STRUCTURES AND OBJECTS: A DEFENSE OF STRUCTURAL REALISM

by

David A. Glick

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements
For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2014

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by David A. Glick
entitled Structures and Objects: A Defense of Structural Realism
and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

Date: 4 June 2014

Richard Healey

Date: 4 June 2014

Terry Horgan

Date: 4 June 2014

Jenann Ismael

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.
I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

Date: 4 June 2014

Dissertation Director: Richard Healey

STATEMENT BY AUTHOR

SIGNED: David A. Glick

ACKNOWLEDGEMENTS

There are many whose help has made this dissertation possible. First and foremost, I wish to thank my advisor Richard Healey. Without his patience and encouragement this work would have never been written. There are many others I must thank as well. Jenann Ismael, Shaughan Lavine and Terry Horgan have provided me with invaluable feedback on various stages of this work. I also wish to thank John Stachel, Jonathan Schaffer, Paul Teller, Oliver Pooley, Shamik Dasgupta, and Anjan Chakravartty for stimulating conversions on the material in Chapter 5 and Christian Loew for valuable discussions on a broad range of related issues. Finally, I could not have completed this project without the constant support of my friends and family, and especially, Hannah.

DEDICATION

For Hannah.

TABLE OF CONTENTS

TABLE OF CONTENTS – *Continued*

ABSTRACT

What stance should we take toward our best scientific theories? Traditionally, there have been two answers: realism and antirealism. Structural realism is an attempt to find middle-ground between these two views. Rather than accept everything our best theories seem to say about the world, the structural realist endorses only what those theories tell us about the *structure* of the world. I argue that switching the focus to structure allows the realist to better deal with problems of theory-change, and to better make sense of contemporary physics. I go on to offer a specific version of structural realism based on an understanding of structures as networks of relations between objects that are nothing more than *places* in structures. My view allows that there are objects and relations, but reverses the usual order of dependence: objects depend on relations rather than the other way around.

CHAPTER 1

Introduction

When John Worrall revived structural realism the late 1980's he claimed that the view could break the impasse between scientific realists and antirealists. Much progress has been made since then, but challenges remain. The aim of this dissertation is to provide the outline for a defensible version of structural realism. There are three main parts to this project. First, I hope to clarify the motivation for the view originating in the history of science. Second, I take up the question of how to be a structural realist. That is, I attempt to carve out a coherent alternative to traditional scientific realism and antirealism. Third, I apply the view to modern physics by arguing that it can explain a curious feature of how quantum mechanics and general relativity model the world.

Before turning to these specific tasks, I will say something about the general orientation from which my project begins.

1.1 Naturalistic Metaphysics

My own attraction to structuralist approaches toward science follow from an engagement with naturalized metaphysics; saying what the world is like in light of our best science. Such a project requires at least two things: (1) a relation between scientific theories and the world and (2) an understanding of naturalism that makes room for metaphysics.

I take (1) to require a minimal kind of scientific realism. From the beginning of the next chapter, I make it clear that the motivation I offer for structural realism appeals only to those who endorse the inference—common to many versions of scientific realism—from the success of science to the approximate truth of scientific theories. Without such a link between theories and the world, the project of naturalistic metaphysics cannot get started. For the classical instrumentalist, for example, theories tell us nothing about the world beyond which measurement results we should expect.

I provide no arguments in this dissertation directed at the denier of the realist inference from success to truth. To borrow a term from van Fraassen (2008), scientific realism is best conceived of as a *stance*: it is the attitude or cluster of attitudes that commit us to take science to tell us what the world—including the unobservable features of it—is really like. The fact that realism provides an explanation of success provides support for this stance, but it is unlikely to sway those with an alternative stance (such as van Fraassen's own *empiricist stance*).

The project undertaken here begins with the realist/naturalist stance and argues that structural realism is the best way to achieve the aims it envisions. Structural realism allows the realist to account for the success of science without commitment to a problematic ontology. This is the essence of Worrall's claim that structural realism is the "best of both worlds" in the realism/antirealism debate.

1.2 Problems for Structural Realism

Despite the appeal of Worrall's argument for structural realism, many have found the view unpalatable. One hope is that the present dissertation will answer some of the outstanding questions and address some important concerns regarding the view.

1.2.1 What is Structure?

One common worry concerns how the notion of *structure*—so central to the view—is to be understood. After all, there are many different conceptions of structure, and some authors (including Worrall in his original article) are unclear about what is intended by the term. Perhaps the most straightforward conception of structure comes from set theory. In that context, a structure is a domain of objects and the relations defined over them. If, however, structural realism involves a commitment to set-theoretic structures, it's hard to see how the view differentiates itself from traditional realism. What more is there to be ontologically committed to other than objects and their features?[1]

In order to be a substantive thesis, structural realism must offer more than an endorsement of ordinary, set-theoretic structure. The distinguishing feature of structural realism is not a commitment to structure per se, but rather, a commitment to the *ontological priority* or *fundamentality* of structure over its constituent parts. This makes structuralism a kind of *holism* in contrast to atomism. While many views recognize the existence of structures, the structural realist is unique in regarding *whole structures* as fundamental.

Given this understanding, the novelty of structural realism does not consist in some new notion of structure, but an increased importance given to it as ordinarily conceived. I have cast this importance in terms of ontological priority, but one could alternatively understand it epistemically: viz., that our knowledge consists, in the first place, of knowledge of structures.

Even though the distinguishing feature of structural realism need not be its conception of structure, one must still have *some* understanding of structure for the view to be complete. For the most part, in what follows I will take *structures* to be set-theoretic constructions composed of a domain of objects together with the relations defined over them. There will be times when I appeal to other conceptions of structure as well. For example, I will occasionally treat structure informally as a *network of relations between things.* I also make use of category and group-theoretic conceptions of structure. I maintain that science can provide us will models that

---

[1] One might think that the restriction to *relations* differentiates structural realism from traditional scientific realism. Yet, if we allow a structure to include both polyadic and *monadic* relations, then any property seems to fit the set-theoretic framework.

contain structures in all of these senses. As such, we should be willing to attribute all of them to the physical systems we model, that is, to (parts of) the world.

1.3 Naturalism

An overarching motivation of this project is naturalism; the need to understand the world as given to us by science. Naturalism takes many forms in philosophy, so it is worth saying something about how I am thinking of it here. As a starting point consider Quine's formulation (which appears again in Chapter 4):

> [Naturalism is] the recognition that it is within science itself, and not in some prior philosophy, that reality is to be identified and described. (Quine 1981, 21)

Quine's formulation emphasizes that science (in contrast to "prior philosophy") should be taken to be the source of our metaphysical understanding of the world. Yet, one may worry whether such a version of naturalism is too strong to allow for the sort of philosophical work that I attempt in what follows. This is a difficult issue. I do, for example, appeal to theory change in the history of science to undermine the idea that our ontology should be simply identified with the superficial commitments of our best scientific theories. Isn't this to place philosophy before science, precisely what Quine's naturalism rejects?

As I read Quine's doctrine (which is not Quine's own reading), it commits the naturalist to a fairly strong kind of realism: (only) science can tell us about the nature of reality. Yet, such a view is subject to criticism by antirealist arguments from the history of science. In particular, the problem of theory change taken up in Chapter 1 is a powerful argument against realism.

Thus, while I endorse a version of naturalism close to Quine's, I also regard challenges to it as worthy of reply and, in fact, it is these challenges which lead to the limited version of realism defended here.

I also take naturalism to include a place for metaphysics. Even if one regards science as the proper source of knowledge about the world, there remains the important work of determining what exactly it is telling us about the nature of reality. It is simply not possible to "read one's metaphysics off of physics" in any straightforward sense, and hence, there is a worthwhile project of naturalistic metaphysics: saying what the world is like if our best scientific theories are true of it.

Interestingly, arguments for structural realism operate at both the meta-level and the level of first-order metaphysics. As mentioned above, the arguments in the first part of the dissertation are aimed at showing that we can have confidence about the physical structures described by our best scientific theories. That is, while the history of science undermines a commitment to the ontology of objects and properties associated with current theories, a similar argument does not succeed at the level of structure. Thus, we can be reasonably confident that scientific models capture the structure of the world.

In addition to this meta-level argument in favor of structural realism, there is an argument that operates at the level of first-order naturalistic metaphysics. The essence of this argument is that if we take contemporary physics as our metaphysical guide, we arrive at a structuralist ontology. In Chapter 5, I argue that the permutability of entities in models of quantum mechanics and general relativity suggests that the physical objects that these theories concern—points and particles—occupy their places in structure *essentially*. This provides one example of how contemporary physics recommends an ontology in which whole structures are fundamental.

Thus, there is an interesting convergence in the structuralist approach. Structural realism is often considered to be a form of selective skepticism: theory change undermines the traditional realist's commitment to the objects and properties of our best scientific theories, but this skepticism does not affect the commitment to structures. Yet, if the preceding is correct, then even a *traditional realist* should endorse a structuralist ontology in the context of contemporary physics. In this sense, structural realism is not a form of skepticism at all.

One may conclude from this that these two arguments for structural realism are in conflict, but I do not think this is the case. Instead, there is a common explanation for both the structuralist ontology of contemporary physics and structural continuity across theory change. If scientific theories can only tell us about structure, then we should expect our ontology to come to reflect this fact. In Newtonian mechanics, for example, the apparent ontological commitments of the theory include individual objects that may be taken to exist independently of the relational structures they enter into. Thus, structural realism recommends skepticism toward those objects. Below I argue that contemporary physical theories, by contrast, do not posit anything "over and above" structure. This reflects a positive development in physics. Rather than positing ontological elements unlikely to survive theory change, contemporary physics eliminates such commitments in favor of a structuralist ontology. This is what we would expect to find if it really is structure that is responsible for the success of scientific theories.

1.4 Other Structural Realisms

This dissertation is certainly not the first place where a defense and articulation of structural realism has been attempted. In this section I will briefly outline three extant versions of the view and highlight certain difficulties they face. My goal here is not to provide a refutation of these views, but rather to motivate the need for another approach to structural realism.

1.4.1 Worrall's Ramsey Sentence Realism

As mentioned above, John Worrall's 1989 paper provided much of the impetus for contemporary work in structural realism. Worrall's version of structural realism argues that theory change in the history of (physical) science shows that the only viable form of scientific realism is one that is limited to the mathematical structure of our best theories. This is a sentiment I largely share and argue for in the first part of the dissertation.

I part ways with Worrall, however, when we move beyond motivation to the development of structural realism. Some key aspects of Worrall's view that I find problematic are the following: (1) the understanding of structural realism as an epistemic thesis, (2) the emphasis on the form of mathematical equations in characterizing structure, and (3) the claim that the structure of a theory is captured by its Ramsey sentence.

The first claim takes it that the reason our best theories can only reveal the structure of the world has to do without our limitations as knowers. I argue against such an understanding in Chapter 4. On my view, if science can only reveal the structure of the world, then we shouldn't posit anything over and above structure as part of the physical world. This means that the limitation to structural knowledge is no limitation at all; all there is to the world is structure.

The second claim focuses on the (approximate) preservation of mathematical equations as the cornerstone of continuity through theory change. In Chapter 3, I make the case that it is structures themselves—not mathematical equations—that should be taken to approximate one another. In the case I discuss there, special relativity and Newtonian mechanics are continuous in that their models have the same structure at the classical limit. The general problem with Worrall's approach is that the form of an uninterpreted mathematical equation says *very* little about the world.[2]

The third claim concerns how structure is to be understood. For Worrall, "a theory's full cognitive content is captured by its Ramsey sentence" (2007, 147). The Ramsey sentence (T^*) of a theory (T) is generated by replacing all of the theoretical predicates that occur in the theory with variables and quantifying over them (in

---

[2] This is another form of Newman's problem mentioned below.

second-order logic). T and T^* will have all the same empirical consequences, but the latter avoids commitment to unobservable properties like "being an election," for example. The problem with such an approach is that it threatens to reduce structural realism to antirealism.

If the theory T is empirically adequate then it follows that its Ramsey sentence T^* is true as long as it has the same cardinality as T. Thus, if the structure of T is identified with T^*, then structural realism is a very weak realism indeed. The only claim it makes beyond those shared by antirealists is one about the cardinality of the world. This hardly seems to capture what we intend by the "structure of the world," but, as I discuss in Chapter 2, this is not the only reason to reject this conception of structure.

The scientific realist aims to explain the success of science by appealing to the approximate truth of its theories, but if the truth of a theory is identified with the truth of its Ramsey sentence, this realist explanation is undercut. What we are trying to explain is the empirical adequacy of a theory T, so clearly appealing to the truth of its Ramsey sentence T^* is of little help. After all, the only thing the truth of T^* contributes is (at best) a cardinality claim and surely this is insufficient for explaining the instrumental success of T.

Thus, while Worrall has provided important *motivation* for structural realism, his own version of the view doesn't distinguish itself from scientific antirealism. Neither view is capable of underwriting an adequate explanation of successful applications of scientific theories.

1.4.2 Ladyman and Ross: *Every Thing Must Go*

Perhaps the best known contemporary defense of structural realism comes from Ladyman and Ross's book *Every Thing Must Go* (2007). This book contains a wide-ranging and evocative discussion of structural realism, but ultimately I find the position developed there unacceptable. Here I will offer the briefest sketch possible in order to articulate the shortcomings of this important version of structural realism.

Ladyman and Ross (along with David Spurrett), begin the book by criticizing contemporary analytic metaphysics. They argue that by ignoring genuine science and basing their arguments solely on intuitions, metaphysicians are inept when it comes to their stated aim of discovering the nature of reality. The positive conclusion of this part of the book is that any legitimate metaphysical claim should respect the Principle of Naturalistic Closure (PNC), which is stated as follows:

> Any new metaphysical claim that is to be taken seriously at time t should be motivated by, and only by, the service it would perform, if true, in showing how two or more specific scientific hypotheses, at least one of which is drawn from fundamental physics, jointly explain more than the sum of what is explained by the two hypotheses taken separately. (*ibid.*, 30)

Ultimately, the PNC is too narrow of a construal of naturalistic metaphysics. Above I offered my own understanding of the project of naturalistic metaphysics that allows for a much broader range of legitimate claims. For example, claims that quantum theory shows the world to be non-local are a legitimate topic for naturalistic metaphysicians to pursue, even if such claims do not unify the sciences in the way required by the PNC.

Ladyman and Ross are sympathetic to the motivation for structural realism from the history of theory change in science, but disagree with Worrall that this should motivate *skepticism* about the real objects and properties in the world. Instead the authors advocate an ontic version of structural realism (OSR) according to which there is only structure. In articulating their version of OSR, however, Ladyman and Ross are less than precise about what the view amounts to.

At times they seem to be advocating eliminativism about objects, while at other times they seem to merely claim that objects are inseparable from the structures in which they are embedded. In Chapter 4, I argue that the former (eliminativism) is untenable while the latter is a claim I endorse. Yet, even if we understand Ladyman and Ross to be offering a non-eliminativist version of OSR, their view faces a major conceptual problem.

Central to Ladyman and Ross's metaphysics is the notion of a "real pattern" introduced by Dennett (1991). In fact, they claim that, "everything that exists is a real pattern, and that that is all there is to say about what there is" (122).

So what is a real pattern? In Dennett's original discussion, he defines a real pattern as a way of encoding some data that improves on the "bit map" description, where the latter is a complete enumeration of every bit of data (1991, 33). Yet, Ladyman and Ross wish to claim "it's real patterns all the way down" (2007, 228). The idea is that the data themselves resolve into real patterns upon closer analysis. As I argue below (Chapter 2), such an approach fails to support the structuralist idea that the fundamental ontology of the world is structural. For every structure, understood in terms of real patterns, there must be some more fundamental level of data in which the structure is found. This means that the world cannot be fundamentally structural as the structuralist wishes.

Thus, while my project shares with Ladyman and Ross's the goal of developing a metaphysical picture motivated by historical and physical arguments for structural realism, we disagree about the basic outline of this metaphysics. I regard real patterns as a useful metaphysical tool—and indeed one that may play an important role in understanding emergent ontology—but ultimately there must be more to the world than real patterns.

1.4.3 Esfeld and Lam's Moderate Structural Realism

One view that comes close to the metaphysical proposal that follows is Michael Esfeld and Vincent Lam's "Moderate Structural Realism" (Esfeld and Lam 2008). According to this view, there is a mutual dependence between objects and the relations they enter into. The result is that structures—networks of relations between objects—are just as basic as the objects that enter into them. Moreover, objects are characterized entirely by the relations they enter into and thus fail to possess any intrinsic properties.

This last claim has proved to be the most difficult aspect of the view to defend. Consider, for example, the properties of an electron. On one hand, there are properties that many interpretations of quantum theory attribute to an electron on the basis of its state ("state-dependent properties"). On the other hand, there are properties—like mass, charge and quantity of spin—that characterize electrons generally ("state-independent properties"). While the former are (arguably) susceptible to a structuralist treatment, the latter seem to resist such a treatment. That is, state-independent properties seem to provide a clear case of *intrinsic* properties possessed by elementary particles. Esfeld and Lam regard this as a serious problem and have since modified their view to allow for intrinsic properties.

Esfeld and Lam's new version of Moderate Structural Realism begins with a claim about how we are to understand the distinction between properties and relations (on one hand) and the objects that bear them (on the other). The authors adopt the view of John Heil (2005) that properties and relations are "ways of an object existing." This makes the distinction between properties and objects *conceptual* rather than *ontological*. With this new view of properties and relations in place, Esfeld and Lam update Moderate Structural Realism to claim that "at least some central ways in which the fundamental physical objects exist are relations so that these objects do not have any existence—and in particular not any identity—independently of the structure they are part of" (2011, 157). Crucially, this allows that there are objects with intrinsic properties because these objects have "ways of existing" that correspond to intrinsic properties. Yet, because properties are simply "ways of existing," these intrinsic properties do not constitute anything additional ontologically.

The new view, however, seems to lose track of the original structural realist position. The original idea is that relational structure is fundamental, ontologically, but the updated version of Moderate Structural Realism seems to do away with structure altogether in favor of objects. Ontologically speaking, the new view seems to claim that there are only objects, but these objects have "ways of existing" that are fundamentally relational. To call this view "structural realism" is to stretch the term beyond its breaking point. Yes, relational structures are required as essential components of objects' identity, but these structures are understood as aspects of the objects themselves.

The difficulties with the view are not merely terminological. In order to accrue motivation from the history of theory change in science, a version of structural realism must account for the ontological transience of physical theories. However, the present view seems unable to do so in virtue of recognizing objects with intrinsic ways of existing. Below I will argue that the history of theory change in physics recommends a view according to which there is only structure. Objects, to the extent that we can recognize them, must be located within structures. Such a view encounters the problem of intrinsic properties that led Esfeld and Lam to modify their original view, but unlike them, I am optimistic about the prospect of accounting for seemingly intrinsic properties in structuralist terms. In the case of state-independent properties, for example, there is an approach to categorizing elementary particles going back to Wigner (1939) that locates them in a relational structure. According to this approach, elementary particles are classified as irreducible unitary representations of the Poincaré group. This group-theoretic classification scheme places elementary particle *types* into a network of relations and thereby provides a path toward a structuralist understanding of state-independent properties.

In summary, Esfeld and Lam take the problem of intrinsic properties to undermine an attempt to give a purely structuralist metaphysical account of fundamental physics. This leads them to propose a new version of Moderate Structural Realism that allows for intrinsic properties understood as “ways of existing.” This, however, is to give up on structural realism in its original form. Instead, I claim we should attempt to recover apparently intrinsic properperties by appeal to higher-level structures, such as the group-theoretic structures appealed to by Wigner.

1.5 Plan

There are three parts of this dissertation. The goal of the first part is to provide motivation for structural realism generally. In Chapter 2, I take up the issue of theory change in the history of science that led Worrall to structural realism. There I argue that the history of science supports a version of realism in which entire theoretical models correspond to the world. Moreover, this “global correspondence” is best understood as models (more or less) accurately capturing the structure of the world. In Chapter 3, I look at a specific case: the transition from Newtonian mechanics to special relativity. I argue that this transition, while discontinuous at the level of (object-based) ontology and ideology, is continuous when one focuses on the structure of these theories.

The second part of the dissertation addresses the issue of how to best develop structural realism. Even if one grants that there are good reasons to consider structural realism, some challenge that the view either (1) collapses into antirealism or traditional realism, or (2) is simply incoherent. Chapter 4 is an attempt to respond to such challenges by laying out a plausible version of the view. First, I argue that structural realism should be understood as a metaphysical thesis by providing a new argument for OSR based in naturalism. Second, I argue that OSR must be understood as non-eliminative about objects. Third, I propose a metaphysics according to which there are only structures, but objects may be recovered as *places* in structures. This view, I maintain, allows the structural realist to use the motivation from parts one and three to advocate for a unique view of scientific models and the world they describe.

