Perhaps the way to save SR is to embrace the pluralism of Ruetsche’s coalesced structures approach within a structural realist framework. Ruetsche suggests

¹³A symmetry is defined as an automorphism of the Weyl (C*) algebra $T : \mathcal{A} \rightarrow \mathcal{A}$ that preserves all expectation values.

that different contexts require different, inequivalent versions of QFT. This view is unacceptable for the traditional realist in so far as it is committed to contextually-determined incompatible ontologies of the world. But is it equally incompatible with SR?¹⁴

To begin, it is worth noting that there are varieties of pluralism that *are* compatible with traditional realism. Anjan Chakravartty (2011), for example, argues that pluralism about the metaphysical nature of objects and about how dispositions of objects manifest themselves are both compatible with scientific realism. But, such pluralism won't help in this context. The central issue in the interpretation of QFT is that *there are no candidates* for the fundamental objects the theory describes, not ambiguity about their nature or the powers they manifest.

Similarly for SR, the kind of pluralism that is needed is not merely pluralism about the features of structures or how they are represented, but pluralism about which structures constitute our most fundamental description of reality. This is particularly relevant in considering a kind of pluralism that emerges in connection with the version of SR developed by Ladyman and Ross (2007). Ladyman and Ross develop an ontic version of SR where structures consist of “real patterns” (Dennett, 1991). As Dennett himself emphasizes, it is a feature of such patterns that there can be several seemingly incompatible patterns present in the same collection of data. However, this version of pluralism is not easily applied to the present case in so far as it requires an underlying level of “data” in which the plurality of patterns is to be located.¹⁵

While it may be possible for the (epistemic) structural realist to posit an underlying non-structuralist ontology which supports the presence of multiple structures, the structuralist looking to QFT as a guide to fundamental reality will be unhappy with such an approach. Fortunately, there is another way

¹⁴French (2012; 2014) also asks this question and, as I will below, answers in the negative. His reconciliation differs from mine in that he emphasizes “structural features of QFT that cut across different representations” (2014, 323). On the pluralist view I suggest, by contrast, the inequivalent structures provided by QFT serve as our most fundamental description of physical reality. Thus, the pluralist structuralist does not seek to find a larger structure in which these inequivalent structures may be embedded, but rather, is content with a plurality of irreducible and inequivalent structures.

¹⁵Ladyman and Ross deny the claim that real patterns require underlying data. For them, “it’s real patterns all the way down.” (Ladyman and Ross, 2007, 228)

for the structural realist to embrace pluralism about structure.

The key to this view is to appreciate that SR, like any plausible form of realism, takes the knowledge of the world provided by science to be *incomplete* and *approximate*. Hence, no theory should be viewed as providing a perfectly accurate picture of the world. In the case of SR, this means that the structure of theoretical models should not be taken to capture the structure of the world exactly, but only approximately and/or incompletely.¹⁶ This looseness of match allows the structural realist to regard different structures as capturing different aspects of the structure of the world. Even seemingly incompatible structures may in fact be compatible with a single fundamental reality if they are less than perfectly accurate pictures of it.

Dennett emphasized that multiple real patterns—even seemingly incompatible ones—can be present in a single data set. Similarly, the structural realist can allow for a theory that attributes multiple structures to the world by taking them to capture the world’s structure only approximately and incompletely. Just as Dennett notes that the presence of noise in real patterns doesn’t weaken their claim to reality, so too can a pluralist version of SR claim that the structures provided by a such theory reveal genuine aspects of the world’s structure in spite of their limited accuracy. Applied to the case of QFT, the pluralist may regard unitarily inequivalent representations as each describing different structures that are really present in the world, while none of them, taken individually, provides the complete, exact structure of the world.

Recall the two aspects of the problem of unitary inequivalence for SR: underdetermination and lack of common representation. An example of the first sort of worry is provided by the Unruh effect, where the Minkowski and Rindler representations offer unitarily inequivalent descriptions of the vacuum state. The pluralist structural realist will allow that *both* representations capture aspects of the structure of the situation, even though neither should be taken to be giving us *the* structure. Each representation provides us with a structure that approximates (parts of) the structure of the world, but we cannot know where the inaccuracies lie, for doing so would require unmediated access to the structure of the world.