The third and final part of the dissertation brings to bear considerations from contemporary physics. In Chapter 5, I consider quantum mechanics and general relativity and argue that a curious feature of these theories is best explained by a structuralist metaphysics. That curious feature is what John Stachel (2002) calls "general permutability": that the basic entities in the models provided by these theories can be permuted without coming to represent a distinct state of affairs. The best explanation for this fact, I claim, is that the objects these theories describe—spacetime points and elementary particles—are to be understood as essentially elements of a relational structure. A particular spacetime point p, for example, must occupy the place it does in a geometric structure in all possible worlds in which that structure is present. The view I advocate there, Minimal Structural Essentialism, gives content to the idea that objects are individuated by structure alone.

I conclude, in Chapter 6, by explaining how all of these pieces fit together to provide a unified framework for structural realism. I also briefly touch on the scope of the project and how quantum field theory (QFT) might provide further support for this version of structural realism.

CHAPTER 2

The Realist Trap

2.1 Introduction

The goal of this chapter is to introduce structural realism as a promising version of scientific realism. Scientific realists who endorse the inference from the success of science to the approximate truth of its theories expose themselves to challenges from the history of science. In particular, there are well-known arguments associated with the work of Thomas Kuhn and Larry Laudan that threaten to undermine the realist position. I aim to show that structural realism—the view that our best theories can tell us only about the *structure* of the world—is the best escape from the trap set by these antirealist arguments.

This chapter will proceed as follows. Section 2.2 will outline the inference from the success of a theory to its truth that lies at the center of many versions of scientific realism. Next I will introduce the problem of theory change in section 2.3. There seem to be numerous cases of past theories that were successful but not true, apparently undermining the realist inference. Indeed, as Kuhn notes, there seems to be no progress in science when one focuses on the ontological commitments of our most successful scientific theories. This situation gives rise to the *realist trap* discussed in section 2.4: truth (even approximate truth) seems to require getting the ontology of the world right, but the history of science seems to undermine our faith in the objects and properties currently recognized as posits of our best theories.

After presenting the realist trap in detail and considering various ways out, I propose that the trap should be avoided by separating truth from a commitment to the ontology of the theory. In particular, in section 2.6 I propose *global correspon-*

dence—a relation that obtains between *whole* models and the world—as the best way to achieve the desired separation. Structural realism is then put forward as a specific version of global correspondence. The remainder of the chapter (section 2.7) is devoted to dealing with various problems with structural realism as a solution to the realist trap.

2.2 The Realist Inference

There is considerable disagreement over the definition of scientific realism. I don't intend to enter into this debate here, but rather, I will focus on an inference that plays a central role in many realist views. The *realist inference* is an inference from the instrumental success of a scientific theory to its (approximate) truth. This inference is at the heart of the so-called "no miracles" argument introduced by Putnam (1975a, 73)[1], but the two should not be identified. Putnam's "no miracles" argument alleges that "realism is the only philosophy of science that doesn't make the success of science a miracle" (1975a, 73). If we assume that (1) there must be a non-miraculous explanation for success and, (2) realism requires the (approximate) truth of scientific theories, then Putnam's argument is an instance of the realist inference. Yet, the realist inference is more general than Putnam's argument. All that is required for the realist inference is the claim that the success of a scientific theory provides evidence for the claim that it is approximately true.

Putnam's argument does, however, suggest a certain *motivation* for the realist inference. The reason one should take successful theories to be true is that truth *explains* success. This motivation for the inference is usually cast in terms of inference to the best explanation (IBE) as follows.

1. Theory T is successful.

---

[1] The "no miracles" argument is also known as the "ultimate" argument for scientific realism (van Fraassen 1980; Musgrave 1985).

2. The best explanation of T's success is that T is approximately true.
3. IBE [2]
4. Therefore, T is (probably) approximately true.

If it is successful, this simple argument provides justification for the realist inference, but it does not go so far as to argue that realism is the only philosophy of science capable of explaining the success of science (non-miraculously). In fact, it doesn't mention realism at all, only the approximate truth of a theory. Most scientific realists who accept the realist inference justify it by something like the argument involving IBE sketched above; success implies truth because the latter explains the former.[3]

Some regimentation is required for the realist inference to be the substantive thesis its backers take it to be. There are many conceptions of success and truth, some of which give the realist inference an unintended meaning. How, then, ought we understand *success* in the context of the realist inference? Most realists who endorse the inference take success to consist in the fact that a particular scientific theory has given rise to (1) successful predictions, (2) technological advancements and, (3) a unified understanding of the world. Alan Musgrave (1985) and others following him (e.g., Psillos (1999)) have emphasized the role of *novel* predictive successes. These authors maintain that genuine successes occur when an unexpected phenomenon is predicted by a scientific theory and then found to occur. For example, Eddington's

---

[2] There are wrinkles that confront any attempt to provide a precise definition of IBE. For our purposes we may characterize IBE, roughly, as the following claim: the theory which best explains the evidence—when compared with all other candidate explanations—is probably true.

[3] One reason for avoiding the use of IBE is the worry that the principle begs the question against the antirealist, who denies its cogency. One defender of the realist inference who explicitly rejects the use of IBE is John Worrall (2007). He claims instead that the inference is simply a defeasible statement of our default position toward scientific theories.

1919 observational confirmation of the degree of deflection of light by the Sun as predicted by general relativity is a *novel* predictive success. Richard Boyd (1990), on the other hand, emphasizes the reliability of scientific methodology at a general level rather than focusing on the success of individual scientific theories. For our purposes here, we will adopt a picture of success closer to Musgrave's than Boyd's. The success of science consists in individual theories satisfying (1)-(3) above.

The next important point of clarification concerns truth. Among the many accounts of truth there are some that must be excluded from consideration as far as the realist inference is concerned. There are various conceptions of truth that view it as a fundamentally *epistemic* notion by tying it to notions such as verification, assertability or justification. That we are (or will someday be) warranted in asserting the contents of a successful scientific theory, for example, does not capture what the realist intends by "success implies truth." In addition, deflationationary understandings of truth (e.g., Horwich (1998)) fail to provide an account of truth robust enough for the realist inference. If one thinks that the disquotational rule "p" is true iff p exhausts what there is to say about truth, then the realist inference will not have the desired force.

What is intended by the realist inference is that success merits belief in the truth of theories in the *correspondence sense*. Here "correspondence" should be taken quite broadly, namely, as the thesis that truth consists in some relation between theories and the world. As we will see in section 2.5, there are many ways of developing this idea, but all share a conception of truth as a theory-world relation. It is worth noting that one can endorse the realist inference without subscribing to a correspondence view of truth generally. The conclusion of the inference is that our best theories and the world “correspond” in the relevant sense, which one can hold even if one subscribes to a alternative account of truth in other contexts (say, ordinary everyday utterances).

There is another constraint on the notion of truth that appears in the realist inference: it must be explanatory. The most natural way to motivate the realist inference is as an inference to the best explanation. When viewed in this way our conception of truth must be robust enough to offer an explanation of the success of science. If, for example, truth is understood in strictly verificationist terms (i.e., T is true just in case it is empirically adequate), while this would count as a “correspondence” view in the weak sense above, it fails to offer an adequate explanation of success. The empirical adequacy of a theory cannot explain success as characterized above, for empirical adequacy forms part of the realist’s explanandum. Of course, what constitutes an “adequate explanation” may also be challenged and, like truth, certain accounts of explanation may undermine the spirit of the inference. Again, it is worth emphasizing that the goal here is not to defend a particular account of truth or explanation, but rather to characterize the realist inference. Once characterized, one may accept it or reject it, and either option is consistent with whatever view of truth or explanation one favors generally.

To summarize, the realist inference may be stated as follows.

> *Realist inference*: The instrumental success of our best contemporary scientific theories constitutes evidence that these theories are approximately true.

Recall that *success* is understood to include primarily novel predictive success (but may also include advances in technology and unification in our understanding of the world) and *truth* is understood in terms of a correspondence relation R holding between a theory and the world capable of supporting an explanation of theoretical success.

In the sections that follow I will refer to those who endorse the realist inference simply as "realists" and those who deny it as "antirealists." In so doing I do not intend to take a stand on what scientific realism is, though it does seem that the vast majority of self-described scientific realists would be "realists" in my sense. Entity realists like Ian Hacking (1982; 1983; 1985) and Nancy Cartwright (1983) are notable exceptions. Entity realists are generally regarded as scientific realists because of their commitment to many of the unobservable entities that occur in the context of contemporary science (e.g., electrons), but they reject any commitment to the truth of the theories which purport to deal with those entities. It is worth discussing briefly how the belief in unobservable entities—which is often taken as a hallmark of scientific realism—relates to the realist inference with which I am concerned.

There is a natural path from the realist inference to the belief in unobservable entities which perhaps explains why most realists hold both. If one takes our most successful theories to be approximately true (by application of the inference), and approximate truth requires that the entities mentioned in the theory exist, then one arrives at the commitment to unobservable entities. I will have more to say about this argument below, but for now I simply wish to note that while the realist inference may be sufficient for belief in unobservable entities, entity realism shows that it certainly isn't necessary. Again, in what follows, "realism" will simply be identified with the endorsement of the realist inference. This means that entity realists who deny the inference (like Hacking and Cartwright) will not count as realists in my sense.

Now that we have a workable notion of realism, I will turn my attention to challenges faced by the scientific realist.

2.3 The Problem of Theory Change

One of the most troubling lessons to fall out of *The Structure of Scientific Revolutions* (Kuhn 1962/1996) is that the history of science is full of radical shifts in our theoretical commitments. Theories that were once viewed with near certainty are now regarded as outright refuted. To take one example, the wave theory of light associated with Augustin-Jean Fresnel was very well established in the first half of the 19th century and the luminiferous ether it posited was thought to be an essential component of reality. We no longer think there is any luminiferous ether and hence Fresnel's theory, which describes light in terms of the properties of this ether, cannot be anything but false.

This conclusion itself isn't worrisome—it is not news that abandoned theories are false—but an extension of this line of reasoning to today's theories leads to anti-realist skepticism. We seem to be in an epistemic position not substantially different from that of earlier thinkers with respect to our scientific theories. Today's realists take their theories to be true in light of the significant successes they have achieved, but earlier realists made the same inference about Frensel's theory and were wrong. As it is often put, in light of the history of science it is unlikely that our current theories will prove final. Of course, if theory changes were continuous, we could perhaps salvage some approximation of truth for our theories, but Kuhn argues this is not the case. His best known view is that scientific revolutions amount to changes in paradigms that are radical and discontinuous. This is a controversial view, to say the least, but the following observation is less so.

> I do not doubt, for example, that Newton's mechanics improves on Aristotle's and that Einstein's improves on Newton's as instruments for puzzle-solving. But I can see in their succession no coherent direction of ontological development. On the contrary, in some important respects, though by no means all, Einstein's general theory of relativity is closer to Aristotle's than either of them is to Newton's (1962/1996, 206).

The idea that the ontology of the world—what our best theories tell us there is—has changed radically and discontinuously finds good support in the history of science. Yet, even this claim is big trouble for the realist. The realist finds evidence for the truth of a theory in its instrumental success, but the past is full of successful theories that purport to refer to things that we no longer think exist. Furthermore, there is little reason to think that the ontology of our current theories will prove final. Hence there are two problems facing the realist:

1. How is it that success indicates truth if past theories were successful but not (even approximately) true?
2. How can we be justified in thinking our current theories are (even approximately) true given that they aren't likely to be final?

It is worth noting that both problems arise even if one only grants Kuhn the observation that ontology changes discontinuously with theories. The reason past theories cannot be even approximately true is that they describe things that don't exist. (A theory of how many angels fit on the head of a pin cannot be even approximately true if there are no such things as angels!) The same holds true for our current theories. The issue is not so much that they will be succeeded by new theories, but rather that their ontology, the entities they posit, are unlikely to be carried over. Thus, the same fate befalls them: they are not even approximately true.

The challenge to realism under discussion is closely related to the well known arguments of Larry Laudan (1981; 1984a). Laudan provides several counterexamples—including Fresnel's wave theory of light—to the realist inference and argues that realism can give no account of why such theories were successful. There are two reasons why I have chosen to focus on Kuhn's observation rather than Laudan's argument. First, Laudan understands realists as being committed to the claim that the central terms of successful theories genuinely refer. This emphasis on the reference of *terms* is problematic for reasons that will emerge below. Second, Kuhn's observation provides a clear and simple statement of the challenge facing realists. The heart of Laudan's challenge is captured by questions (1) and (2) above. These questions follow straightforwardly from Kuhn's observation and hence we can avoid some of the complexities of Laudan's argument by focusing instead on Kuhn's simple observation.[4]

---

[4] Another argument in this vicinity is the so called "pessimistic induction" that argues (roughly) that most current theories are likely false because most past theories were false. Once again, this argument adds unnecessary complexity to what is needed to generate the problem for scientific

2.4 The Realist Trap

There are several ways for the realist to respond to Kuhn's challenge. Perhaps most straightforwardly, one can deny that history casts doubt on current theories because the latter bears little similarity to earlier theories. We can appeal to improved methods of testing predictions or more sophisticated mathematical tools to argue that earlier theories weren't up to today's rigorous standards. If this is the case, then the fact that they proved false and their ontologies unreal ought not affect our stance toward today's well-established science. Moreover, they were never "successful" in the relevant sense, so the realist inference is safe from counterexample as well. Yet, such a claim, while it might be plausible with respect to some earlier theories (e.g., the phlogiston theory[5]), faces difficult cases such as Fresnel's wave theory of light mentioned above. That theory was precise and mathematical in its formulation and faced stringent tests even by today's standards.[6] Nevertheless, there is no ether and hence the theory is false. This should give us pause when considering what the world is really like in light of current science.

Suppose that the realist grants that there were earlier scientific theories, not significantly dissimilar from current ones, that were successful but are now abandoned.

---

realists. For example, Magnus and Callender (2004) argue that the pessimistic induction commits the base-rate fallacy by making an unwarranted statistical claim about the population of current theories on the basis of evidence about the population of past theories.

[5] Hardin and Rosenberg (1982), for example, regard the phlogiston theory as "pre-paradigmatic science." Of course, one can find successes of the theory as well, but whether these are enough to meet the realist's requirement is debatable. The point here is simply that one may be able to eliminate certain counterexamples to the realist inference by denying that they meet the threshold of "successful." The phlogiston theory provides one plausible example.

[6] The Arago spot experiment, a failed attempt to refute Fresnel's wave theory on the basis of a seemingly absurd prediction, is often regarded as a paradigmatic case of an experimental test of a theory.

Is there any way to avoid the antirealist implication? Unfortunately, it seems not. Recall the two problems described above: (1) the inference from success to truth faces counterexamples in these abandoned theories and (2) current theories are cast into doubt by the generalization from these counterexamples to current theories. The realist would like to avoid both (1) and (2) by weakening the notion of truth from full strength to approximate truth. If Fresnel's theory was approximately true then the core intuition that success indicates (some bit of) truth is preserved and the generalization to all theories weakens the realist's claim only slightly. The problem with approximate truth, as indicated earlier, is that it seems to be untenable given that (i) truth depends on ontological matching between the theory and the world and (ii) abandoned theories got the ontology of the world wrong.

We might try to deny (ii) by reconsidering the reference of terms so that the term "ether" in Fresnel's theory refers to aspects of the electromagnetic field (or perhaps the quantum field) rather than a substance permeating all of space.[7] In this way we can preserve the approximate truth of theories like Fresnel's, but we do so at a cost. First, we must adopt a revisionary theory of reference; Fresnel and others using the term "ether" were not referring to the material substance they thought they were. Many of the beliefs Fresnel and other scientists had about ether were incorrect, even some that seem purely definitional (e.g., "ether is an elastic solid in which light waves vibrate.")[8]

A more worrying consequence of this strategy is that it suggests an overly per-

---

[7] This strategy is taken up by Hardin and Rosenberg (1982) who draw on the causal theory of reference developed by Kripke (1980) and Putnam (1975b).

[8] Of course, we could render such statements true by additional semantic revision, but this only serves to underscore the extent of the operation needed.

missive view of approximate truth (see Laudan (1984b)). If we can save Frensel's ether theory by charitably reconsidering the referent of "ether" what's to prevent us from saving the phlogiston theory in a similar manner? Phlogiston was supposed to be the substance responsible for phenomena we now know as oxidation and reduction. According to this theory, combustion produces phlogiston, so that an enclosed flame eventually goes out because the air has become completely "phlogisticated." In this context, phlogiston played a similar role to the absence of oxygen, hence, we might propose that "dephlogisticated air" refers to oxygen.[9] The problem is not that this strategy is difficult to believe, but rather that if we extend charity to all non-referring terms in past theories we lose any meaningful distinction between false theories and those with some degree of approximate truth. If one is happy with the consequence that phlogiston theory is approximately true, replace it with whatever theory one wishes to regard as false. Using a causal theory of reference one will be able to reconstruct the theory in such a way that the theory is saved.[10]

Kuhn's problem, it seems, has created a trap for the realist. The realist wants to assert the truth of current theories, but to do so she must regard their posited entities as real. History casts doubt on scientific ontology by presenting numerous compelling cases of successful theories wedded to forgotten entities. Yet, to have any part of truth it seems a theory must make reference to real entities, so the

---

[9] Indeed, Kitcher (1993) claims that there are contexts in which such an understanding is appropriate.

[10] There are more sophisticated approaches along these lines as well. Psillos (1999) adopts a combination of causal and descriptive theories of reference that aims to strike a balance between permissiveness and triviality. On his view, Frensel's term "ether" successfully refers, but "phlogiston" does not. The worry with such approaches is that what a scientific theory is taken to mean seems to depend entirely on whether, or to what extent, the theory agrees with what we now believe.

truth of past theories as well as today's are undermined. Simply put, truth requires ontology, but history undermines just that, so realism can't work.

2.5 Truth and Correspondence

We've already explored one attempt to avoid the realist trap by blocking the inference from the history of science to ontological skepticism. The idea was to argue that past theories were not *really* committed to what they appeared to be committed to because the real meaning of certain terms was not apparent. Another way out of the trap would be to question the connection between truth and ontology. If there was a way to endorse the truth of theories without being committed to their ontological posits, then perhaps we could avoid the trap that way. Of course, there is a long history of philosophers associating truth with the existence of a theory's posits. Quine (1948) famously proposed a theory of ontological commitment that requires one to endorse all of the posits of an accepted theory and countless others have made similar remarks. One obvious reason for this tight connection is that successful reference of the terms that occur in our theories seems to be a prerequisite for truth. But this assumption can be challenged.

One way to create some space between truth and ontology is truthmaker theory. This view's point of departure is the intuitive idea that for each truth, there is something in the world that makes it true. But crucially, it claims that truthmakers and facts need not be related one-one. A single truthmaker can make true many propositions. This allows the truthmaker theorist to pair down her ontological commitments as compared to a standard correspondence theorist.

> One of the benefits of truthmaker theory is to allow that "x exists" might be made true by something other than x, and hence that "a exists" might be true according to some theory without a being an ontological commitment of that theory. (Cameron 2008, 4)

Thus, truthmaker theory offers the following escape from the realist trap. Successful theories are (approximately) true, even if there is reason to be skeptical toward their apparent ontological commitments, because the truthmakers for the theory needn't include its ontological posits. Hence we might regard Newton's second law as true, while denying there are really any Newtonian forces or masses, because something else makes it true. Of course, the immediate question is just what the truthmakers could be if not the entities that are ostensibly referred to? Cameron (2008) has in mind cases such a metaphysical nihilist who will say that "tables exist" is made true, not by tables themselves, but by simples arranged tablewise. Extending this approach to real cases in science is more difficult. In the case of past theories, it is tempting to say that what makes them true is what is actually the case according to our current best theories. For example, Newton's second law is made approximately true by facts about the dynamics of our world as given by contemporary physics.

There are several problems with this approach. First, it seems to reduce to the reference-changing approach discussed above and hence inherits its problems. In this case we say the terms have the same meaning, but are made true by what is really the case, but this difference is largely superficial. It is still the case that nearly any theory can be saved in this way and that scientists have many false beliefs about the content of their own theories. A second problem is that this puts us in no position to question the ontological commitments of current theories. If what makes T true is what there is according to our best theories, then if T is one of the best theories, its ontological commitments are binding. Thus, we cannot take seriously the skepticism toward the ontology of our current theories recommended by generalization from historical cases. Truthmaker theory itself is certainly up to the task of separating truth from ontological commitment in a way that would avoid the realist trap, but without some procedure for finding the truthmakers for scientific theories, the view is incomplete. The most obvious approach to finding truthmakers fails and there is no clear replacement.

A similar approach to truthmaker theory has been advanced by Horgan and Potrč (2000, 2008) in service to their monistic ontological picture. "Indirect correspondence" as they understand it is "semantic correctness under contextually operative semantic standards" (2008, 37). In other words, a statement is true just in case there is a certain state of the world that makes it appropriate to assert the statement in the particular context it occurs. The work Horgan and Potrč have in mind for indirect correspondence is similar to Cameron's motivation. They want to be able to preserve the truth of common sense ontological claims (e.g., "there is a table") in the setting of a revisionary ontological picture in which common sense objects aren't real. Horgan and Potrč's indirect correspondence is a species of truthmaker theory on my characterization of the latter (but not theirs), but what makes a statement true depends on the context in which it occurs and often involves contributions from various aspects of the world rather than a single truthmaker.

Once again it is clear that indirect correspondence can permit the realist to escape from the trap by regarding theories as true without being committed to their ontological claims. The notion of a truthmaker that is diffused throughout the world is also a welcome addition in light of the difficulty in trying to locate truthmakers for scientific theories other than their posits. Rather than force and mass making Newton's second law true, for example, we can now say that aspects of the entire world allow for the semantic correctness of the law in certain contexts (i.e., those of relatively low speeds and short distances). In fact, by the standards of indirect correspondence it may be possible to say that Newtonian mechanics is not just approximately true, but is in fact true *simpliciter* when deployed in appropriate contexts. Yet, as with truthmaker theory in general, indirect correspondence does little to help specify just what makes scientific theories true. We can say that general features of the world contribute to a theory's truth, but it's hard to see how to go beyond that to a view that would satisfy the realist's explanatory needs.

Another problem faces both truthmaker theory and indirect correspondence (as well as related approaches) as accounts of the realist's theory-world relation.[11] Both of these proposals deal with the truth of linguistic entities—sentences, propositions and the like—while many philosophers of science now regard scientific theories as non-linguistic. The model-theoretic or "semantic" view of theories (Giere 1988; van Fraassen 1980; Suppe 1989; Suppes 2002) takes theories to be collections of models that are applied to target systems in the world. On this view, models may or may not be mathematical, but they are certainly not linguistic in the sense of consisting

---

[11] It is important to note that these views are only being rejected as accounts of the theory-world relation and not for other purposes. For example, the realist may wish to use these relations to account for the truth of everyday utterances once issues of metaphysics are settled. I am sympathetic to deploying these relations in that context (I agree with French and McKenzie (2012) that such tools are essential to making sense of structural realism), but maintain that a different relation is needed for describing how theories (at least those of fundamental physics) relate to the world.

of a collection of sentences, axioms expressed in first-order logic, or anything of that sort. Thus, when we talk about the truth of a theory on this view we are talking about a relation between a model and its target. While it may be possible to describe that relation in terms of truthmakers ("there is something in the world that makes this model true of it"), this doesn't seem to be the most natural way to describe the situation. Neither does indirect correspondence easily adapt to the case of theoretical models. To the extent that it does, the semantic correctness of a given model appears to be too weak to satisfy the realist's need for explanation.[12]

Is there a way, then, to escape the realist trap consistent with a model-theoretic view of theories? That is, can we find a notion of truth for models that is sufficient for explaining success but doesn't commit us to the model's ontology? Here I think Kuhn offers a useful hint later in the same passage quoted above.

> Perhaps there is some way of salvaging the notion of 'truth' for application to *whole* theories. (1962/1996, 206, my emphasis)

The notion of whole theories is important here. Where simple correspondence went wrong for the realist was in making the truth of the whole theory depend on the reference of its singular terms. Even the more nuanced truthmaker theory and indirect correspondence approach maintain the piecemeal approach to truth. In the case of truthmaker theory it is typically assumed that there are truthmakers for each individual claim a theory makes, even if they are sometimes shared by more than one claim. Indirect correspondence is also aimed at rendering true bits of theories, even

---

[12] The realist seeks an explanation for why a given theory T is successful. The fact that it—or one of its models—is the correct model to deploy in some context does not provide such an explanation. In order to do that one would have to answer the further question: why is it appropriate to use this model in this context?

though it does so by appealing to global properties of the world. Moreover, even traditional model-theoretic (semantic) approaches have assumed that truth comes via local correspondence.

Giere (1988), for example, argues that the similarity between models and the world is what accounts for the truth of certain models. Yet, in trying to specify just what that notion comes to, it seems that we need to appeal to a "local" notion of correspondence again. The claim that a model and its target are similar seems to require specifying the respects in which they are similar (or perhaps identical). But aren't these respects just the ontology of the system? One might propose instead that the similarity consists in the matching of the phenomena in the world with elements in the model. But, unless we mean something more than mere observable phenomena, this hardly gets us a realist position. What is needed is a way to say why these phenomena occur in both the model and the world, not *that* they occur. Clearly though, if we must reduce similarity to a pairwise matching of elements of the model with elements in world, we have made little progress toward our goal of separating truth from ontology.

We might propose to take similarity as basic and reject the need for any further analysis. But on its own similarity seems to be far too vague to serve as a replacement for correspondence. Nearly anything can be regarded as similar to anything else is some respect, so it is insufficient to the task of explaining why a theory works to say that its models are similar to the world. If, however, similarity could be understood in terms of a more developed notion of model-world correspondence that is both specific and global, perhaps we can make sense of Giere's view and give the realist a way out of the trap.

2.6 Global Correspondence

The guiding idea behind global correspondence is to make sense of the notion of entire models matching the world (or some bit of it). The problem is that when we try to specify the nature of this "matching" we seem to be forced back into a form of local correspondence. The only way, it seems, to understand how models correspond to the world is in terms of some of their elements representing real things in the world. But in going this route all the familiar obstacles to local correspondence—changes in theoretical ontologies, problems of reference, etc.—reemerge regarding our models. Yet, rejecting the idea that certain elements of our best models represent things in the world seems to leave us without any means of specifying what is meant by global correspondence. Moreover, as was mentioned above when Giere's notion of similarity was discussed, such a vague notion as global correspondence can hardly support the kind of realist position we are looking for. So, is there a way to develop global correspondence that is both specific and global?

It is here that we can usefully borrow from the philosophy of mathematics. The problem we are struggling with is one that has long plagued that discipline. The philosopher of mathematics who is interested in claiming the truth of mathematical theories seems to be committed to a variety of abstract objects such as numbers and sets. While there have been many proposals for avoiding such commitments, one in particular is of special interest in our case: mathematical structuralism. Here is how Shapiro characterizes the view:

> Rather than focusing on the nature of individual mathematical objects, such as natural numbers, the structuralist contends that the sub-

> ject matter of arithmetic, for example, is the structure of any collection of objects that has a designated, initial object and a successor relation that satisfies the induction principle (1997, 61).

If mathematics isn't about objects, but the structures they appear in, then we can avoid many of the problems of a mathematical ontology. There is no Platonic number one that is referred to by "1," nor does this refer to a set or any other individual object. Rather, numerals stand for places in the natural number structure. We may talk about this structure abstractly, or in terms of concrete systems that instantiate it, but either way mathematical objects are merely elements or aspects of a certain structure.

How might this help in the case of science? The thought is simple. If we view science as concerned with structures, then we similarly avoid the problems associated with a commitment to unobservable objects. Specifically, structuralism offers a path to global correspondence that doesn't reduce to local correspondence. In the case of mathematical structuralism, a natural number system is anything with the right structure; what makes it a model of the natural numbers isn't any sort of local correspondence between its elements and the numbers, but rather the structure it exemplifies as a whole. In keeping with this line of thought, we might suggest that global correspondence consists in a model capturing only the structure of the world. This is a feature of the entire model and one that doesn't require any ontological matching at the level of individual objects. This position is known as structural realism or scientific structuralism.[13]

---

[13] I will take the term "structural realism" to indicate the thesis that science tells us about the structure of the world only. In other words, the structural realist holds that the representational

Of course, we might wonder if this proposal is specific enough. We rejected Giere's notion of similarity because it seemed too weak to vindicate realism as we understand it. But unlike similarity, shared structure can be given a formal analysis in terms of isomorphism. Two systems have the same structure just in case they are isomorphic, that is, there is a structure-preserving mapping from one to the other. If we take structures to be set-theoretic constructions, then they are isomorphic iff there is a mapping of the domain of one set onto that of the other that preserves the relations between elements. The appeal to elements of the domain might appear to reduce the global correspondence to local correspondence yet again, but this is not so. Isomorphisms map elements of one set to elements of another without any regard to the nature of these elements. Hence, the elements of the two structures need not be similar, much less identical, for those structures to be isomorphic. The notion of isomorphism can also be applied to category-theoretical structures and graphs if these should better represent mathematical models.

2.7 Problems with Structural Realism

Structural realism (SR) is a promising way out of the realist trap, but it faces several challenges that threaten its viability. I will address three pressing questions here: (1) can isomorphism replace local correspondence? (2) can SR make sense of theory change? (3) how many structures does the world have? My goal is to make a case

---

content of successful models is exhausted by their structure. I will reserve the term "scientific structuralism" to pick out the view that the world described by our best scientific theories is fundamentally "structural." According to the scientific structuralist, science suggests that fundamental reality is populated by irreducible structures. The two views are related, but not necessarily so. The ontic structural realist, for example, endorses structural realism and scientific structuralism while the epistemic structural realist only endorses the former. I will have more to say about this in Chapter 4.

for SR as a tenable option for the realist.

2.7.1 Can Isomorphism Replace Local Correspondence?

According to SR, the sense in which our theories are approximately true is that they are roughly correct about the structure of the world. Above we shifted the focus from theories to models and hence the structural realist's claim is now that models are approximately true when they share the structure of the world (i.e., such models and the world are isomorphic). On the standard set-theoretic definition, a structure $S = \langle D, R \rangle$ contains (i) a set of objects, or domain D and, (ii) a set of relations R on D. We may think of a model as a particular structure picked out by a scientific theory.[14] Thus, the structural realist claims that models and the world are isomorphic: the members of the domain D_M can be mapped onto the domain D_W of the world one-to-one while preserving all of the relations R.

This set-theoretic understanding of structures gives rise to an immediate difficulty: isomorphism is a relation between structures, but structures are logical/mathematical objects, so structural realism appears committed to viewing the world as a logical/mathematical entity. It is not the intention of the structural realist to conclude that the world is a mathematical object,[15] yet this seems to follow from the claim that the world and our best models of it share a common structure. One possible way to avoid this Pythagorean conclusion is to consider data models,

---

[14] I am intentionally avoiding specifying the exact relationship between theories and models. Originally, a model was understood in the logical sense of an interpretation that makes true all of the sentences of the theory. Obviously, if we view theories as non-linguistic entities, this understanding can't be quite right. A model, for our purposes, can be thought of as a particular instance of a given structure. Theories specify what models there are, and hence what structures there are.

[15] Although, Tegmark (2008) and Dipert (1997) are possible exceptions.

which are cleaned-up versions of the phenomena our theoretical models seek to explain.[16] These are mathematical models and hence the structural realist can argue that our theoretical models are (close to) isomorphic to them. However, it is clear that such a maneuver only pushes back out original worry: how can we say data models and the world are isomorphic?

It useful to revisit mathematical structuralism in the service of answering this question. Mathematical structuralists distinguish between two version of the view: *ante rem* structuralism and *in rebus* structuralism. The distinction closely parallels that of universals and tropes in the metaphysics of properties. The *ante rem* structuralist claims that mathematics describes structures understood as abstract objects that exist independently of other objects, while the *in rebus* structuralist holds that the structures described by mathematics are only found in their concrete realizations. Thus, the *in rebus* mathematical structuralist and structural realist appear to be in same boat; both views need to find mathematical/logical structure instantiated in the world. Given that *in rebus* structuralism is generally regarded as a coherent view, it must be at least *possible* for a physical system to instantiate a mathematical structure.[17]

There is, however, a further difficulty with identifying global correspondence and isomorphism. Isomorphism is a purely formal relation, and as such, the constraints it places on physical systems are quite weak. If we consider the structure $S = \langle D, R \rangle$ and ask of some set of objects D' whether it has the same structure, the answer will

---

[16] See van Fraassen (2008) for a discussion.

[17] Another example of isomorphism involving physical objects is found in aesthetics. One position in that field is that art represents its target via isomorphism. Both the piece of art and the target are typically physical objects, so it is assumed here that physical objects can instantiate structures (French 2003).

always be yes so long as they have the same cardinality ($|D| = |D'|$). The reason why is that the relations, if defined extensionally, are simply n-tuples of objects from the domain ($R_i =< a_1, ..., a_n >, a_i \in D$). Thus, isomorphism between D and D' requires (1) a mapping f of the members of D into D' and (2) that if $< a_1 ... a_n >$ stand in relation R_i then $< f(a_1) ... f(a_n) >$ stand in relation $f(R_i)$. But there is always *some* relation between objects in the new domain that holds between *any* of its objects.

This makes the claim that the world and a model are isomorphic nearly vacuous and surely insufficient for the realist's explanatory needs. A key assumption for this problematic result is that relations R_i are defined extensionally. If the structural realist accepts this assumption, then isomorphism is indeed too weak to support the realist inference, but the assumption need not be accepted. The structural realist must claim that the relations R_i are not defined extensionally but instead are given a physical interpretation in the representational model. In this way, the structural realist can insist that a relation-preserving isomorphism is a robust form of global correspondence.[18]

2.7.2 Can SR Account for Theory Change?

The problem of theory change is that ontology is often not preserved when theories are replaced, and hence the realist falls into the trap described in section 2.4. The

---

[18] There is a worry with giving up on extensional definitions to save isomorphism, however. Picking out certain relations and requiring representational isomorphisms to preserve them introduces a non-structural element into the structural realist picture. SR cannot be understood as claiming simply that models and the world share a purely formal structure; interpreting some elements of our scientific models cannot be avoided. But this shouldn't come as a surprise; simply claiming that the world is described by, say, the Schrödinger equation without giving any interpretation of the variables it contains is to say very little about the world.

initial problem is simply that abandoned theories were talking about things (e.g., phlogiston, luminiferous ether) that we no longer take to be real. This issue simply does not arise on the structural realist view. All theories are talking about the same thing, namely, the structure of the world. Of course, theories may disagree about what the structure of the world is, but the worry of trivial falsity on account of talking about the wrong kinds of things is avoided. Thus, at the very least, the structural realist can say that past theories (and by extension, today's theories) *may* be approximately correct in describing the structure of the world.

While this constitutes an advance in the realist's position, it is surely not enough to satisfy the realist that today's best theories have the mere *possibility* of being approximately true. To say more, the structural realist must argue not only that all theories are talking about the same thing, but that there is some "coherent direction of development" in the structure attributed to the world (to borrow Kuhn's phrasing). It is important to note that the structural realist need not argue that past theories are fully isomorphic with the theories that replace them. The realist regards science as progressive, so we shouldn't expect new theories to simply restate their predecessor's claims. But what we should expect on a realist view is that past theories were close enough to right about the structure of the world so as to explain their success. In addition, we should expect that successive theories, while not identical in structure, propose structures that build on those that came before rather than developing discontinuously.

Unfortunately, the formal notion of isomorphism doesn't seem up to the task of making precise the relation between successive theories; either they are isomorphic, in which case they are effectively equivalent for the structural realist, or they are not, in which case we have no grounds on which to claim their continuity.[19] It is in this context that da Costa and French (1990, 2003) suggest the idea of partial isomorphism.

> $\mathcal{A}$ is partially isomorphic to $\mathcal{A}'$ when a partial substructure of $\mathcal{A}$ is isomorphic to a partial substructure of $\mathcal{A}'$,
>
> [where a partial structure has the form $\mathcal{A} = \langle A, R_i, f_j, a_k \rangle_{i \in I, j \in J, k \in K}$]. (2003, 49)

In other words, two structures are partially isomorphic when certain of their elements (in this case relations R_i) stand in one-to-one correspondence. Partial isomorphism can come in degrees and in different respects depending on which elements are in correspondence with each other. This framework allows us to make sense of the idea that models are continuous across theory change. If certain key elements of structure are preserved, then there is a clear sense in which the change is continuous.

Consider two familiar examples. In the history of optics, we find that Frensel's equations for describing the reflection and refraction of light can be derived from Maxwell's laws of electromagnetism. Similarly, as we shall see in the next chapter, Newton's laws can be derived as limiting cases from relativity theory. These are two instances of what seems to be an institutional norm in physics: successor theories should be able to recover the governing equations of the theories they replaced.[20]

---

[19] Brading and Landry (2006) argue that we don't require a specific formal account of shared structure to claim that theory change is continuous. The worry, however, is that without a formal framework the claim that successor theories share structure with their predecessors is no less vague than the claim that the theories are "similar."

[20] This heuristic is sometimes called—following Niels Bohr's more limited usage—the "correspondence principle."

The recovery of equations after theory change provides evidence that certain relations are preserved and hence that there is some degree of partial isomorphism between earlier theories and their successors. Moreover, the particular relations that are carried over are precisely those responsible for the instrumental success of the earlier theory. Thus, the realist can appeal to those parts of the models of the abandoned theory to explain why the theory enjoyed instrumental success.

This brief discussion cannot establish the structural realist's claim that science is progressing towards truth, but it goes some of the way toward making that claim plausible. The argument may be sketched as follows: (1) instrumentally successful theories are getting something right about the structure of the world; (2) what they are getting right are the relations encoded by the laws of the theory used in successful predictions; (3) the recovery of laws after theory change gives us reason to think such relations are preserved by subsequent theories; hence, (4) each successive theory incorporates key parts of the structures of earlier theories and builds on them.

2.7.3 How Many Structures Does the World Have?

According to SR, successful models tell us about the structure of the world. But, there are quite a lot of different models, with different structures, which are successful. In some cases, even a single theory may give rise to non-isomorphic models. Psillos provides the following example:

> Take Newtonian mechanics, where $\mathbf{F} = m\mathbf{a}$, and compare it with a reformulation of it, according to which $\mathbf{F}$ always is the vector sum of two more basic forces $\mathbf{F}_1$ and $\mathbf{F}_2$. Here we have two non-isomorphic

> structures, which are nonetheless empirically equivalent. Which of them is the structure of the Newtonian world? (2006, 562)

These two formulations are empirically equivalent (and perhaps equivalent in a deeper sense as well), but are non-isomorphic, so the structural realist seems unable to say that they both capture the structure of the Newtonian world. There are three basic strategies the structural realist might take in response to cases such as this: they might hold that (a) despite appearances, the different formulations are really isomorphic, (b) there is a single correct formulation, or (c) the world has several different structures.

Option (a) is attractive in this particular case. A real example that is similar to Psillos' case is that of the Schrödinger and Heisenberg formulations of quantum mechanics. The formalism of quantum mechanics contains states and observables (roughly, "measurable quantities") of those states. The Schrödinger and Heisenberg pictures differ on how quantum systems evolve, that is, their dynamics. On the Schrödinger formulation, states change over time (in accordance with the Schrödinger equation) and the observables remain constant, while on the Heisenberg formulation, the *observables* are time-dependent and quantum states remain constant. Despite having superficially distinct structures, it can be shown the these two formulations of quantum mechanics are in fact different ways of describing the same underlying structure.[21] We should expect that the same would apply in Psillos' case. For example, we might regard the *real* structure of the classical world as given by the Hamiltonian formulation of classical mechanics and his two formulations as

---

[21] That structure is one in which states correspond to rays in Hilbert space and observables are identified with operators acting on those rays.

just different expressions of that common structure.

However, there is a problem with generalizing this approach to all cases of different formulations of theories. In particular, if we move beyond non-relativistic quantum mechanics to consider quantum field theory (QFT), we find that option (a) is much harder to adopt. In QFT there are unitarily inequivalent representations corresponding to different ways of quantizing a field with infinite degrees of freedom. In this case representations are non-isomorphic and there is no underlying common structure to appeal to (because they are unitarily inequivalent). This presents a difficult challenge for the structural realist opting for (a) as a general strategy.

The problem of unitarily inequivalent representations in QFT leads Steven French to conclude that the structural realist should pick one representation ("conventional" QFT) as the correct one in this case (French 2012b, 134). As French points out, scientific realists always must pick the best theory on offer and often this goes beyond the empirical content of that theory. The problem with adopting this approach (option (b) above) in general is that there may be cases where non-isomorphic models are on equal footing. We may encounter situation where, even taking account of super-empirical virtues, our choice of formulation is underdetermined. In such cases the proponent of (b) is compelled to side with one formulation over another without being able to justify her choice. Certainly this is not an ideal position for the structural realist to find herself in.

Finally, we may simply grant that there are many different structures, all of which correspond to the world. Immediately this causes a problem for our understanding of correspondence as isomorphism. Isomorphism is transitive, so if two models are isomorphic with the world, then they had better be isomorphic with each other. Yet, recall that isomorphism had to be abandoned in above in favor of partial isomorphism, and note that partial isomorphism is *not* transitive. Two models may each be partially isomorphic to the world in different respects or to different degrees. One model may capture certain relations in the world, while another captures other relations. Thus, we could have models which are incompatible with each but nevertheless each enjoy correspondence with the world to some degree in some respect.

A way of visualizing the idea the world having multiple structure can be found in Ladyman and Ross's discussion of real patterns (2007, 220–238) (originally proposed in Dennett (1991)). If we think of structures as patterns, then just as there can be many patterns in a given set of data, there can be many structures in a single world. To continue the analogy, these patterns will differ in the features they focus on and the amount of "noise" they contain just as structures differ in the respects and degree to which they are partially isomorphic to the world.