The second worry is that (as a result of Haag’s theorem) free and in-

¹⁶One way to make precise this notion of approximating the structure of the world is the *partial isomorphism* approach developed by da Costa and French (1990, 2003).

interacting quantum fields require inequivalent representations. This is an instance where *context* is of particular importance. In the context of certain applications—say, scattering experiments¹⁷—it is useful to deploy models which describe interacting fields, while in other contexts a free quantum field model may suffice. Here the pluralist version of SR can claim that the structure of the world is exhibited in different ways by free and interacting quantum systems. It is not that the structure of the world *itself* is context-dependent, but rather that it manifests itself differently in different contexts. Thus, the pluralist can allow for context to determine the appropriate representation—and the structure it carries with it—so long as each structure matches the world only “loosely.”

At this point, one might wonder whether this pluralist version of SR still deserves the title of “realism.” After all, on this view we cannot say what the fundamental structure of the world is, but only that it is approximately and incompletely like those structures QFT deploys in its models. Yet, I contend that pluralist SR preserves a key attribute of scientific realism: accounting for empirical success. The same argument used by traditional realists—the “no miracles” argument of Putnam (1975)—can be deployed by the structural pluralist; the reason QFT is so successful is that its models capture how the world really is.¹⁸ Of course, they do so imperfectly, but, to repeat, *any* plausible realist position must make this concession in light of the history of theory change in science.

Poincaré had a view according to which science was capable of revealing only the relations between things, as expressed in mathematical equations, and not the nature of those things. The most we can hope for is that the equations of a theory match the relations between things in the world; any content beyond these relations is mere window-dressing. His view has many affinities with contemporary SR and his remarks on reconciling inconsistent

¹⁷In fact, scattering experiments typically involve the application of both free and interacting models. The “in” and “out” states are represented as states of free fields, while the fields are treated as interacting where the particles are scattered. Such a plurality of models fits well with pluralist SR, which claims that each model (imperfectly) captures aspects of the structure of the situation.

¹⁸Not everyone is convinced by this style of argument, of course. My point is simply that, to the extent that the argument is successfully deployed by traditional realism, pluralist SR can also make use of it. One important challenge for advocates of the argument is to say how approximate truth—or approximate isomorphism—is sufficient for success (Laudan, 1981, 30–32), but pluralism does not introduce any special difficulties here.

theories are particularly apt.

When a physicist finds a contradiction between two theories equally dear to him, he sometimes says: “We will not bother about that, but hold firmly the two ends of the chain, though the intermediate links are hidden from us.” This argument of an embarrassed theologian would be ridiculous if it were necessary to attribute to physical theories the sense the laity give them. In case of contradiction, one of them at least must then be regarded as false. It is no longer the same if in them be sought only what should be sought. May be they both express true relations and the contradiction is only in the images wherewith we have clothed the reality. (Poincaré, 1913, 142)

Poincaré makes two points that are relevant to pluralist SR. First, he claims that contradictory theories may be reconciled if we focus on “what should be sought”—that is, relations between things or structure. The most obvious way this could occur is if the theories agree about structure and disagree only in the way this structure is dressed-up. The pluralist claims that even when those theories (models) *disagree* about structure they may be rendered compatible by taking them to capture the world’s structure only incompletely and approximately. Second, Poincaré claims that inconsistent theories may each express true relations. As I claimed above, in order for SR to appeal to the “no miracles” argument for realism, theories (models) must share structure with the world, even if the match is imperfect. This is not inconsistent with them “expressing true relations” so long as they express those relations only approximately or only express some of them.

In contrast to pluralist SR, consider a traditional realist position that accepts genuine pluralism about fundamental ontology. Such a view is unable to account for predictive success in the same manner as pluralist SR. If our fundamental theory—in this case QFT—is equivocal on what the ontology of the world is, then it cannot be successful in virtue of describing that ontology, even approximately. For the traditional realist, it is only once we have settled what the fundamental ontology *is* that there is room for pluralism about the nature of that ontology (à la Chakravartty). Thus, the structural realist can embrace a robust pluralism that the tradition realist cannot; *structurally* inconsistent models can be taken to be approximately *isomorphic* to the world, but *ontologically* inconsistent models cannot be taken to be approximately *true* in the same way.