2.8 Conclusion

I do not pretend to have established the truth of SR here, instead my aim is merely to motivate the view and argue for its promise. SR is an attractive way of developing the idea of global correspondence and provides a way for someone who accepts the realist inference to avoid the realist trap. And while problems remain for the structural realist, I hope to have shown that there are ways out as well.

CHAPTER 3

Case Study: The Transition from Newton to Einstein

3.1 Introduction

Newtonian mechanics (NM) is an extremely useful scientific theory with myriad applications. This is true despite the fact that NM has been replaced by relativity as our best theory of the large-scale structure of the physical world. This creates a problem for the scientific realist who regards our current best theories as approximately true. Given that NM has been replaced, how can we be justified[1] in our continued use of it? The standard answer is that NM is a limiting case of the approximately true theory that replaced it: relativity theory.[2]

Yet, it is not obvious what this claim comes to. In what sense is NM a *limiting case* of relativity theory and how does that justify its use? In this chapter I will attempt to shed some light on the limiting relation that holds between NM and the special theory of relativity (STR). My contention will be that structural realism provides a perspective from which one can defend the standard reply; NM can be

---

[1] By "justification" here I intend the following. According to a realist understanding, we have good reason to use relativity theory because we have strong empirical evidence of its approximate truth. Yet, NM is no longer among our best physical theories, and hence we do not have strong evidence for its approximate truth (in fact, we seem to know it is false). The question of justification is this: what reasons can we offer to legitimize our use of NM, an apparently false theory?

[2] One might think there is a simpler answer: there is no alternative. In many cases modeling a physical situation in relativity theory is too complicated, and so NM is the only *available* option. Yet, without some claim about the relation between NM and relativity, there is no reason to think that NM—a theory that has been superseded—will be of any use. *Being a limiting case* is supposed to be just such a relation.

seen as a limiting case of STR when we focus on the *structures* each theory describes.

Of course, it is *general* relativity that has actually supplanted NM as our current best theory of the large-scale structure of the world,[3] but many of the points below regarding the relation between NM and STR will carry over to general relativity. After all, the Minkowski spacetime of STR is a model in general relativity, albeit an oversimplified one. This suggests that focusing on the conceptual issues that arise in the relation between NM and STR will have relevance for general relativity or, at least, that these issues will not *go away* with the move to general relativity.

In addressing this particular case, I aim to add to the many examples of inter-theoretic relations already discussed by structural realists. Perhaps best known of these is John Worrall's (1989) discussion of the transition from Fresnel's ether theory of light to Maxwell's electromagnetism. Holger Lyre (2004) extends Worrall's argument by illustrating structural continuity between Maxwell's electromagnetism and quantum electrodynamics. Simon Saunders (1993) attempts to show the same for the case of Ptolemaic and Copernican astronomy. James Ladyman (2011) takes on the case of phlogiston and contends that phlogiston theory made true claims about the structure of chemical processes. Finally, Jonathan Bain and John Norton (2001) argue for structural continuity in theories of the electron. Each of these cases provides additional evidence in favor of the view that *structure*, rather than object-based ontology, should be emphasized by scientific realists. This case study also has that moral.

The chapter will proceed as follows. Section 2 will present the standard textbook

---

[3] In fact, we know that even general relativity cannot be exactly right. Ultimately, a theory of quantum gravity will be needed to combine the insights of quantum mechanics and relativity, but there is no such universally accepted theory of quantum gravity at this time.

account of the limiting relation between STR and NM based in recovering Newton's laws of motion from their relativistic analogs. Challenges to the standard account will be presented in section 3. In section 4, a model-theoretical approach to the relation between STR and NM will be advocated. Section 5 will argue that this approach is best understood from the perspective of structural realism.

3.2 The standard approach

> Faced with the results of special relativity, we should in principle rewrite all our mechanics accordingly. But we know that this is not necessary. The Newtonian scheme, although it is strictly correct only in the limit of vanishingly small velocities, works beautifully in an enormous variety of situations. This ... is because the greatest velocities that we encounter in the dynamics of ordinary macroscopic objects are still minute compared to the velocity of light ($v < 10^{-5}c$).(French 1968, 161)

This quotation reflects a common sentiment. In light of STR, we must regard NM as correct only in the limit of "vanishingly small" velocities, but most applications approximate this limit, so their use is justified as well. To begin with the first part of the claim, in what sense is NM correct in the limit?

The standard textbook presentation centers on the coordinate transformations deployed in each theory. In NM, two inertial reference frames[4] S and S' moving at a constant velocity v (in the x direction) with respect to each other are related by

---

[4] An inertial reference frame is, roughly, a Cartesian coordinate system associated with a non-accelerating trajectory through spacetime. More precisely characterizing *inertial* states of motion is difficult in this setting, which provides a reason for the model-theoretic approach advocated below.

the Galilean transformations.

Galilean transformations

$$x' = x - vt \tag{3.1}$$

$$y' = y \tag{3.2}$$

$$z' = z \tag{3.3}$$

$$t' = t, \tag{3.4}$$

$$\tag{3.5}$$

where x, y, z are orthogonal spatial coordinates, and t stands for time.

In STR, inertial reference frames moving at a constant velocity with respect to each other are related instead by the Lorentz transformations.

Lorentz transformations

$$x' = \gamma(x - vt) \tag{3.6}$$

$$y' = y \tag{3.7}$$

$$z' = z \tag{3.8}$$

$$t' = \gamma(t - \frac{vx}{c^2}), \tag{3.9}$$

where γ is the Lorentz factor: $\gamma = (1 - v^2/c^2)^{-1/2}$.

The chief difference in the two sets of transformation rules is the factor γ that appears only in the Lorentz transformations. Yet, it is clear that if we take the limit where $v/c \rightarrow 0$, then $\gamma \rightarrow 1$, and the Lorentz transformation rules reduce to their Galilean counterparts.[5] This goes some of the way toward connecting the idea of

---

[5] Note also that in the limit the additional factor in the time translation rule also goes to zero

the limit of "vanishingly small velocities" with the recovery of NM from STR.

The transformations take on additional importance when supplemented with a further claim:

> *Invariance*: the same dynamical laws hold in all inertial reference frames.

If we assume Invariance, then the transformation rules act as constraints on the dynamical laws. The laws of NM must be invariant under the Galilean transformations, and the dynamical laws of STR must be Lorentz-invariant.[6] Consider, for example, Newton's second law of motion.

$$\mathbf{F} = \frac{d\mathbf{p}}{dt} = \frac{md\mathbf{v}}{dt}, \tag{3.10}$$

where $\mathbf{F}$ is the net force on a body, $\mathbf{p}$ is its momentum, m is its mass, and $\mathbf{v}$ is its velocity.

The kinematic magnitudes— $\mathbf{p}$ and $\mathbf{v}$—must be given relative to some inertial reference frame, but, assuming Invariance, regardless of which frame is chosen, the same laws relating them will hold.

In the case of STR, there is a Lorentz-covariant[7] analog to Newton's law given by the following four-vector equation:

$$F^{\mu} = \frac{d\pi^{\mu}}{d\tau}, \text{ with } \pi^{\mu} \equiv m_o \frac{dx^{\mu}}{d\tau}, \tag{3.11}$$

($vx/c^2 \rightarrow 0$) so that the Galilean time transformation $t' = t$ is recovered at the limit.

[6] To be more precise, Invariance requires that the laws be *covariant* under the relevant transformations. Each of the equations below contains quantities that are frame-dependent (i.e., non-invariant), hence these equations cannot be invariant. Yet, we can require that if the equation holds in one inertial frame, it will hold in all inertial frames (with the frame-dependent quantities changing accordingly). This is what is meant by describing an equation as *covariant.*

[7] See previous footnote.

where F^μ is four-force, π^μ is four-momentum, τ is proper time and m_0 is rest mass.

One may also find a three-vector equation analogous to Newton's second law:

$$\mathbf{F_r} = \frac{d}{dt}(m_r \mathbf{v}), \tag{3.12}$$

where $\mathbf{F_r}$ is a 3-component relativistic force, m_r is relativistic mass and $\mathbf{v}$ is the frame-dependent 3-velocity.

Equation 3.12 resembles Newton's second law very closely, but one important difference between the relativistic and the Newtonian version (equation 3.10) is that the term m_r that appears in the former is quite different in character from the mass term m that appears in NM.

Mass in NM is an invariant quantity: it takes the same value in all inertial reference frames. Relativistic mass m_r, on the other hand, varies with velocity, and hence is a non-invariant, frame-dependent quantity. Yet, there is an invariant mass in STR, the rest mass m_0, which is related to relativistic mass by $m_r = \gamma m_0$. Notice that, as before, $\gamma \rightarrow 1$ when $v/c \rightarrow 0$, so the frame-dependent m_r takes the same value as the invariant rest mass m_0 at the limit. This is a good thing insofar as one thinks that frame-dependent quantities are artifacts of our representational apparatus rather than genuine facts about the world.

The standard approach simply equates m_0 with Newtonian mass m (as well as τ with t) and takes Newtonian and relativistic force (F^μ or $\mathbf{F_r}$) to become equivalent as $v/c \rightarrow 0$. Of course, actual applications of NM do not occur at the limit, but the assumption of the standard approach is that v/c is so small in most applications that it may be approximated as 0.

3.3 Problems with the standard approach

3.3.1 The limit conditions

Before discussing difficulties with the standard approach, it will be useful for what follows to distinguish four varieties of equivalence for scientific theories. First, *empirical equivalence* occurs when theories say the same thing about the "observable"[8] realm. Second, *formal equivalence* requires that the mathematical/formal constituents of the theories have the same form. Third, *physical equivalence* occurs when the theories say the same thing about the physical world (including the "unobservable"). Finally, *theoretical equivalence* is when the theories are in fact different formulations of the same theory.

The standard approach maintains that at the limit where $v/c = 0$ the dynamical equations of STR reduce to those of NM. This is combined with an assumption that these equations play a central role in the descriptive content of the theories (i.e., what STR and NM tell us about the world). Thus, the standard approach holds that at the limit NM and STR are both *formally* and *physically* equivalent. Even if we grant all of this (for the moment), a further problem remains. As mentioned above, actual applications do not occur *at* the limit, but merely *near* it.[9] So what justifies such applications?

The standard approach maintains that when velocities are very small with re-

---

[8] It's hard to say just what counts as observable, but the rough idea is that empirically equivalent theories can never be distinguished by experiment; they agree on all predictions capable of experimental verification.

[9] One may doubt even that applications are "near" the limit. There may well be some cases where NM is used simply because it is the only available theory for treating certain phenomena, regardless of the speeds (and distances) involved. The realist must either regard these applications of NM as less justified than those near the limit, or else provide a justification which goes beyond the standard appeal to a limiting relation discussed here.

spect to that of light $v \ll c$ STR and NM are nearly equivalent. But what kind of approximate equivalence is there between STR and NM? It seems that, on the standard view, the only kind of equivalence (approximate or otherwise) they enjoy is empirical equivalence.[10] This is problematic for the standard approach if it is aimed at providing a realist justification for our continued use of NM. The realist needs to claim that near the limit NM is approximately *true.* Thus, the relevant notion of equivalence is physical equivalence; near the limit NM and STR should say roughly the same thing about how the world is.

It is hard to see how the standard approach can support the approximate physical equivalence of NM and STR near the limit. Even if we grant (again, for now) that the central equations of the two theories are formally equivalent at the limit, they are not so once we depart from the limit. Equation 3.12 appears to have the same form as Newton's second law, but once we substitute m_0 for m_r we must introduce the Lorentz factor γ, which gives the equation a different form. Equation 3.11, on the other hand, already contains m_0 but only appears formally equivalent to Newton's law before one unpacks the different notions of momentum at work: $\mathbf{p} = \mathbf{mv}$ and $\pi^{\mu} = m_o dx^{\mu}/d\tau$.

In addition, once we consider applications away from the limit of $v/c \rightarrow 0$, other conditions must be imposed even for empirical equivalence/approximation. The Lorentz transformation rule for time, for example, contains the factor vx/c^2 which may be significant even if $v \ll c$. One must further specify that the distances

---

[10] Note that so long as we are considering situations sufficiently close to the limit, the central equations of NM and STR will be not merely *approximately* empirically equivalent, but empirically equivalent full stop. This is because at the level of detail that is experimentally accessible there will be no difference in the predictions of the two theories. As we move further away from the limit, there may be situations in which empirical equivalence is merely approximate.

involved are sufficiently "short" ($x \ll ct$). This illustrates once again that the formal equivalence of the equations *at* the limit disappears once one departs from the limit; at the limit only v/c needs to be considered, while near it one must also consider the distance variable x.

Thus, it seems that, on the standard approach, NM and STR are merely *empirically* equivalent near the limit, not *formally* or *physically* equivalent. As mentioned, this is problematic for the scientific realist. So long as the two theories are physically inequivalent, their agreement on empirical matters is of no help to the realist. Indeed, it makes matters worse for the realist by providing a powerful case of the underdetermination of theory by observation emphasized by antirealists.

The scientific realist needs a way to regard the two theories as approximately physically equivalent near the limit. The standard approach sketched above aims to do this by showing that the dynamical laws of STR reduce to the laws of NM at the limit. But, even if this is the case, *near* that limit the laws of these two theories are formally distinct. Moreover, the only obvious sense in which they "approximate" one another concerns their *empirical* content, not their *physical* content. The dynamical equations (3.10 and 3.11/3.12 above) take a distinct form away from the limit, and hence it is hard to see in what sense their physical content could be approximately equivalent.[11]

---

[11] It is worth emphasizing that I am assuming here a version of scientific realism according to which one looks at the mathematical equations of a theory and attempts to determine their physical content—what they tell us about the world—by interpreting the variables that occur in these equations. On such a picture it's hard to see how one could view equations of different forms as telling us roughly the same thing about the physical world. The solution, I contend, is not to abandon realism, but rather to revise how the realist thinks about the physical content of a theory.

3.3.2 Kuhn's objection

In the last section, we focused on situations away from the mathematical limit $v/c \to 0$. Thomas Kuhn, however, has objected that even at the limit it is wrong to regard either of the relativistic equations (3.11, 3.12) as equivalent to Newton's second law. He says such relativistic analogs to classical laws are,

> a special case of the laws of relativistic mechanics, they are not Newton's Laws... The variables and parameters that in the Einsteinian [laws] represented spatial position, time, mass, etc., still occur in the [relativistic analogs of Newton's laws taken at the Newtonian limit] and they there still represent Einsteinian space, time, and mass. But the physical referents of these Einsteinian concepts are by no means identical with those of the Newtonian concepts that bear the same name. (Newtonian mass is conserved; Einsteinian is convertible with energy. Only at low relative velocities may the two be measured in the same way, and even then they must not be conceived to be the same.)
>
> ...[The apparent derivation of Newton's laws from STR] has not, that is, shown Newton's Laws to be a limiting case of Einstein's. For in the passage to the limit it is not only the forms of the laws that have changed. Simultaneously we have had to alter the fundamental structural elements of which the universe to which they apply is composed. (Kuhn 1962/1996, 101–102)

Kuhn's objection is that the physical magnitudes which occur in Newton's laws—m, F, t—are distinct from their relativistic counterparts, and hence the apparent "derivation" of Newton's second law within STR is an illusion. Following Kuhn's own emphasis, much of the discussion has focused on the concept(s) of mass which occurs(occur) in STR and NM. As we've seen, the notion of relativistic mass that occurs in equation 3.12 is m_r, which does differ considerably from the Newtonian concept m.

Yet, we have also mentioned the relativistic concept of rest mass m_0 which, like m in NM, is an invariant quantity in STR. This gives us some reason to think that m_0 is simply Newtonian mass. Moreover, the *total* rest mass of a closed system is conserved. Why, then, does Kuhn claim that mass in relativity is not conserved? Presumably the answer lies in the relationship between mass and energy given by Einstein's famous equation $E_0 = m_0c^2$. This allows for a case where, for example, a massive particle decays into massless photons, so that the initial rest mass of the system is not conserved but is "converted" into energy. However, this way of speaking is misleading. If the system is truly *closed* then the rest mass associated with the total system *will* be conserved. This is simply because rest mass is proportional to the total (rest) energy of the system, which must be conserved if the system is closed.

Yet, there is an aspect of rest mass which distinguishes it from Newtonian mass. Consider again a massive particle that decays into some number of photons. The rest mass of an individual photon is 0, yet, we have just claimed that the total rest mass of the system (if it's closed) is conserved. This means that the *additivity* of rest mass must fail in STR: the total rest mass of a composite system is not, in general, the sum of the rest masses of its parts. Of course, if the components of the system are all at relative rest, additivity does hold, but it's not clear that this eliminates Kuhn's worry.

Consider, for example, the following remark from one of Kuhn's allies on this matter, Paul Feyerabend.

> Of course, the values obtained on measurement of the classical mass and of the relativistic mass will agree in the domain D', in which the classical concepts were first found useful. This does not mean that what is measured is the same in both cases: what is measured in the classical case is an *intrinsic* property of the system under consideration; what is measured in the case of relativity is a relation between the system and certain characteristics of D'. (1962, 122)

Feyerabend's point is that mass in STR, unlike its counterpart NM, is a function of the state of motion (in some inertial reference frame) of the system. In the case of rest mass this is perhaps less obvious, but in order for rest mass to be the appropriate quantity, a system must be *at rest*, so the concept builds in a reference to the state of motion of the system.[12] This would suggest that even though the rest mass m_0 of a composite system is additive at the limit of $v/c \to 0$, it is not identical to the Newtonian mass m *conceptually* because the additivity of m_0 depends on the state of motion of (the components of) a system, unlike m.[13]

---

[12] Alternatively, one may view m_0 as "non-kinetic mass" and insist that there is no reference to the state of motion of the system, even implicitly. But, the only time it is appropriate to use this notion of mass rather than one that takes into account kinetic energy (m_r) is when a system is at relative rest. Thus, it is not clear that thinking of m_0 as "non-kinetic mass" rather than "rest mass" removes the implicit appeal to the state of motion of the system.

[13] For his part, Kuhn's own view seems to be that equation 3.12 is *the* fundamental force law in STR and hence m_r is the relevant concept of mass. On this view, while m_r and m_0 take the same value at the limit, it is always the former that one uses in STR, even at the limit. The discussion here is meant to show that even if one grants the legitimacy of m_0 in STR, there remain Kuhn-style difficulties with identifying it with m in NM.

There have been many replies to Kuhn's objection. For example, Hartry Field (1973) argues that Kuhn is wrong to suppose that "mass" refers to different things in NM and STR just because certain claims made *about* mass in NM are incorrect from the perspective of STR. According to Field, what we know is that at least one of the following two claims of NM is false: (a) momentum = mass * velocity, and (b) mass is invariant (under changes of inertial frame of reference). Which of (a) and (b) is false, however, Field takes to be indeterminate. Instead he proposes that "mass" in NM partially denotes m_r and partially denotes m_0. This has the effect of making both (a) and (b) indeterminate in their truth value, while their conjunction is false. This gives us a situation in which "mass" in NM ("m") always refers to a mass in relativity (either m_r or m_0), but there is no fact of the matter as to which. There will still be certain claims, like the conjunction of (a) and (b), which are false regardless of which relativistic mass is taken to be the referent, but this doesn't undermine the view.

> [T]here is *certainly* nothing incoherent in the position that Newton was referring to mass even though he had a great many false-but-approximately-true beliefs about it; indeed, the fact that most of Newton's beliefs involving 'mass' come out approximately-true-though-strictly-false if we regard them as being about mass looks like evidence *for* ["mass" having the same referent in both NM and STR] rather than *against* it. (Field 1973, 465)

John Earman and Author Fine (1977) disagree with Field's claim about the indeterminacy of reference of "mass" in NM, but agree that the term refers to the same thing in both NM and STR. Earman and Fine take issue with Field's claim that "proper mass and relativistic mass play about equally important roles in special relativity theory" (1977, 468) and claim instead that there is a single concept of mass in STR, namely, rest mass m_0. The view that rest mass m_0 is the fundamental notion of mass in STR now seems to be the dominant view (at least among philosophers of physics)[14], and it naturally suggests identifying m_0 and m in NM.

Field and Fine and Earman agree that talk of mass in NM refers to the same thing as mass talk in STR, despite certain claims *about* mass in NM being false. Yet, their focus on the reference of mass terms is somewhat misleading as far as our present aims are concerned. The question we are considering is whether NM is a limiting case of STR, and in particular, whether the dynamical laws of STR reduce to Newton's laws at the classical limit. For such a reduction to occur, the laws of STR and NM must be equivalent in their descriptions of the physical world. Kuhn objects that this requirement is not met because the mass terms in the two theories "refer to different things" even at the limit. Now, one could argue (as Fine and Earman do) that if the mass terms refer to different things, then at most one such term can succeed in referring to mass in our world. If we assume that "m_0" successfully refers, then the Newtonian "m" fails to refer. But if "m" fails to refer then all claims involving it are false (or perhaps lack truth-value[15]), which is unacceptable to the

---

[14] Among physicists, the case for taking m_0 to be the only meaningful concept of mass in STR has been made by Lev Okun (1989). Although, see Sandin (1991) for an attempt to justify the value of relativistic mass m_r.

[15] Strawson (1950), for example, argues that reference failure does not entail the falsity of a sentence containing the non-referring term. In this case, taking such a view would mean that claims involving mass in NM are neither true nor false. This is equally unsatisfactory for the scientific realist, however, who must maintain that certain of these claims are true, or at least approximately so.

scientific realist. Thus, it seems that (at least from a realist perspective), we should take Newtonian and relativistic mass terms to be coextensive.

Kuhn's objection, however, centers on the disctict *concepts* of mass that occur in NM and STR. One way to illustrate this distinctness is to consider possible worlds described by the two theories. In the world described by the laws of NM, "m" describes one property, while in the STR world "m_0" describes a distinct property. There may be interesting trans-world relations between these two properties (e.g., they take the same value at the classical limit), but they are conceptual distinct nonetheless. If we follow Kuhn in this understanding of mass terms in NM and STR then the dynamical laws of these theories are physically inequivalent, even at the classical limit. If we insist on considering only the actual world—in which, let's assume, mass is m_0—then we immediately come up against questions of realism when we ask whether Newton's "m" refers. Therefore, a better understanding of Kuhn's objection in this context is that the mass terms in NM and STR *have distinct meanings* or *are associated with different concepts*. After all, such conceptual distinctness is all that's needed to undermine the physical equivalence of the two theories' dynamical laws.

Following Quine, we may distinguish between a theory's ontology and ideology, where the latter determines "what ideas are expressible in the language of the theory" (1951, 15). In this terminology, Kuhn's objection is that NM and STR differ in their ideology, even at the limit where $v/c \to 0$. We may note, without mentioning the issue of reference, that the two theories' mass terms do not *express the same ideas*, and, hence, expressions that contain "m" must be distinguished from those that contain "m_0." In order for the standard approach to succeed, the dynamical laws of the two theories must be physically equivalent at the limit. But, even if they express the same relations between variables, they aren't physically equivalent so long as the mass terms involved are conceptually distinct.[16] Field and Fine and Earman claim that "m" in NM refers to mass in STR, but (1) such an understanding begs the question against those who, like Kuhn, deny the realist framework, and (2) even if the terms are coextensive, they may still be conceptually distinct, which is enough to undermine the physical equivalence of the two theories at the classical limit.

Another aspect of the problem is the following. If one takes m_0 as the only meaningful notion of mass in STR, then the relativistic equation 3.12 and Newton's second law (3.10) have the same form *only at the limit*. Given that actual situations do not occur *at* the limit (as the previous subsection emphasized), it's hard to see how applications of NM are justified according to this view. The advocate of the standard approach faces a dilemma: one must either recognize m_r in the formulation of the Lorentz-invariant analog to Newton's second law, or recognize that any distance away from the mathematical limit of $v/c \rightarrow 0$ the dynamical laws of NM and STR have a different form.

Finally, throughout this discussion we have followed Kuhn in taking equation 3.12 to be the central force law in STR, but there are good reasons to prefer the 4-vector equation 3.11. Yet, switching to this law does not eliminate Kuhn's worry. First, even though m_r does not occur in the 4-vector equation, we have seen that there

---

[16] This is not how Quine himself would understand the situation. For Quine, ideology refers to the *extension* of predicates. But in this case, the extension of the predicate "mass" in the actual world is arguably the same in both NM and STR. Thus, to sustain the present point, we must understand ideology more broadly than Quine does. One way to do this would be to allow the extension of predicates to include possible worlds as well (Lewis 1983).

is reason to question the identification of m_0 and Newtonian mass. Second, even if one does identity m and m_0, other variables generate similar worries. For example, Newtonian absolute time t is replaced with special relativistic proper time τ. But the two concepts t and τ have little in common. τ is the invariant spatiotemporal interval between timelike-separated events. As such, it is not a *sui generis* measure of temporal passage like t in NM, but rather a special case of a more basic distance measure, the spatiotemporal interval. Moreover, this interval has features quite distinct from (Newtonian) duration. For example, distinct events in STR can have no "distance" between them if they are lightlike-sperated. In NM, by contrast, the duration between distinct non-simultaneous events (including those connected by a light signal) is always non-zero. Similar challenges could be leveled at the other kinematic terms that must be identified to regard the two theories as physically equivalent (even approximately) near the limit.

To summarize, the standard textbook approach faces serious difficulties. The fact that certain dynamical equations come to take the same form at the limit where $v/c \to 0$, does not show that NM is a limiting case of STR in a sense suitable to the scientific realist. Real applications of NM do not occur at the limit, but only close to it. Even at the limit, it is difficult to maintain that the variables which occur in the two equations are conceptually equivalent.

These considerations leave the original challenge unanswered: how can we understand the sense in which NM is a limiting case of STR such that the two theories are (approximately) physically equivalent when speeds are small and distances are short?

3.4 The model-theoretic approach

The model-theoretic approach focuses on the models associated with a theory rather than its laws. In the case of NM and STR, the relevant models are spacetime models.[17] There is an immediate problem in comparing the models of NM and STR, however. The former theory (NM) is traditionally formulated in terms of three spatial dimensions and one temporal dimension while STR is typically formulated in terms of a four dimensional spacetime. Yet, this problem is easily avoided. Both NM and STR can be described in terms of a background of four-dimensional spacetime. Moreover, both theories can be developed without appealing to reference frames or coordinate systems at all. Such an *intrinsic* presentation lays bare the real differences between the models specified by the two theories.

My approach here is to look at the intrinsic features of models in NM and STR and argue that the former are a limiting case of the latter. I will assume that to be a limiting case in the sense needed by the scientific realist is for the theories to be *physically equivalent at the limit*, in this case, where $v/c \rightarrow 0$. But what is the nature of this physical equivalence? Here is where general philosophical orientation leads to a divergence of views. One position, associated with traditional scientific realism, regards two theories as physically equivalent just in case they agree on what there is and how it evolves according to the dynamical laws of the theory. This was the idea underlying the traditional, equation-based, approach above.

Another option, which is associated with traditional forms of antirealism, holds

---

[17] Alternatively, an advocate of the model-theoretic approach may focus on state-space models. To adopt this strategy, one would have to identify the relevant state-space models and show how they are physically equivalent at the Newtonian limit and approximate each other nearby. I will not take up such a task here and focus only on spacetime models in what follows.

that empirical equivalence is the only intelligible sense in which scientific theories may be equivalent. Finally, there is the view favored by the structural realist: two theories are physically equivalent just in case their models are isomorphic and approximate each other to the extent that (the structures of) their models are similar.[18] The previous sections of this chapter have exposed difficulties with approaches that assume the traditional realist understanding of physical equivalence. In what follows I will argue that focusing on the structure of NM and STR allows one to overcome these difficulties without abandoning scientific realism altogether. Hence, our goal is to illustrate the sense in which models of NM and STR are isomorphic at the limit of $v/c \rightarrow 0$ and are similar in structure near the limit.

3.4.1 The difference between models in NM and STR

It will be useful to distinguish three different spacetime structures for what follows. NM was originally developed (by Newton) in a background of fixed absolute space and time. Let the version of spacetime inspired by this be *Newtonian spacetime*: a flat four-dimensional differentiable manifold with a Euclidean distance metric and temporal duration between *every* pair of points throughout time. In Newtonian spacetime velocity, acceleration and simultaneity are all absolute. In *Galilean spacetime*, by contrast, the Euclidean distance metric applies only to pairs of points *at the same time*. There is no preferred inertial frame and hence no absolute velocity, but acceleration and simultaneity remain absolute. Finally, in *Minkowski spacetime* we

---

[18] The notion of "models with similar structures" can be made more precise by appealing to the concept of *partial isomorphism*. Two structures are partially isomorphic when they share *some* of the relations which characterize them. See section 2.7.1 above and, for a more detailed discussion, da Costa and French (2003).

still have a flat four-dimensional differentiable manifold of points, but the Euclidean spatial metric and unique temporal metric are replaced by a single Minkowski metric, which is thoroughly spatiotemporal. As a result, simultaneity is no longer absolute. The distinction between accelerated and non-accelerated frames remains absolute, however.[19]

Galilean spacetime contains all that is needed to formulate NM, hence we will take it, rather than Newtonian spacetime, to give the spacetime structure of models of NM.[20] Minkowski spacetime gives the structure of models in STR. What, then, is the difference between these two spacetime structures? The primary difference is that the planes of simultaneity of Galilean spacetime have been replaced by the light-cone structure of Minkowski spacetime.

The distinct spacetime structures of NM and STR underlie the problems with the standard, equation-based approach criticized above. So long as the terms in the dynamical equations of STR refer to elements of Minkowski spacetime, and those of NM refer to elements of Galilean spacetime, the laws can only be regarded as empirically or (at best) formally equivalent.

---

[19] While there is an absolute (i.e., invariant) difference between frames which are accelerated and those which are not in terms of whether they are associated with a geodesic trajectory, the *magnitude* of 3-acceleration is frame-dependent and that of 4-acceleration is covariant with invariant magnitude zero.

[20] It might seem that we've ignored the Kuhn's objection by moving from Newtonian to Galilean spacetime. The concepts of position and velocity are not the same in the two spacetimes, so this objection goes, one cannot simply move to Galilean spacetime and claim that we are still dealing with NM. This, however, mistakes the aim of the model-theoretic approach. The aim here is to compare the *structures* of the two theories, not their ontologies or ideologies. To that end, it is entirely appropriate to consider the models which display the structure of the theories most perspicuously. In the case of NM, Newtonian spacetime contains redundant structure in the form of a preferred inertial reference frame, and thus Galilean spacetime better captures the *structure* of NM.

3.4.2 The limiting process: from models in STR to models in NM

We are now in a position to consider the question of what is required for NM to be a limiting case of STR in the model-theoretic framework. First, at the limit where $v/c \rightarrow 0$ the models should be physically equivalent. Second, near the limit, models of the two theories should "approximate" one another by having similar structures.

In fact, both conditions are met when one considers the models of STR and NM. The distinctive feature of Minkowski spacetime is its light-cone structure, but at the limit where $v/c \rightarrow 0$ the light cones flatten to become the simultaneity planes characteristic of Galilean spacetime. In order to see this, we must further specify the limit as $c \rightarrow \infty$. In terms of a typical two-dimensional spacetime diagram, we begin with the future and past light cones at ±45° from the x-axis (which represents the speed of light = ${c_0}^{21}$), then as we increase c the angle decreases until the sides of both light cones are coincident with the x-axis at $c = \infty$. Thus, as we approach this limit we continuously move from the standard light-cone structure of Minkowski spacetime to the simultaneity plane structure of Galilean spacetime (see figure 3.1 below).

One considerable advantage to the model-theoretic approach is that one approaches the limit continuously. On the standard approach, the *form* of the dynamical equations changes discontinuously as one approaches the limit of $v/c \rightarrow 0$. On the model-theoretic view, Galilean spacetime is a natural end point of a continuous process; the light cones get progressively wider until they become entirely flat planes of simultaneity.

---

[21] I will use c_0 when necessary to indicate the speed of light in a vacuum: $c_0 = 299,792,458m/s$. Without the subscript, c will simply be shorthand for "the speed of light (in a vacuum)" and hence $c = c_0$ according to STR.

Figure 3.1: Light-cone structure as $c \longrightarrow \infty$. (I assume units where $c_0 = 1$.)

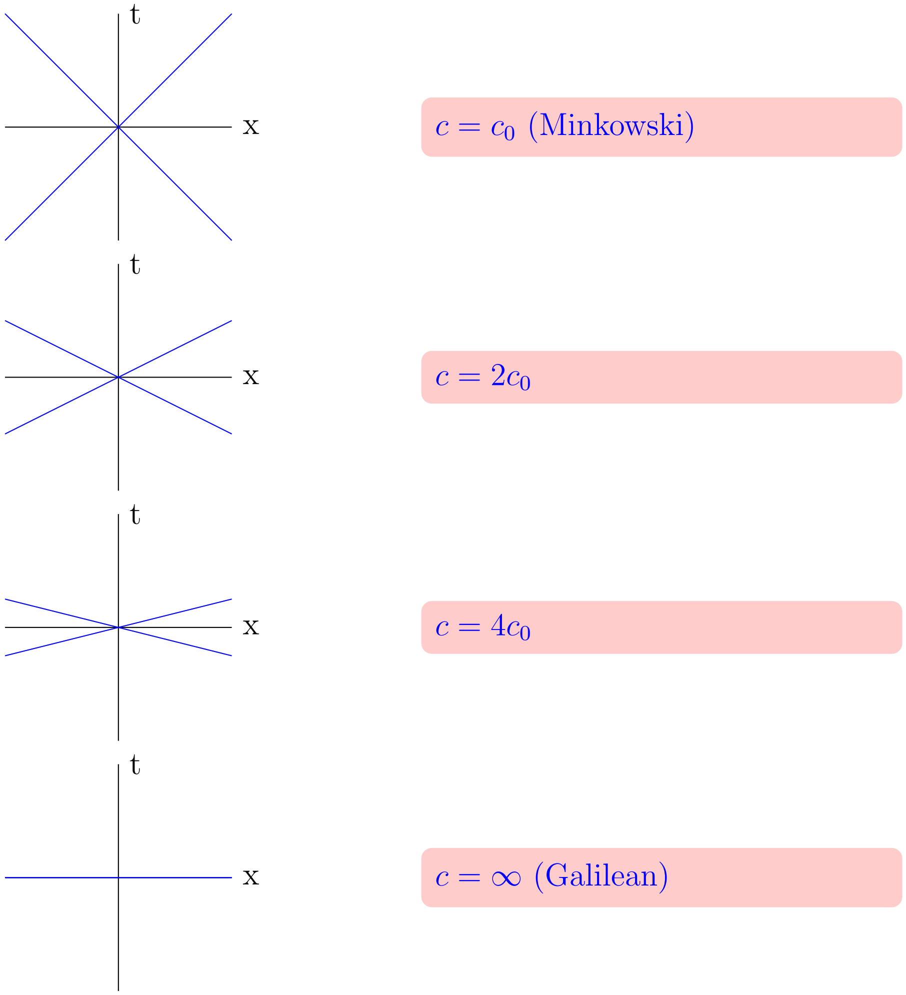

Let's now turn to two possible objections with the aim of clarifying the approach. First, in the preceding sections I have always characterized the limit as $v/c \longrightarrow 0$, but on the present approach it seems that we must further specify the limit as $c \longrightarrow \infty$. Yet, while this limit is formally adequate, it is conceptually distinct from taking $v \longrightarrow 0$. The intuitive motivation for viewing the limit as $v \longrightarrow 0$ is that, as the relevant speeds get slower, STR comes to more closely resemble NM. Viewing the speed of light as increasing to infinity seems less well motivated. While in many applications of NM speeds (in some relevant frame) are low ($v << c$) and distances are short ($x << ct$) there seems to be no corresponding commitment to the speed of light being very large. In NM, the speed of light is a mere empirical fact rather than a central parameter. Understanding the classical limit in terms of c tending to infinity seems to have little to do with the application of NM we are intending to justify.

Second, one alleged benefit of the model-theoretic approach is that the limit in which light-cone structure becomes the Galilean simultaneity plane structure is approached *continuously*. This is important because actual applications of NM occur *near* the limit, so we need a way of making sense of *approximate* physical equivalence between NM and STR near the limit. If models of STR continuously approach the limit of physical equivalence, then models near the limit will be approximately equivalent. One of the defects of the standard approach was that the form of the dynamical equations changes discontinuously once one departs from the limit. If the dynamical laws ever take the same form it is only *at* the limit, not anywhere away from it.

Of course, there is a sense in which the model-theoretic limiting procedure is continuous: the value of c increases continuously from c_0 to infinity, and the opening angle of the light cones changes from 90° to 180°. Yet, the limit represents a point of discontinuity. The only place in which there are *planes* of simultaneity is at the limit, and the boundaries of the past and future light cones overlap only at the limit. So does the model-theoretic approach really do any better than the standard approach at describing a continuous approach to the limit?

Now, to address the objections in turn let us return to the first issue. For the majority of the chapter we have taken the classical limit at which NM and STR converge to be $v/c \rightarrow 0$, but now it seems we must take the limit to be $c \rightarrow \infty$. But what does this limit—which concerns only the speed of light—have to do with applications of NM? The limit of $v \rightarrow 0$ makes sense in that many applications concern systems at low speeds and short distances, but few if any applications specify that c is very large. However, when we examine the velocity-based limit (v-limit) carefully, its superiority over the speed-of-light-based limit (c-limit) breaks down.

First, the notion of *velocity* is problematic in Galilean and Minkowski spacetimes, both of which lack a preferred inertial reference frame that is objectively *at rest*. Speed or velocity is always defined with respect to some reference frame, and thus, what counts as "low speeds" ($v << c$) depends on the frame one chooses; every physical system is moving slowly in some frames and extremely quickly in others. As Tim Maudlin has emphatically put the point,

> [I]n Relativity, just as in Galilean spacetime, there simply are no such speeds. There is no physical fact about how fast the earth is moving right now; it is no more correct to say that it is not traveling near the speed of light than it is traveling at 99 percent of the speed of light. To understand Relativity, we have to expunge all ideas of things having speeds. (2012, 68)

Perhaps we might resolve this issue by taking the surface of the Earth as the relevant (roughly) inertial reference frame we are seeking. After all, most of the applications of NM in which we take ourselves to be justified occur at low speeds *with respect to the surface of the Earth.* Yet, the concept of "low velocity relative to Earth" cannot inform our understanding of the classical limit if we take the model-theoretic perspective. On the model-theoretic approach, NM is limiting case of STR just in case the models of NM and STR are equivalent at the limit (and similar nearby). But only a fraction of these models will contain the Earth as we understand it. Thus, the "limiting case" claim cannot be evaluated if the limit itself makes reference to the Earth. When the model-theoretic approach was introduced above it was noted that it considers only *intrinsic* features of the global models of spacetime. Velocity (or speed) is an *extrinsic* feature in that it must be specified with respect to some reference frame (in this case, the frame associate with the Earth) and hence is unsuitable to the task of defining the classical limit from the model-theoretic perspective.

Second, suppose, in spite of the preceding, that one does find talk of velocity acceptable. Is the v-limit really more relevant to applications of NM than the c limit? After all, the way "low speed" is characterized is in terms of c: $v << c$. Yet, this is equivalent to $c >> v$, which is a claim about c as much as first inequality was about v. The cases in which we are justified in applying NM—those in which the speeds involved are very small (with respect to c)—are equally well described as cases in which the speed of light is very large (with respect to v).

Finally, we may also note that the c-limit has another benefit over the v-limit, namely, that it is applicable *at the limit.* Above it was claimed that actual applications do not occur *at* the limit where $v/c = 0$, and this is true if we take this to mean $v = 0$; we typically apply NM to moving systems. But, at the limit where

$c = \infty$ we may still consider applying NM, and would be successful in doing so. In a possible world in which $c = \infty$ the laws of NM could still apply to moving bodies. This suggests that the c-limit is actually *more* relevant at the limit than the v-limit.

The second objection above challenges the claim that the structural understanding of the limit is more *continuous* than the traditional alternative. Both the traditional, equation-based approach and the model-theoretic approach take NM and STR to be physically equivalent *at* the limit where $v/c \to 0$, but there is an important difference in the two approaches *near* the limit. As I emphasized above, the equation-based approach faces the following problem: applications of NM do not occur at the limit (modulo the last remark about the c-limit), and yet equations 3.11 and 3.12 differ from Newton's second law (equation 3.10) away from the limit. This makes it hard to see in what sense beyond the empirical equivalence NM "approximates" STR near the limit according to this view.

The model-theoretic approach, however, gives a clear meaning to such approximation. Near the limit ($c \approx \infty$) the light-cones of Minkowski spacetime will be very nearly the simultaneity planes of Galilean spacetime. Notice that this differs from the empirical equivalence of the equation-based approach in that it concerns what the global structure of the world according to these theories, not merely what they predict. In particular, the model-theoretic approach describes a series of worlds that begin with those closely resembling the special relativistic world and come to resemble the Newtonian world more and more closely. Creating such a collection of worlds is problematic in the equation-based approach because the form of the equations change when one reaches the Newtonian limit and the ontology and ideology of the worlds differ. This is what is meant by saying that the model-theoretic approach allows for the limit to be approached *continuously.*

3.5 Structural realism

So far, I've argued that the model-theoretic approach is superior to the traditional approach as a way for the scientific realist to endorse the claim that NM is a limiting case of STR. I have also claimed that this has implications for one's general stance toward scientific theories. In particular, I've claimed that the model-theoretic approach supports structural realism. So what is the connection between structural realism and the model-theoretic approach? The model-theoretic approach views NM as a limiting case of STR insofar as there is common structure shared by the models of both theories at the limit, and the models approximate each other nearby. This makes precise the structural realist's central thesis that theories should be taken to tell us only about the structure of the world.