This is simply a synchronic application of the argument SR deploys diachronically in the context of theory change. In the latter case, SR urges that despite discontinuity at the level of traditional ontology, scientific theories exhibit significant continuity of structure. This allows the structuralist to claim that earlier theories (e.g., Fresnel’s ether theory of light) and later theories (e.g., Maxwell’s electromagnetism) are both approximately true in so far as they both capture the structure of world, more or less. This option is not obviously available to the traditional realist. If truth for the traditional realist requires getting the ontology of the world right, then, when earlier theories differ from later ones in their ontological commitments they cannot both be true, even approximately so. This is, of course, a point made convincingly by Larry Laudan (1981). Turning to the synchronic case, the same considerations seem to show that traditional realism cannot regard ontologically inconsistent theories as both true, even approximately so. And again, the structuralist may claim a different result: structurally distinct theories may each approximate the structure of the world.

I hope to have shown that pluralist SR is a viable form of realism. If that’s right, then there is another solution to the problem of unitarily inequivalent representation in QFT. Even if we grant that different Hilbert space representations ascribe different structures to the world, we may nevertheless take them each to be (approximately and incompletely) getting at the structure of the world. I don’t claim that pluralist SR is required to solve the problem—it is quite possible that one of the strategies outlined in section 3.5.1 will pan out—but it makes inequivalent representations less threatening than they appear to be at first glance. In fact, it may now seem that unitarily inequivalent representations make QFT just the theory structural realists were hoping for.

4. Conclusion

The testing of SR requires the consideration of not only historical cases, but also contemporary and future ones. I have argued that SR leads to the expectation of empirically successful theories, now or in the future, that do not posit a traditional ontology of objects and properties. The argument has two main steps. The first is to note that SR takes *structure alone* to be responsible for the empirical success of a scientific theory, and regards as surplus those aspects of theories that go beyond structure. The second is the observation of a general tendency in physics to eliminate surplus elements

that fail to contribute to the empirical success of the theories in which they occur. Together, these lead to the expectation of theories without traditional ontological commitments.

I claim that QFT is plausibly such a theory. While it has an unassailable record of empirical success, questions over its ontological commitments remain vexed. Both particle and field interpretations face major challenges, which creates a problem for traditional realism. If a theory is successful in virtue of faithfully describing the world, then a successful theory from which no such description is forthcoming is worrying. Thus, QFT seems to vindicate SR's expectation for an empirically successful theory lacking a single traditional ontology. Yet, SR also makes a *positive* demand, namely, that theories should tell us about the structure of the world. It is in the context of this expectation that the problem of unitarily inequivalent representations occurs. The presence of such representations seems to show that QFT ascribes inequivalent structures in different contexts, or sometimes even in the same context. Several strategies of reply to this problem were considered and another was added: pluralist SR. According to the pluralist, each representation is successful in virtue of getting at the structure of the world, albeit in a necessarily incomplete and approximate manner. Pluralist SR offers a promising alternative to the other strategies of reply and thereby positions the structuralist better with respect to the problem of inequivalent representations.

So is SR vindicated in its prediction? Ultimately the answer depends on the status of debates in the interpretation of QFT. If the question of its fundamental ontology is hard, but ultimately soluble, then QFT fails to meet the negative expectation of SR. It is possible, of course, that the particle or field interpretation (or some other traditional ontological interpretation) is ultimately correct, but the current state of the debate suggests otherwise. Another possibility is that QFT may be superseded by a more complete theory (e.g., one that includes a theory of quantum gravity). If this successor theory clearly does provide a traditional ontology, this would present a difficulty for SR, while not necessarily a refutation of the view. In this case, the structuralist would be surprised that physics developed a theory that reintroduces traditional ontological posits when it is unnecessary to do so, but she might attempt to dismiss them as merely heuristic devices. At any rate, it wouldn't change the fact that physics does provide, in QFT, a theory of just the sort SR leads us to expect.

The general point is that SR allows for—indeed, *recommends*—theories

whose only ontological commitments are structural, while the traditional realist must require substantially more. This lowering of the ontological bar illustrates one sense in which SR makes a concession to the antirealist. It also allows the structuralist to come closer to the attitude of those working physicists who see QFT as an exemplary, rather than defective, scientific theory. If the current trajectory of physics is maintained, it seems unwise to hope for theories capable of an ontological interpretation of the sort required by traditional realism. This suggests that scientific realists would do well to embrace SR as their best hope for realism going forward.