The traditional realist, however, takes theories to tell us more. In particular, theories should provide us with an ontology of individual objects and an ideology of concepts. Yet, in this case such an approach leads to the problems associated with the traditional approach discussed above. The issue underlying these problems is how NM can be viewed as a limiting case of STR when the two theories disagree about what objects there are and which concepts should be used to describe them. Given the traditional realist framework, it's hard to see how this challenge can be overcome. Unlike structures, the is no obvious sense in which individual concepts or objects may approximate one another in the way two structures may.

Another way of articulating the connection is worth revisiting. Above I claimed that (1) for theory T to be a limiting case of theory T' is for T and T' to be *equivalent* at the limit, (2) two theories are equivalent (at the limit) just in case their models have the same structure (i.e., are isomorphic.) The second of these claims has been challenged by Hans Halvorson (2012).

In particular, Halvorson argues against the *model isomorphism criterion for theoretical equivalence*: "If theories T and T' are equivalent then each model of T is isomorphic to a model of T'" (2012, 187). He proceeds by considering three possible meanings of "isomorphism" and showing that each of them either equates inequivalent theories or fails to equate equivalent ones. Equinumerosity, for example, is too weak a notion of *isomorphism* to capture theoretical equivalence; it recommends the identification of theories that do not seem to be equivalent.[22]

My argument above, however, argues that taking model isomorphism to be the criterion for equivalence—rather than the more traditional understanding—provides the scientific realist with a way to justify the claim that NM is a limiting case of STR. As figure 3.1 shows, the distinctive structural feature of Minkowski spacetime, its light cone structure, becomes the simultaneity plane structure of Galilean spacetime at the limit, and hence there is a clear sense in which the models of NM and STR are isomorphic at the limit. But if Halvorson is correct, model isomorphism should not be taken as grounds for theoretical equivalence.

Before responding to Halvorson's argument directly, it is worth noting that my claim in this chapter is weaker than his target, namely, the full theoretical equivalence of theories. The structural realist does not contend that STR and NM are equivalent theories, but rather that they are physically equivalent—i.e., they say the

---

[22] See (2012, 191) for Halvorson's example of a pair of theories which do not seem to be equivalent but whose models are equinumerous.

same thing about the world—only *at the limit*. I neither contend that *all* of the two theories' models are isomorphic—only a subset of them are—nor do I assert theoretical equivalence full-stop. So, there is a sense in which the *model isomorphism criterion* simply does not apply here, and hence Halvorson's criticism of it leaves my argument unscathed.

Yet, if we construct two related theories, STR-L and NM-L, which are the usual theories with their domain of application restricted to at the limit of $v/c \to 0$, the problem can be reintroduced. In this case, the structural realist would claim that the two theories are fully equivalent in virtue of the fact that *all* of the models of one theory are isomorphic to models in the other.[23] Hence, Halvorson's argument is a concern.

To establish that the model isomorphism criterion is inadequate, Halvorson relies on the notion of *definitional equivalence* as a sufficient condition for theoretical equivalence. Yet, definitional equivalence is a syntactic notion that applies only to theories formulated in first-order logic.[24] The model-theoretic approach advocated here rejects expressing theories in first-order logic and emphasizes the role of (spacetime) models. I have argued that STR-L and NM-L are *inequivalent* on the traditional equation-based approach and, indeed, they are definitionally inequivalent as well. Halvorson is correct to note that different conceptions of equivalence will disagree about which theories are equivalent, but this is unsurprising. It is

---

[23] Kuhn (1962/1996, 92–93) rejects viewing NM as a theory limited to the domain of application in which it is successful. Kuhn argues that taking such a construction to be NM is *ad hoc* and opposed to the spirit of science. But what Kuhn is objecting to is a defense of the claim that NM is a limiting case of STR that relies on understanding NM as a domain-limited theory along the lines of NM-L. I agree with Kuhn that this is not the way to understand being a limiting case, instead I claim that NM is a limiting case of STR insofar as STR-L and NM-L are equivalent.

[24] For a precise characterization of definitional equivalence see Halvorson (2012, 191).

a substantial thesis that STR-L and NM-L are equivalent, and one that a realist must argue for. But the structural realist is happy with the fact that these theories are equivalent as it allows one to justify the use of NM on realist grounds. Thus, in this case, there seems to be no reason to prefer definitional equivalence to the model isomorphism criterion. Whether the structural realist can accommodate the implications of the model isomorphism criterion when it is applied to other cases remains to be seen.

3.6 Conclusion

The question with which we began this case study was: how can the use of NM—a superseded theory—be justified for the scientific realist? The traditional account of the standard reply that NM is a limiting case of STR does not seem to withstand scrutiny. The model-theoretic approach allows us to see that the models of both theories share a common structure at the limit. This approach is exactly the one recommended by structural realism: two theories are equivalent when their models are isomorphic.

CHAPTER 4

Ontic Structural Realism and the Problem of Objects

4.1 Structural Realism

Structural realism (SR) is the view that we should be realists with respect to the structural content of scientific theories. That is, when these theories are successful we are justified in regarding them as providing an approximately faithful description of the structure of the world. This differs from traditional realism in that any objects and properties—to the extent that these are thought of as something over and above structure—posited by these theories are not among the structural realist's commitments. Of course, how one understands "structure" will play a crucial role in developing SR, and I will have something to say about this below (sections 4.2 and 4.3). If, for the time being, we take a structure to be a network of relations, SR is the claim that science can tell us only about *relations*, not about the *relata* which enter into those relations. Thus, SR attempts to merge realism's commitment to correspondence between our theories (or models) and the world while rejecting the commitment to the entities posited by these theories.

Once we begin to consider SR in detail, the view quickly subdivides further. Perhaps the most central distinction is due to French and Ladyman (2003) who ask: why is it that we can only know about the structure of reality? In particular, what is the nature of this restriction to structure? Epistemic structural realism (ESR) answers that the restriction to structure reflects our limitations as knowers. Our theories—being artifacts of human creation—can only get at structure, not the full picture of fundamental reality. Typically, ESR has been taken to regard the world as being made up of a fundamental ontology of objects and properties much like the traditional realist view. The difference is that fundamental reality is deemed unknowable by ESR, and hence we are stuck with incomplete, structural knowledge of the world. For this reason (i.e., that it posits an unknowable reality) ESR has sometimes been associated with a Kantian or NeoKantian orientation.

In contrast to ESR, ontic structural realism (OSR) holds that the reason we can only know the structure of reality is because *that is all there is to know.* There is no unknowable realm of objects and properties that underlies the structure revealed by our best theories, rather, there is only structure at the fundamental level of reality. OSR strikes many as a radical view; it seems to require a drastic revision of our current understanding of metaphysics and the nature of structure. I will have more to say on these issues later. At this point, OSR is perhaps best captured by the (admittedly unsatisfying) slogan “structure is all there is.”

Recently, other parts of the logical landscape have been charted. ESR, for example, need not be committed to the positive existence of objects and properties of the kind imagined by traditional realism. It might instead be developed as maintaining agnosticism towards such objects (Morganti 2004) in the same sense constructive empiricists like van Fraassen are agnostic toward unobservable entities and structures. This view occupies a middle ground between OSR (atheism about objects) and traditional ESR (theism about objects), but is still a form of ESR in that the restriction to structural knowledge is epistemic. In sum, agnostic ESR is unwilling to accept OSR’s claim that our knowledge of the world is (in principle) complete, but doesn't deny it either. It is an open question for agnostic ESR whether there is an underlying ontology of objects and properties giving rise to structure (traditional ESR) or not (OSR).

Another interesting region of logical space that has been explored recently concerns weaker forms of OSR. OSR has traditionally been developed as a form of eliminativism about objects according to which there are no objects in the fundamental ontology of our world. Recently, several authors have advocated for weaker, non-eliminativist versions of OSR that allow for objects which are *dependent on structure* in some sense. It's not clear just how far one can weaken the restrictions of OSR before slipping back into a traditional realist position. For example, Anjan Chakravartty characterizes OSR in a way that allows for non-eliminative varieties:

> At the heart of OSR there is one thing on which all proponents of different stripes can agree: an emphasis on structures or structural relations at the expense of things putatively standing in those relations—their relata (Chakravartty 2012, 3).

If all OSR requires is an *emphasis* on structure over that which is structured, certainly eliminativism is not the only game in town. Ultimately, I will argue that the structural realist is best served by a non-eliminativist version of OSR. My argument will proceed as follows. First, in section 4.2 I will argue that the standard argument against ESR fails and offer a replacement. Next, in section 4.3 I will discuss the Problem of Objects facing OSR: how can there be structure without objects? My position here will be that the Problem of Objects is not easily dismissed. In section 4.4 I will turn to non-eliminative OSR and propose a version of the view in which objects can be recovered as places in many structures.

4.2 The Case Against ESR

Well-known early forms of structural realism were versions of ESR. Russell, for example, was an early advocate of the view. According to his structuralist view (Russell 1911, 1927/1992) we only have direct epistemic access to the contents of experience and all else must be derived from this experience. To the extent that we can know about the external world (as opposed to the internal world of experience), this knowledge is restricted to the structure of that world. Maxwell (1970a,b, 1972) developed Russell's view further by casting it in terms of Ramsey-sentences. This allowed for the precise articulation of the claim that our knowledge of the external world is purely structural.

To arrive at the Ramsey-sentence of some theory **T**, we begin by formalizing **T** in first-order predicate logic. Then we divide the terms that occur in **T** into logical terms, observational terms and theoretical terms, where observational terms are those that apply to observable things and any remaining nonlogical terms are theoretical. Finally, we replace the theoretical terms of **T** with variables, then quantify over them. Suppose, for example (from Maxwell 1972), our theory is given by the following expression in first-order logic:

$\forall x((Ax \wedge Dx) \rightarrow (\exists y Cy))$,

where A and D are theoretical predicates that mean that x is a radium atom and x radioactively decays, respectively, and C is an observational predicate that means a click is registered on a Geiger counter. The Ramsey-sentence for **T** in this case would be the following second-order logical expression:

$$\exists X \exists Y \forall x ((Xx \wedge Yx) \rightarrow (\exists y Cy)).$$

Ramsey-sentences enable one to express a theory without the use of theoretical terms. This provides a useful explication of Russell and Maxwell's commitment to "concept empiricism," the claim that the only meaningful concepts are those arrived at empirically. By "Ramsifying away" problematic theoretical predicates, we arrive at a version of **T** that is expressed in meaningful terms. Moreover, Ramsey-sentences appear well-suited to the needs of ESR. They aim to go beyond the empirical content of the theory, but only by specifying that there are *some* predicates that have the same structure as the specific theoretical terms being replaced. The view is an epistemic version of SR because it doesn't attempt to do away with the commitment to the unobservable entities posited by our theories. Such entities are not referred to by theoretical terms, but are nevertheless required to make the relevant Ramsey-sentence true. We are no longer committed to the radium atom as such, for example, but we are committed to *that which has the properties X and Y which entail the Geiger counter click*.

There is, however, a well-known objection to this version of ESR. First raised by Newman (1928) in response to Russell's view, and later developed by Demopoulos and Friedman (1985) and Ketland (2004), this objection purports to show that the kind of "structural knowledge" offered by Ramsey-sentences is trivial. The essence of Newman's objection is that any collection of things can be organized so as to have any arbitrary structure. For example, if we claim that there are three objects $[a, b, c]$ that enter into the relations $[R_1(ab), R_1(ac), R_1(bc)]$, this tells us no more than that there are three objects. Elementary set-theory tells us that *any* three objects will instantiate a relation R_i as well as any other relations requiring no more than three objects.[1] Of course, there are also observational (or "internal") predicates on the Russell-Maxwell view, but these don't help avoid the charge of triviality regarding the theoretical ("external") content of theories. This leads Newman to conclude that if we begin with a theory that is empirically adequate then the most it can tell us about the unobservable/external realm is its cardinality.

Newman's objection was directly applied to the Ramsey-sentence approach to ESR by Demopoulos and Friedman (1985) and given a formal treatment by Ketland (2004). The upshot of these versions of Newman's objection is that the Ramsey-sentence of a theory **T** is true *iff* it is empirically adequate and **T**-cardinality correct. Thus, it seems that Ramsey-sentence ESR of the sort envisioned by Russell and Maxwell leaves us with nothing more than antirealism combined with a cardinality claim.[2] Clearly this is unacceptable as a form of realism, structural or otherwise. Getting the cardinality of the world right is surely insufficient as an explanation of the predictive success of an empirically correct theory, which is one of the central motivations for scientific realism. Newman's objection constitutes a very serious objection to ESR. The question now is how the objection can be met or avoided.

Initially it may seem as though Newman's objection is unique to the Ramsey-sentence approach and as such can be avoided by giving up on Ramsey-sentences. This line of reasoning is particularly attractive in light of more general problems with Ramsey-sentences. In order to implement the Ramsey-sentence procedure one

---

[1] We are assuming that R_i is defined extensionally, for example, $R_1 = [\langle ab\rangle, \langle ac\rangle, \langle bc\rangle]$. Clearly any relation R_i we define in this way will be satisfied by any three objects.

[2] Even the cardinality claim doesn't hold generally. If T is non-categorical, then it will have models of varying cardinality. The truth of its Ramsey sentence T^* would tell us only that the world has the same cardinality as *some model* of T (Frigg and Votsis 2011, 249–250).

must view theories as linguistic entities that can (in principle) be formulated in first-order predicate logic. Such a view is associated with the "syntactic" view of theories in contrast to the "semantic" view according to which theories are associated with a collection of models. Many contemporary philosophers of science have abandoned the syntactic view for independent reasons and with it the possibility of Ramsey-sentence ESR. If the scope of the objection is limited to Ramsey-sentence formulations, then it simply does not arise for those epistemic structural realists (ESRists) that prefer the semantic view.

French and Ladyman (2003) suggest such a response to Newman's objection (although from the perspective of OSR), but on its own this seems insufficient. After all, Newman's original formulation of the objection wasn't cast in terms of Ramsey-sentences, but rather concerns any attempt to make claims about the structure of the unobservable world. Moreover, as Ainsworth (2009, 17) shows, we can easily adapt Ketland's proof (cast in terms of Ramsey-sentences) to the semantic view. Thus, while we may wish to jettison the Ramsey-sentence approach on independent grounds, doing so will not avoid Newman's objection. In fact, the generality of Newman's objection threatens *any* attempt to specify the structure of the unobservable in observable terms.

If the objection cannot be easily dodged, perhaps we should meet it head on. Zahar (2004); Cruse (2005); Melia and Saatsi (2006) each argue against the Newman objection in defense of ESR. Cruse, for example, argues that if one adopts a weak version of the observational/theoretical (O/T) distinction, the objection can be avoided. Ketland, in his proof, assumes what Cruse calls the strong O/T distinction: observational predicates apply only to observables and theoretical predicates apply only to unobservables. As Putnam (1966) observes, such a division is implausible. Red blood cells are unobservable, but clearly “red” is an observational predicate. Thus, Cruse suggests that we adopt a weak O/T distinction, according to which observational predicates can apply to unobservables and theoretical terms can apply to observables.

Sure enough, if we replace the strong O/T distinction with the weak O/T distinction, Ketland’s proof fails. It follows that we can know more about the unobservable realm than its cardinality. But, as Ainsworth (2009) notes, this knowledge seems decidedly non-structural. Consider, for example, the predicate “larger.” This is a term that can be applied to both observables (“watermelons are larger than apples”) and unobservables (“protons are larger than electrons”), but notice that if we specify the predicate when applying it to unobservables we go beyond merely giving the structure of the situation. We have gone beyond merely considering the abstract structure of an uninterpreted theory to attributing properties and relations to the world. Presumably, this is what SR is trying to avoid.

Perhaps more troubling, a version of Newman’s objection can be run on the weak version of the O/T distinction as well. Rather than consider “observational” and “theoretical” terms, we now consider predicates whose meanings are known (i.e., weakly observational terms) and those that are not (i.e. strongly theoretical terms). The upshot will be that any theory $\mathbf{T_k}$ which is empirically adequate and correct about what the known predicates say about the unobservable world is guaranteed to be true. That is, if we have a theory that is empirically adequate, and is correct in the properties attributed to the unobservable realm by mixed predicates, its truth depends only on the cardinality of the unobservable realm. Thus, the ESRist is faced with a dilemma: either we have knowledge of the unobservable realm that is non-structural or we can only know its cardinality. The former gives us knowledge of the unobservable realm, but that knowledge is not of its structure. The latter also tells us something about the unobservable realm (how many objects it contains), but not what the realist had hoped.

So, what are we to make of ESR in light of Newman's objection? We have seen that the problem lies not with the Ramsey-setence approach itself nor can it be avoided by weakening the O/T distinction. The Newman objection relies on two crucial assumptions: (1) predicates are defined extensionally and, (2) structure is to be captured in observational terms.[3] Russell and Maxwell were committed to (1) and (2) by their concept empiricism, but contemporary ESRists need not be so committed. Structural realists, epistemic or otherwise, allow that we can know the relations between objects whether they are observable or not. Often, SR places great significance in the mathematical structures involved in predictively successful scientific theories. These mathematical structures are taken to represent real relations in the world. These relations need not be observational nor extensionally defined, and knowledge of them is not "non-structural" in that these relations *constitute* the structure of the world. Thus, it is the assumptions (1) and (2), and the concept empiricism that motivates them, that an ESRist must deny to avoid Newman's objection. If there is a tenable notion of structure that can do this, ESR is not undermined by the objection; if there is not, ESR is in trouble, but so too is

---

[3] We have just seen that a version of the Newman objection can be raised even if some theoretical terms are included in the base vocabulary. So, although (2) was true for the best known versions of the objection, it isn't required. What *is* required is a distinction between knowable and unknowable predicates (or the corresponding distinction in the semantic view).

OSR.

Newman's objection, then, places important constraints on the notion of structure used by *any* form of SR, epistemic or ontic. Another important upshot of the previous discussion is that Newman's objection cannot be used by ontic structural realists as a means of ruling out ESR. The objection casts a significant shadow over the version of ESR envisioned by Russell and Maxwell, but its starting point is not so much ESR as concept empiricism. To the extent that ontic structural realists (OSRists) can avoid the objection by casting structure in terms that reject the assumptions above, an epistemic structuralist can do the same. This is unsurprising. Any version of ESR will have two components: (1) an account of the structural content of a theory and (2) an "epistemic" stance toward that structural knowledge (i.e., that structural knowledge reflects our inability as knowers to get the full metaphysical story). Newman's objection places severe constraints on the form of (1), but it doesn't directly concern (2). Any version of SR must address (1) and hence must deal with Newman's objection, so appealing to it as a means of blocking ESR is a strategy that is bound to backfire for the OSRist.

4.2.1 A General Argument Against ESR

When Ladyman and Ross directly address ESR (2007, 124–128) they offer two different lines of criticism: one which centers on Newman's objection and another specifically targeting Maxwell-style Ramsey-sentence ESR. The latter shows, at most, that this particular conception of ESR is problematic, a conclusion that the contemporary ESRist can readily accept. As we have seen, Newman's objection also fails to tell against ESR any more than OSR. Thus, these considerations can only impact part (1) above—how structure is conceived—and hence do not constitute a general argument against ESR. Elsewhere in the same book, however, one finds the suggestion of a different line of reasoning.

> ...such a gap between epistemology and metaphysics is unacceptable. Given that there is no a priori way of demonstrating that the world must be composed of individuals with intrinsic natures, and given that our best physics puts severe pressure on such a view, the PNC [a naturalistic principle that requires all metaphysical claims be in the service of science] dictates that we reject the idea altogether. (2007, 154)

As mentioned above, there are two forms of ESR, which I will call "Kantian" ESR and "agnostic" ESR, respectively.[4] The Kantian variety claims that there are objects ("individuals with intrinsic natures" in Ladyman's terms) which instantiate the relational structure described by physics. The agnostic version aims to stay neutral on the question of whether structure is ontologically basic or not; it takes the attitude that perhaps Kantian ESR is right, perhaps OSR is right, we are not justified in taking either stance.

Returning to the passage quoted above, Ladyman and Ross are suggesting that Kantian ESR is unacceptable on naturalistic grounds because it proposes a gap between epistemology and metaphysics. What there is should not outstrip what we can in principle know about. Does naturalism really preclude such a gap? And if so, does this undermine Kantian ESR?

---

[4] The only sense in which the former bears any resemblance to Kant is in that the world "as it is in itself" is unknowable. This two-tiered metaphysical picture conjures Kant's distinction between the noumenal and phenomenal realms. I do not intend to conjecture on whether Kant was a structural realist, much less a Kantian ESRist.

It must be noted that the principle evoked by Ladyman and Ross, the PNC, is a much stronger thesis than mere naturalism. It holds that *any metaphysical claim whatsoever* must show “how two or more specific scientific hypotheses, at least one of which is drawn from fundamental physics, jointly explain more than the sum of what is explained by the two hypotheses taken separately” (2007, 50). But, in the interest of making this argument against ESR broadly appealing, we should ask whether garden variety naturalism can support it.

Unfortunately, there is no consensus among philosophers on just what naturalism comes to. Indeed, at times it seems there are as many versions of naturalism as there are advocates for it. One common distinction is between metaphysical and methodological varieties of naturalism. The former is a thesis about ontology. It says that the world is natural—it contains nothing supernatural—where “natural” has something to do with what science tells us there is. Methodological naturalism concerns the practice of philosophy. It holds that philosophy and science are continuous in method. The same methods used by science—things like using empirical evidence to support theories—are to be adopted by philosophy. One (somewhat idiosyncratic) statement comes from Quine, who defines naturalism as,

> the recognition that it is within science itself, and not in some prior philosophy, that reality is to be identified and described. (1981, 21)

The contrast with prior philosophy is important. Methodological naturalism is particularly hostile to *a priori* methods in philosophy. This bears on the status of Kantian ESR because, according to that view, there exist unknowable entities. Given that these entities are unknowable by science (which can only furnish structural knowledge), there must be some *prior* reason for positing them. Two suggestions for such a reason are: (1) intuitions about the constitution of reality and, (2) the conceptual necessity of objects to realize structures. If these count as "prior philosophy" in the relevant sense, then it seems we can rule them out on the basis of a methodological naturalism in the spirit of Quine's. This suggests a general argument against Kantian ESR with the following form: *The existence of unknowable objects required by Kantian ESR is unacceptable on naturalistic grounds because it requires appeal to methods beyond those of science.*

But perhaps this is too quick a dismissal of Kantian ESR. Quine's naturalism is quite idiosyncratic and, at any rate, it's hard to say which conceptual issues can be dismissed as "prior philosophy." Thus, the two motivations for Kantian ESR deserve closer examination. The first motivation draws on our metaphysical intuitions to recommend Kantian ESR. Intuitions about the nature of reality have recently been drawing criticism from those with broadly naturalistic orientations. Experimental philosophy has called into question the role of intuitions in philosophy in general, and the growing literature on "metametaphysics" has targeted metaphysics specifically.[5] Another line of thought takes its cue from human evolution. While our metaphysical intuitions have served us well in navigating the macroscopic world, there is little reason to think that the microscopic world should conform to them.[6] Thus, the fact that we tend to think of the world intuitively in terms of individual objects provides little warrant for recognizing such objects as part of our fundamental ontology.

---

[5] See Chalmers et al. (2009) for a sampling of work on metametaphyiscs.

[6] David Wallace puts this point more colorfully. "Our intuitions about what is "reasonable" or "imaginable" were designed to aid our ancestors on the savannahs of Africa, and the Universe is not obliged to conform to them" Wallace (2010, 69).

The second motivation is more difficult to dismiss (and, accordingly, will receive a more thorough treatment below). If structures require objects by conceptual necessity, then it might seem that Kantian ESR is the only viable form of SR. Moreover, given ordinary conceptions of structure (e.g., networks of relations, set-theoretic structures), it *does* seem as though objects are required. Perhaps even this motivation can be written off as "non-naturalistic," but this seems unlikely. It is not an empirical motivation, but neither is it unknown to science. Luminiferous ether, for example, was motivated in part by the felt conceptual need for a wave to have a medium to wave *through*. One option would be to extend the evolutionary line as follows: just because our concept of structure requires objects does not mean that actual structures in the world do. Our concept was the result of the evolutionary pressures of our macro-environment and hence does not apply outside of that context. Along these lines, Ladyman and Ross criticize the apparent necessary connection between structure and objects as merely an artifact of philosophical theorizing.[7] Such replies are unappealing for their vacuousness. If one claims that the world contains structures unlike any we can imagine (in that they don't depend on objects), then until such a notion of structure can be characterized *in some way* we loose any grip of what is meant by "structure" in this context.

Of course, ontic structural realists claim to have found ways of characterizing structure that avoids the necessity of objects (at least objects *qua* "individuals with intrinsic natures"). If those attempts are successful, the second motivation for Kantian ESR disappears. Thus, *if* we can debunk the apparent need for objects to

---

[7] "We think it better to attempt to develop the metaphysics presented in this book than to continue to use off-the-shelf metaphysical categories inherited from the ancient Greeks that are simply not appropriate for contemporary science" (2007, 125).

underlie structure, there is very little left to recommend Kantian ESR over OSR. Beginning in the next section I will argue that OSR can in fact deal with the apparent need for objects underlying structures. For now, let's assume that this issue can be resolved with the aim of presenting the naturalistic argument against ESR.

Table 4.1: A general argument against ESR

1. **Structural realism**: science can only tell us about the structure of the world.
2. **Naturalism**: science is our best/only guide to reality.
3. Therefore, **OSR**: all there *is* is structure.

This argument rules out Kantian ESR by eliminating the gap between epistemology and metaphysics exploited by the position. Unlike appeals to Newman's objection or attacks on Ramsey-sentences, this argument precludes the very possibility of a view according to which there are objects and properties of which science is forever ignorant.

Yet, even if the above argument succeeds, perhaps the right position to take toward the putative objects and properties in question is one of agnosticism rather than atheism. After all, naturalism is not so stringent as to require that unknowable objects are *impossible*. This line of reasoning is attractive, but misleading. Even the OSRist can grant that they may be wrong—individual objects very well might exist—by adopting a version of fallibilism. What agnostic ESR really needs is the stronger claim that we are not justified in saying *whether or not* there are unknowable objects underlying structure. Here the ontic structuralist can appeal to the naturalistic argument above: there is no naturalistically acceptable motivation for

Kantian ESR, so we are justified in rejecting that view as a description of the world. The motivations for Kantian ESR are ruled out by taking the naturalistic perspective and hence there is no good reason to withhold making a judgement about the status of Kantian ESR either.

Thus, it seems ESR in either form (Kantian or agnostic) is unacceptable on naturalistic grounds. There is no good reason to remain neutral on whether structure is ontologically basic so long as the conceptual link between structure and objects can be broken. The onus thus falls on OSR. If it can be made coherent, it offers the best prospect for a naturalistically-motivated form of SR.

4.3 OSR and the Object Problem

It is with this task in mind that we now turn to OSR. As with ESR, OSR also comes in several forms of differing strengths. The most extreme view is eliminativist OSR. Certain remarks from Ladyman and French suggest such a view.

> There are no epistemically inaccessible objects laying behind the structures which we can know (French and Ladyman 2003, 39).

> a first approximation to our metaphysics is: ‘There are no things. Structure is all there is’ (Ladyman and Ross 2007, 143).

A few sentences after the latter quotation, Ladyman and Ross describe their position as committed to “eliminativism about self-subsistent individuals, the view that relational structure is ontologically fundamental” (*ibid.*). Eliminativist OSR rejects the existence of objects and instead posits a fundamental ontology of “pure structure." The structural content of theories in fundamental physics provides a complete description of the world. Immediately, however, a problem arises.

On most standard understandings of *structure*, the existence of objects is required to realize a structure. Set-theoretic structure consists of a domain of objects together with relations defined over them. The informal notion of a "relational structure" may not include objects in its specification (it's just a "pattern of relations"), but there must be some relata for the relations to hold between, and hence objects are still required. This is what I'll call the object problem for OSR. *Structures require objects, so OSR's claim that "all there is is structure" must be false.* In other words, structure, to the exclusion of objects, cannot replace objects as the fundamental ontology of the world because structure requires that there are objects.

Clearly the object problem is a real threat to OSR. If sustained, it renders OSR untenable, and thus forces the structural realist into an epistemic stance toward structure. The problem is also a pervasive one. As we'll see, it recurs in many areas when trying to develop a coherent version of OSR. Yet despite this, we might think there are some easy solutions available to the advocate of eliminative OSR.

One reply to the object problem is to simply deny that there is a problem. This view denies there is any difficulty in making sense of structure without objects (or relations without relata), so there is no problem in saying the world is purely structural. We might develop this reply as follows. Our theories attribute a structure to the world via the apparatus of set-theoretic structures, which clearly require objects, but the fact that structure takes this form is merely an artifact of how we choose to represent things. In our representation there must be objects, but in the world there need not be.[8] The *real* structure of the world (which is objectless) cannot be directly represented by us (we lack the requisite conceptual powers), but we can represent it (perhaps approximately) with set-theoretic structure. Objects, then, are merely heuristics that we use to get at structure in our limited representational scheme, and should not be taken to be real constituents of the world. Again we might appeal to the fact that we evolved to think in terms appropriate to the macro-world, not fundamental reality, so it shouldn't be surprising if we cannot represent fundamental reality directly.

The chief worry with such an approach is its lack of specificity. There is very little the eliminative ontic structuralist can say about what fundamental reality is like on this view. We know that "all there is is structure" but there seems to be no way to move beyond this slogan if our representational means are all inadequate. To some extent, this puts eliminative OSR in a similar position to Kantian ESR: on both views there is very little we can know about fundamental reality. Eliminative OSR was supposed to accept the structural content of scientific theories as complete descriptions of reality, but on this reading even the structure of our theories leads to reifying objects (i.e., those which instantiate relations) whose existence the structural realist wishes to deny. Thus, eliminative OSR goes *beyond* accepting theoretical structure as complete and claims that it has surplus content.

Perhaps even this is giving the reply too much credit. The problem was that structure requires objects, and the reply is that only our concept of structure requires

---

[8] Suppose we have a structure $S = [D, R]$. This mathematical model contains objects and relations, but they are *mathematical* objects used to *represent* the world. It is a further question whether there is anything in the *world* that corresponds to the objects and relations in the model. On the proposal presently under consideration, the relations in this model correspond to relations in the world, but there is no similar correspondence between the *objects* in the model.

objects, but the real structure of the world does not. There is a clear sense in which this is not an adequate reply to the question. If the world has some structure that doesn't require objects then it is eliminative OSR's burden to say just what that structure is. Otherwise, the claim is vacuous. Ladyman and Ross draw an analogy with universals in an effort to show that we can make sense of the notion of relations without relata, but this doesn't obviously help. Universals are thought of either as abstract objects (on the "Platonic" view) or features of concrete reality (on the "Aristotelean" view). On the Platonic reading, there is an abstract relation that exists independently of all the concrete relata it applies to, but there still must be *abstract* relata that enter into the relation. The universal corresponding to the two-place relation R_{ab}, for example, is an abstract object on this view, but it requires the abstract objects a, b as further posits. On the Aristotelean reading, the universal is identical to all of its concrete instances which all are between concrete relata. So, in the case of universals, there must be relata; the only question is whether they are abstract or concrete.[9]

4.3.1 Bain's Category-Theoretic OSR

There are some more sophisticated ways of meeting the object problem facing eliminative OSR. Jonathan Bain (2011), for example, argues that the object problem does not arise if we adopt a category-theoretic understanding of structure instead of a set-theoretic one. A category **C** consists of a pair of classes: the class $(A, B, C, ...)$ of **C**-objects and the class $(f, g, h, ...)$ of **C**-morphisms between **C**-objects (e.g.,

---

[9] Universals can be thought of as the referents of predicates and hence can pick out properties or relations. I have focused on relations as this is more relevant to the discussion of structure that is our focus here.

$f : A \rightarrow B$). It is important to distinguish **C**-objects from the members of a set and, indeed, from ordinary physical objects. In the category **Set**, sets are the **C**-objects and functions are the morphisms between them. On the set-theoretic framework, members of sets represent physical objects, but in the category-theoretic framework the members of sets correspond to elements of a **C**-object. Thus, categories resemble entire theories, while **C**-objects resemble their models. This leads Bain to claim that,

> for any given structured set, there is a category in which the objects are that type of structured set and the morphisms are functions that preserve the structure of the set (see Lawvere and Shanuel (1997) for elementary examples). This suggests that the intuitions of the ontic structural realist may be preserved by defining "structure" in this context to be "object in a category." (2011, 3)

Bain's suggestion is that eliminativist OSR should take structures to be **C**-objects such as the structured sets mentioned above and defined in terms of isomorphism. Now, so long as we are talking about the category **Set**, "objects" remain as elements of the sets, now viewed as **C**-objects, yet Bain emphasizes that in *other* categories, such elements can be eliminated. There are two steps in this elimination process. First, elements of **C**-objects must be characterized precisely in category-theoretic terms. This is done by identifying an element of a **C**-object with a *terminal object* in **C**, where the latter in defined as a **C**-object 1 such that for every other **C**-object A there is one and only one **C**-morphism $A \rightarrow 1$. In the case of **Set**, the terminal objects are the singleton sets and an element a of a set A can be identified with the morphism $\{a\} \to A$. Crucially, terminal objects can be characterized entirely in terms of $\mathbf{C}$-morphisms and hence should not be thought of as *elements* of structures in the sense of "internal constituents" that eliminative OSR seeks to avoid commitment to.[10]

Thus, Bain's category-theoretic framework seemingly allows the eliminative OSRist to define structures ($\mathbf{C}$-objects) in such a way that eliminates reference to individual objects (i.e., relata or members of a set). Category theory offers a very interesting framework from which to address questions about the structure of physical theories, and it does seem that it can do so without appeal to the sort of formal elements that are typically taken to represent physical objects. Yet, as Lam and Wüthrich (2013) emphasize, rather than providing a framework for eliminative OSR, Bain's approach is something altogether different from OSR as traditionally understood. The starting point of OSR is the claim that *structure* provides a complete description of reality, where "structure" means a network of concrete relations between physical objects. The goal was to keep the relations that comprise the structure and eliminate the relata, but by moving to a category-theoretic framework we seem to have changed what we mean by "structure" such that the original relations are gone. Instead of talking about physical relations between objects (e.g., "having opposite spin from"[11]) we are now talking about $\mathbf{C}$-morphisms between $\mathbf{C}$-objects. In order to recover those relations that the OSRist takes to be genuine features of the world, one would have to look beyond the resources available in category theory.

---

[10] See Bain (2011, §2.1), Lam and Wüthrich (2013, 9).

[11] This is a favorite example of the sort of non-supervenient relation that is intended to motivate OSR. See, for example, Ladyman and Ross (2007, 137), Saunders (2003, 2006a,b).

Category-theory undoubtably offers an important framework for viewing physical theories but it does not solve the object problem for OSR. While it allows one to eliminate direct reference to physical objects, it does so only at the cost of ascending to a level of abstraction without obvious metaphysical implications. Wedding eliminative OSR and category theory in the way Bain suggests, we seem to get the result that the world is (or perhaps is represented by) a **C**-object. But unless we fall back onto a set-theoretic understanding, it's entirely unclear what sort of metaphysical picture this gives us. In sum, the characterization of *structure* provided by Bain's category-theoretic OSR is insufficient to the task of describing the physical world in a way that saves the concrete physical relations emphasized by SR.

4.3.2 Monism and OSR

Another potential solution to the object problem should be considered. The crux of the problem is that ordinary ways of thinking about structure—in terms of objects, properties and relations—seem incompatible with the eliminativist agenda of OSR. Monism, however, has the potential to square the two by taking structure to be a feature (or collection of features) of the world. Of course, monism cannot serve as the basis for a truly *eliminativist* OSR given its recognition of one object, yet, if this is the only concession the eliminativist needs to make, much of the view remains intact. The objects that eliminativist OSR wishes to dispense with are chiefly those taken to be *fundamental building blocks* of the world, for example, elementary particles and points of spacetime. As far as such physical objects are concerned, both monism and eliminativist OSR agree on their non-existence.[12] Monism has a major asset over eliminitivist OSR, however, in that it can offer a simple account of what is meant by the slogan of OSR: "all there is is structure." According to *monistic OSR*, there is just the world, and "structure" is the name for the complexity of features it possess.

To give monistic OSR any real substance a further challenge must be met. One must offer an understanding of the features of the world (i.e., its structure) in such a way that is specific but does not invoke additional objects (beyond the world itself). Jonathan Schaffer's (2010) favored approach is to attribute to the world *heterogeneous distributional properties.* Schaffer uses the toy example of *being polka-dotted* to illustrate the idea. This property describes a heterogeneous world—its color varies spatially—but it does so via a single complex property rather than a multitude of color properties instantiated at each location. Schaffer suggests that all variation in the world may be accounted for by analogous distributional properties. To account for the specific details of variation in the world, Schaffer appeals to physical configuration space. In the case of *being polka-dotted*, we can capture the exact way in which the world is polka-dotted by considering a configuration space that contains all possible colors at all possible locations (Schaffer proposes a five-dimensional space: $< x, y, hue, saturation, brightness >$) and selecting a subset of these.

> For the monist, the general fact that the world is heterogeneous is due to the world's instantiating the determinable property of *being hetero-*

---

[12] *Priority* monists like Jonathan Schaffer allow that such objects exists, but not *fundamentally.* Schaffer's view will be discussed below.

> *geneous*. The specific way that the world is heterogeneous is due to the world's instantiating the determinate property of tracing such-and-such a curve through physical configuration space. Thus the one whole can be parturient (Schaffer 2010, 60).

Unfortunately, the notion of a configuration space achieves specificity at the cost of reintroducing a plurality of objects into one's ontology. In order to be a *configuration* space, there must be something that is configured. In the case of *being polka-dotted* that something would be the spatial regions at which we are assigning color attributes. The existence of objects such as spatial regions is not a problem for Schaffer's own view, *priority* monism, because it allows for objects other than the world which are metaphysically dependent on it (i.e., are non-fundamental). Yet, this disqualifies the present approach as a version of eliminative OSR by requiring just the sort of objects that the view aims to eliminate.

One prominent version of monism that aims to avoid a commitment to objects other than the world is due to Terry Horgan and Matjaž Potrč (2008). On their view, there is only one thing that exists—the "blobject"—but it has a very rich structure which accounts for all the variation and apparent multiplicity in the world. *Blobjectivism* (Horgan and Potrč's brand of existence monism) may be viewed as a version of eliminative OSR insofar as there are no concrete objects beyond the blobject and hence structure alone must play the role of explaining and accounting for phenomena. When confronted with the question of how to account for variation within the blobject, Horgan and Potrč draw an analogy with "jello-world" (2008, 168–172). Jello-world is a world that consists solely of gunky, jello-ish stuff that is completely continuous and forms a single extended part-less object, the world. The intuition they develop is that the jello-world could have structure, in the sense of spatiotemporal variation, without having any proper parts. The jello-world may be more dense in certain areas, but because we have assumed it doesn't contain parts, we should not understand this as a region of the jello that has the property of being more dense.

> the actual world could be like the hypothetical jello-world, a physical blobject that lacks real parts, and yet still exhibits genuine structural complexity—that is, genuine variability in how magnitudes are locally, spatiotemporally instantiated (2008, 172).

In other words, the world may have a rich spatiotemporal structure—different magnitudes are instantiated in different locations—without having parts. This seems initially puzzling; if there is spatiotemporal structure in this sense, it seems there must be *locations* or *regions of spacetime* to which the different magnitudes attach, in which case there are objects other than the blobject after all. Horgan and Potrč consider this objection to their view but argue that it can be resisted if we think not in terms of properties and relations being instantiated at regions of spacetime, but rather of the blobject instantiating properties *in a manner that corresponds to that spatiotemporal region.* Rather than saying the property *mass M* is instantiated at spacetime region R, we ought to say "the property *mass M* is instantiated R-ishly by the blobject" (2008, 177). Thus, the key to understanding how a part-less world can have structure lies in the idea of instantiating properties in a spatiotemporal manner.

Blobjectivism is closer to the ideal of eliminative OSR than Schaffer's monism in that no objects beyond the blobject are recognized. Structure, on this view, consists in properties which are regionally instantiated by the blobject. Yet, as a solution to the object problem, blobjectivism faces many of the same challenges as Schaffer's monism. Both views must, for example, revise the concept of *structure* considerably from the traditional understanding. Structure, for the monist, is not a network of relations, but rather a *property* of the whole. For Schaffer, structure is a heterogeneous distributional property, while for Horgan and Potrč it is a collection of properties which are regionally instantiated by the blobject. The worry associated with revising our understanding of structure in this way is that we lose our grip on the concept. In this particular case, regional instantiation seems either to be a mere redescription of regions instantiating properties, or some new kind of instantiation relation. If the former, then we haven't made any real progress in eliminating objects, and if the latter, we have moved little beyond the slogan of "all there is is structure."

A further issue faced by monism is what we take the one object to be. Typically, monists refer to it as "the world" or "the cosmos" or "spacetime," but given that physics has yet to decide whether spacetime is fundamental, or whether there is a "multi-verse," it is hard to pick out the "one" with which monists are concerned. Ultimately, monism provides an interesting framework from which to consider SR, but it does not offer a solution to the object problem amenable to the eliminativist OSRist. Even ignoring the one object posited by monism, there remain difficulties in understanding the structure of the world in object-less terms.