References

- Baker, D., 2009. Against field interpretations of quantum field theory. *The British Journal for the Philosophy of Science* 60 (3), 585–609.
- Baker, D., 2010. Symmetry and the metaphysics of physics. *Philosophy Compass* 5 (12), 1157–1166.
- Baker, D., 2011. Broken symmetry and spacetime. *Philosophy of Science* 78 (1), 128–148.
- Caves, C., Fuchs, C., Schack, R., 2007. Subjective probability and quantum certainty. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 38 (2), 255 – 274.
- Chakravartty, A., 2011. Scientific realism and ontological relativity. *Monist* 94 (2), 157.
- Curiel, E., 2014. Classical mechanics is lagrangian; it is not hamiltonian. *The British Journal for the Philosophy of Science* 65 (2), 269–321.
- da Costa, N. C., French, S., 1990. The model-theoretic approach in the philosophy of science. *Philosophy of Science* 57 (2), 248–265.
- da Costa, N. C., French, S., 2003. *Science and Partial Truth*. Oxford University Press, Oxford.
- Dennett, D. C., 1991. Real patterns. *The Journal of Philosophy* 88 (1), 27–51.
- Earman, J., 2011. The unruh effect for philosophers. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics* 42 (2), 81–97.

- Earman, J., Fraser, D., 2006. Haag's theorem and its implications for the foundations of quantum field theory. *Erkenntnis* 64 (3), 305–344.
- Fraser, D., 2009. Quantum field theory: Underdetermination, inconsistency, and idealization. *Philosophy of Science* 76 (4), 536–567.
- French, S., 2011. Metaphysical underdetermination: why worry? *Synthese* 180 (2), 205–221.
- French, S., 2012. Unitary inequivalence as a problem for structural realism. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 43 (2), 121 – 136.
- French, S., 2014. *The Structure of the World: Metaphysics and Representation*. OUP Oxford.
- Fuchs, C. A., 2010. Qbism, the perimeter of quantum bayesianism. arXiv preprint arXiv:1003.5209.
- Halvorson, H., Clifton, R., 2002. No place for particles in relativistic quantum theories? *Philosophy of Science* 69 (1), 1–28.
- Halvorson, H., Müger, M., 2007. Algebraic quantum field theory (with an appendix by michael müger). In: Elsevier (Ed.), *Handbook of the Philosophy of Physics—Part A*. Jeremy Butterfield and John Earman, Amsterdam, pp. 731–922.
- Healey, R., 2012. Quantum theory: A pragmatist approach. *The British Journal for the Philosophy of Science* 63 (4), 729–771.
- Howard, D., van Fraassen, B., Bueno, O., Castellani, E., Crosilla, L., French, S., Krause, D., 2011. The physics and metaphysics of identity and individuality. *Metascience* 20 (2), 225–251.
- Huggett, N., 2000. Philosophical foundations of quantum field theory. *The British Journal for the Philosophy of Science* 51, 617–637.
- Ismael, J., van Fraassen, B. C., 2003. Symmetry as a guide to superfluous theoretical structure. *Brading and Castellani (2003)*, 371–392.
- Ladyman, J., Ross, D., 2007. *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press, Oxford.

- Laudan, L., 1981. A confutation of convergent realism. *Philosophy of Science* 48 (1), 19–49.
- Malament, D., 1996. In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. In: *Perspectives on quantum reality*. Springer, pp. 1–10.
- Maudlin, T., 2012. *Philosophy of Physics: Space and Time*. Princeton University Press.
- North, J., 2009. The “structure” of physics. *The Journal of Philosophy* 106 (2), 57–88.
- Poincaré, H., 1913. *The Foundations of Science: Science and Hypothesis, The Value of Science, Science and Method*. Science Press.
- Psillos, S., 1999. *Scientific Realism: How Science Tracks Truth*. Routledge, New York.
- Putnam, H., 1975. *Mathematics, Matter and Method*. Cambridge University Press, Cambridge.
- Redhead, M., 2003. The interpretation of gauge symmetry. *Symmetries in Physics: philosophical reflections*, 124–139.
- Ruetsche, L., 2002. Interpreting quantum field theory. *Philosophy of Science* 69 (2), 348–378.
- Ruetsche, L., 2011. *Interpreting Quantum Theories*. OUP Oxford.
- Teller, P., 1997. *An Interpretive Introduction to Quantum Field Theory*. Princeton series in physics. Princeton University Press.
- Vogel, M., 2009. The anomalous magnetic moment of the electron. *Contemporary Physics* 50 (3), 437–452.
- Wallace, D., 2011. Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 42 (2), 116–125.
- Worrall, J., 1989. Structural realism: The best of both worlds? *Dialectica* 43 (1-2), 99–124.