4.4 Non-Eliminative OSR

The failure of eliminative OSR to resolve the object problem has led some to consider a weaker form of the view that would embrace the requirement of objects while preserving the importance of structure. Non-eliminative OSR (NEOSR) attempts to walk the line between eliminative OSR and traditional realism by allowing that there are objects, but that these objects are "thinner" than the traditional realist's objects. A useful analogy comes from graph theory. Euler helped to lay the foundations for graph theory with his famous treatment of the problem of the Seven Bridges of Königsberg.

> The East Prussian city of Königsberg (now Kaliningrad) occupies both banks of the River Pregel and an island, Kneiphof, which lies in the river at a point where it branches into two parts. There were seven bridges that spanned the various sections of the river, and the problem posed was this: could a person devise a path through Königsberg so that one could cross each of the seven bridges only once and return home? (Alexanderson 2006, 567)

Euler showed that there is no solution to the Seven Bridges problem by presenting the layout in abstract terms. Each land mass is represented by a point-like node and the bridges are represented by lines (called "edges") connecting the nodes. Once represented by such a graph, the problem becomes clear. In order to traverse each bridge exactly once an even number of edges connecting each node is required (one to enter the land mass and one to exit). As it turns out, each node is connected by an odd number of edges ("is of an odd degree"), and hence no solution to the problem is possible.

Figure 4.1: The Seven Bridges problem

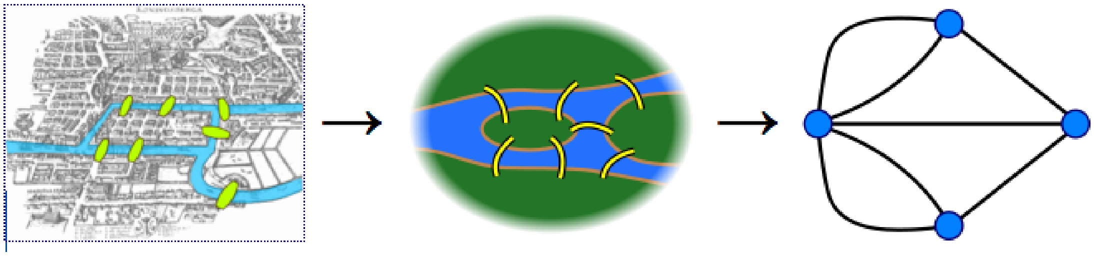

Image source: Wikipedia (2013)

There are several aspects of graph theory that are attractive from the perspective of NEOSR. First, graphs provide a way of capturing the intuitive notion of "the structure" of a situation. The graph of the Seven Bridges problem, when compared to a map of the bridges themselves bears a resemblance that is naturally described as sharing a common structure. Second, nodes in a graph provide a useful model for what might be called a "thin" object. The nodes in the graph of the Seven Bridges problem are *nothing more* than places in the graph structure. They are objects that are necessary for representing the structure of the problem, but are quite different from the land masses they represent in that they lack all non-structural properties.

Of course, graph theory offers only an *analogy* for NEOSR,[13] but the lesson is clear. Objects, according to NEOSR, are real but exist merely as elements in a relational structure. Typically, physical objects are thought to be independent of the relations they enter into, but NEOSR denies this. Rather, objects on this view are "first and foremost" places in a structure. Stewart Shapiro, in his (1997) defense of

---

[13] Cf. Dipert (1997) who takes the comparison quite literally.

mathematical structuralism, makes a useful distinction between two different ways of thinking about objects in structural terms. On one hand, the "places as offices" view holds that places in a structure are *offices to be occupied by objects*. For example, $\{\varnothing\}$ occupies the "one-office" in a particular realization of the natural number structure. On the other hand, the "places as objects" view takes the places themselves to *be* objects. For example, "the Vice President is President of the Senate" treats the office of the Vice Presidency as an object in itself. Shapiro advocates the latter "places as objects" view for mathematical objects and it is this perspective that fits well with the goals of NEOSR when physical objects are concerned.[14]

Putting these ideas together, we arrive at an ontological picture in which "thin" objects exist as places in a structure, defined by their relational properties. Such a view aims to preserve the central claim of OSR that "all there is is structure" by regarding objects as mere *elements of* structure. In doing so, it avoids the object problem by granting that structures require objects and relations require relata. How to move from this general picture to a specific proposal, however, is less clear. A significant obstacle is that the analogy with mathematics belies the ways in which physical objects are different from abstract or mathematical objects. It is a common refrain that physical objects, unlike mathematical ones, cannot be purely "formal" or "structural" in the way that mathematical objects are.

Chakravartty (2012) objects to NEOSR for reasons related to this general worry. First, Chakravartty characterizes NEOSR broadly as a view in which objects are ontologically less fundamental than structure. Second, he takes this to require that

---

[14] See also, Resnik (1997). Resnik argues that mathematical objects "are themselves atoms, structureless points, or positions in structures" (1997, 201).

objects depend on the relations they enter into for their identity. Third, he argues the dependence of objects on relations for their identity requires that such objects lack all other (i.e., non-relational) properties.[15] Finally, Chakravartty concludes that "an object with no intrinsic features at all... is not an object at all" (2012, 197). In other words, the idea of a physical object with a purely relational identity is incoherent.

Chakravartty gives two main arguments for the claim that such relational physical objects are impossible. The first begins with a discussion of graph theory similar to that above. Nodes in a graph *do* seem to be examples of relational entities; one can exhaustively characterize a particular node by appeal to the relations it enters into with other nodes (*if* these relations are unique). However, nodes are mathematical (or at least abstract) entities. In the case of physical objects, Chakravartty argues such a relational characterization is insufficient. "[T]he labels in graphs ... do not furnish identity criteria for objects or properties that might *occupy* their nodes in such a representation; they simply identify *locations* in a structure!" (2012, 199).

However, this criticism ignores Shapiro's distinction above. On the "places as offices" view, Chakravartty's criticism may apply, but not on the "places as objects" view. Perhaps he regards the "places as objects" view as inapplicable to *physical* objects, but this claim requires further argumentation. Simply dismissing the application of the "places as objects" view to physical objects is to beg the question against NEOSR.[16]

---

[15] Chakravartty offers the following in support of this claim: "So long as the relata *have* genuinely intrinsic features—qualitative properties, dispositions, what have you— ...these intrinsic features keep popping up as plausible candidates for determining their identity" (2012, 197).

[16] Chakravartty suggests that in virtue of being *concrete*, physical structures require locations to be *occupied*. But the NEOSRist would deny this inference if "being occupied" requires anything

The second argument offered by Chakravartty centers on the notion of relational ("extrinsic") properties themselves. "The very attribution of an extrinsic [relational] property *assumes* that one has a prior grasp, ontologically speaking, of what it is that stands in the relevant relation or relations" (2012, 201). In other words, in order to attribute a relational property there must be *something* that stands in that relation. But, what could this something be in the case of a purely relational object? It could be a collection of previously attributed relational properties, but this obviously leads to a regress. Thus, it seems that there must be some non-relational properties or attributes to get the process started; there must be some prior *thing* that stands in the "first" relation R_0 to something else.

This argument also misses the mark. NEOSR does not claim that one can build objects out of prior relations, but rather, that objects are nothing more than a nexus of relations.[17] When a structure is specified—whether set-theoretically or informally as a network of relations—there is an ineliminable reference to objects or relata. NEOSR recognizes those objects as genuine existences, but emphasizes their dependence on the structure in which they are embedded. That objects must exist prior to their relations does not follow from the fact that we need something to which to *attribute* a relational property any more than the attribution of ordinary properties requires belief in prior substances or "bare particulars." In both cases a *bundle theory* according to which objects are nothing more than the sum of the

---

more than being able to treat that location as a object. On the version of NEOSR currently under investigation, the world is a concrete structure, and places in that structure simply *are* objects.

[17] Even this formulation is too strong. Minimally, NEOSR is only committed to the idea that objects are to be understood as elements of structure.

their properties is possible.[18]

Chakravartty's argument against NEOSR can be resisted at several points. The two main arguments he provides against relational entities are either invalid or beg the question against the NEOSRist. Moreover, one may resist Chakravartty's allegation that NEOSR is committed to relational entities by considering other ways of understanding "objects depending on structures" or "objects existing as elements of structures." Yet, regardless of the efficacy of Chakravartty's argument as a whole, the disanalogy between mathematical objects and physical objects he emphasizes does seem to pose a problem for NEOSR. The problem I have in mind emerges most clearly in the final portion of his paper where he considers applying NEOSR to actual cases from science.

> what makes something a space-time point (if there are such things) as opposed to a subatomic particle is a function of some important intrinsic differences, even if it turns out that in order to individuate one space-time point as distinct from another, one must rely on their extrinsic properties. (2012, 201)

Another way of putting Chakravartty's claim here is that physical objects must be *complete* in a way that mathematical objects may not be. Michael Resnik, a mathematical structuralist, characterizes *incompleteness* as follows.

> Mathematical objects are incomplete in the sense that we have no an-

---

[18] Paul (2002, 2006) offers one contemporary defense of bundle theory. On her view, objects are mereological fusions of property instances (or, in some cases, multiply locatable properties). There is nothing particularly "structuralist" about such a view, but it does illustrate that it is at least *coherent* to view objects as collections of (instances of) properties. Given this, it's hard to see how viewing objects as collections of *relations* would fail to be coherent.

> swers within or without mathematics to questions of whether the objects one mathematical theory discusses are identical to those another treats; whether, for example, geometrical points are real numbers. (Resnik 1997, 90)

For Resnik, this incompleteness derives from the way definitions work in mathematics and fits well with his view of mathematical objects as positions in patterns. "[R]estricting identity to positions in the same pattern goes hand in hand with the failure to have any identifying features independently of a pattern" (1997, 211). Physical objects, however, seem different. To simply deny there is a fact of the matter concerning identity claims across different structures seems highly counterintuitive. To return to Chakravartty's example, surely we should be able to say whether some collection of relations belong to a spacetime point or a subatomic particle. Yet, if physical objects are positions in structures, then such a query cannot be answered.[19]

Thus, while the "places as objects" view is one possible construal of NEOSR, it has the counterintuitive consequence that physical objects are incomplete.[20] It is worth asking, then, whether there are other ways of developing NEOSR that do not have this implication.

4.4.1 Dealing with Incompleteness

The main goal of this chapter is to establish NEOSR as the most promising version of SR. It does not postulate unknowable entities without falling prey to the

---

[19] Alternatively, we may regard all questions concerning trans-structural identity as trivially false.

[20] Although, there may be independent reason to think physical objects are incomplete. Aspects of modern physics may be taken to have that implication. See Chapter 5 below for a discussion.

object problem that faces eliminative OSR. So far, I have assumed that the view is developed—following the example of mathematical structuralism—as one in which objects are identified with positions in structure. But this seems to require that physical objects inherit the *incompleteness* of mathematical objects. Some may view this consequence as unacceptable and thus conclude that NEOSR, to the extent that it entails incompleteness, is equally unacceptable.[21]

There are two problems associated with incompleteness. The first is that physical objects, unlike mathematical objects, seem to have a full compliment of properties, not just those associated with a particular relational structure in which they are embedded. The second problem is that without these additional, extra-structural, properties, it is impossible to to identify the *same* object across structures. Consider an electron, for example. A given electron might occur in a structure that relates it to all other particles as described by its quantum state, and thus the current approach will recommend identifying the electron with a particular place in the quantum state structure. Yet, what makes this place an *electron*? Here, it seems, we need to appeal to the electron's state-independent properties, like its mass and spin parameters, in order to fully specify that it is an electron occupying that place in structure. The problem is that, by recognizing such properties, it seems as though we have slipped back into the "places as offices" view, and this is a view in which there is more than just structure.

Thus, the proponent of NEOSR seems to face a dilemma: either she must main-

---

[21] Some may be perfectly happy with viewing physical objects as incomplete. Resnik suggests that Quine's doctrine of ontological relativity (Quine 1968) could be viewed as embracing incompleteness. For Quine, we can know objects—mathematical or not—only *up to isomorphism*, hence there is simply no fact of the matter about the identity of objects across structures.

tain that objects just are places in structures and accept incompleteness, or else invoke intrinsic properties that undermine the claim to be a form of OSR. One version of NEOSR that has confronted this problem is the *Moderate Structural Realism* of Michael Esfeld and Vincent Lam mentioned in Chapter 1 (Esfeld 2004; Esfeld and Lam 2008). According to this view,

> [t]here are relations as well as objects standing in the relations without there being any ontological priority between them. Relations and objects are both genuine fundamental ontological entities. They are on the same ontological footing, being given "at once" in the sense that they are mutually ontologically dependent on each other. (Esfeld and Lam 2011, 146)

This statement of Moderate Structural Realism allows for the possibility of intrinsic properties in addition to the mutually dependent objects and structure. Yet, positing such properties would mean embracing the second horn of the above dilemma, which Esfeld and Lam wish to avoid. As mentioned in Chapter 1, Esfeld and Lam have since updated their position. In the original formulation of their view (Esfeld and Lam 2008), the authors denounce the appeal to intrinsic properties in favor of a view according to which there are only relational objects and the structures in which they are embedded. The new version of Moderate Structural Realism makes two important changes: (1) the mutual dependence between objects and relations is regarded as *conceptual* rather than ontological, and (2) objects may have intrinsic properties so long as these do not constitute an intrinsic identity.

> [E]ven if some of the ways fundamental physical objects are consist in in-

> trinsic properties, important features of fundamental physics show that there always are central ways in which these objects exist that consist in relations, so that the objects do not have any existence—and in particular not any identity—independently of the structure they are part of. (Esfeld and Lam 2011, 155)

The new version of Moderate Structural Realism attempts to avoid the above dilemma by recognizing intrinsic properties, but not as distinct *ontological* posits. This understanding of properties as "ways in which objects exist" (due to Heil (2005)), apparently allows Esfeld and Lam to embrace intrinsic properties without running afoul of the OSRist dictum that *all there is is structure.* But, from an ontological point of view, their new proposal seems to emphasize *objects* rather than structures. In fact, there is an important sense in which structures do not really exist on this view. Moderate Structural Realism, in its new form, recognizes structures only as ways in which objects exist, but this is to say that really there are only objects that exist in irreducibly relational ways. Thus, Moderate Structural Realism in this form hardly deserves the title "structural realism" given that structures, understood ontologically, simply do not exist.

There is a better way to escape the dilemma above. This is to take seriously the idea that objects are places in structures, but to broaden the structures we consider. Roberts (2011) outlines a view he calls "Group Structural Realism" that takes the important structures in physics to be *group-theoretic.* In particular, it draws on a approach to elementary particles often attributed to Wigner (1939), that treats elementary particles as irreducible representations of the Poincaré group. Wigner recognized that all possible combinations of the crucial quantities for elementary particles—spin, mass and parity—can identified with irreducible representations of the spacetime symmetry group. Wigner's approach allows the structural realist to understand the properties that characterize a particular elementary particle in structural terms; we are now able to say what makes something an electron, say, without appealing to non-structural intrinsic properties. Group representations allow us to capture the state-independent features of elementary particles in structuralist terms.[22]

Wigner's group-theoretic approach to elementary particles is a specific version of more general strategy for addressing the issue of incompleteness; often questions of *what it is* that "occupies" a given place in a given structure can be addressed by considering other, higher-level structures. In the case of elementary particles it seems as though intrinsic properties are needed to fully characterize a given particle, but in fact group-theoretic structures are up to the task. This suggests the following version of NEOSR:

NEOSR Ontologically speaking, there is only structure, but objects may be recovered as places in *multiple* structures.

By identifying places in several structures, this allows the structuralist to not posit anything other than structures, while offering a complete characterization of objects. Often, as in the case of elementary particles, some of the relevant structures will be "higher-level" like the group-theoretic structures appealed to above. In

---

[22] Earlier I criticize the Bain's category-theoretic version of structural realism for being too abstract to capture the concrete physical relations with which structural realists are concerned. Now I am appealing to group-theoretic structures to help characterize physical objects. Yet, the important difference is that while Bain proposes doing away with set-theoretic structures altogether, I am arguing for a role group-theoretic structures *in addition* to ordinary set-theoretic ones.

addition to providing a complete characterization of objects, this version of NEOSR also allows that the same object can occur in different structures at the same "level of description." For reasons that will made clear in the next chapter, we may wish to say the *very same* spacetime point p occurs in distinct geometrical structures. NEOSR, on the present understanding, allows for such a situation.

4.4.2 Conclusion

The above discussion may be summarized as follows. We began we the goal of developing structural realism, the view that science can only reveal the structure of the world, into a coherent position distinct from traditional scientific realism and antirealism. We found that OSR is preferable to ESR because the latter requires a gap between what is (in principle) knowable and what exists that is precluded by our commitment to naturalism. Our consideration of OSR, however, led immediately to the object problem: how can structure be all there is when structures seem to require objects? Taking the object problem serious led us to reject eliminative varieties of OSR in favor of non-eliminative OSR (NEOSR): there are only structures but objects exist as *places* in structures. Finally, in response to worries about completeness, it was proposed that objects should be understood as places in *multiple* structures.

As the next chapter aims to show, this picture is not just recommend by conceptual considerations, but also fits well with what physics tells us about fundamental reality.

CHAPTER 5

Minimal Structural Essentialism: Why Physics Doesn't Care Which is Which

5.1 Introduction

So far, the motivation for structural realism presented here has focused on theory change in the history of science. The upshot of Chapters 2 and 3 was that structural realism can find continuity in the history of science that eludes the traditional realist. In this chapter, a different sort of motivation for structural realism is presented. I will argue that an interesting feature of contemporary physical theories is best explained by a metaphysical picture consistent with the version of NEOSR advocated at the end of the last chapter.

This curious feature is what John Stachel (2002) calls *Generally Permutability.* Very roughly, Stachel claims that the two pillars of modern physics, general relativity (GR) and quantum mechanics (QM), don't seem to care which basic objects possess which properties. More carefully, Stachel characterizes general permutability as follows:

> General Permutability (GP): "every permutation P of the **a** entities in $\mathbf{S}$, $\mathbf{R}(a)$, $\mathbf{R}(\mathrm{P}a)$ represent the *same* possible state of the world," (*ibid.*, p. 242)

where $\mathbf{S}$ is a set of n entities, $a_1, a_2, ..., a_n$, among which there is an ensemble $\mathbf{R}$ of M n-place relations, $R_1, R_2, ..., R_M$. $\mathbf{R}(a)$ is the ensemble of relations with their places filled by the sequence, $a = (a_1, a_2, ..., a_n)$ and Pa is the permutation of the sequence a.

GP is an interpretative principle. It claims that *representations* or *models* that differ only by a permutation of the entities they contain represent the same physical state of affairs. In the case of GR, this amounts to the claim that “swapping” points in spacetime *models* results in redundant representations of the same spacetime. Thus, GP is committed to a *passive* reading of the permutations of model elements: such permutations reveal a formal symmetry of the theory rather than a symmetry of the (physical) world.

Our goal here is to explore the implications of GP for a metaphysics of objects in fundamental physics. In particular, it has been argued—by Stachel himself and more recently by Adam Caulton and Jeremy Butterfield (2012)—that GP recommends a “structuralist” metaphysics. I agree with these authors that structuralism is recommended, but I find their particular metaphysical proposals insufficient. Instead, I argue for a version of structural essentialism according to which an object occupies its place in a structure essentially. This view, I contend, offers a promising structuralist metaphysics that accounts for GP.

This view also has surprising implications for the question of individuality in physics. There is an appealing argument that GP (or something like it) renders elementary particles or spacetime points non-individuals by making it indeterminate which object is which. On my view, this argument is mistaken. The best explanation for GP implies that these basic physical objects are *structural* individuals, that is, objects which are individuated by their place in a relational structure.

At the onset, it is worth clarifying what is meant by a “structuralist metaphysics.” “Structuralism” is a term of art in contemporary discussions, and in the preceding chapters I have not carefully distinguished this term from “structural realism.” In this chapter, however, the target is not structural realism but structuralism. Allow me to explain. Structural realism is the position in the philosophy of science according to which we should be realists only about the structure of our best scientific theories. As we saw in Chapter 4, structural realism further divides into epistemic and ontic varieties (ESR and OSR, respectively). According to OSR, there is nothing more to fundamental reality than structure. This entails the thesis I intend by structuralism: the correct metaphysics of the world is fundamentally structural.

Notice that one can be a structuralist without being a structural realist. Structuralism is an ontological thesis that one may come to by many paths (in the present case, by taking aspects of physical theories seriously). The point is that the structuralism I am concerned with is the same one proponents of OSR are committed to. Saying more about what is meant by the world being “fundamentally structural” is difficult because of divergent views about how to view structure, fundamentality, and objects. One thing on which there seems to be close to a consensus among structuralists, however, is that the traditional “atomistic” picture must be rejected. The view that there are atomic building blocks with intrinsic properties out of which everything is composed must be replaced by a view in which relational structures play an ineliminable role. As Esfeld and Lam put it, “Structural realism in the metaphysics of science is a sort of a holism in contrast to an atomism” (2010, 10) and later, “Structural realism as a metaphysical position is the claim that there are no fundamental intrinsic properties underlying the relations that we can know.

That is to say, all there is to the fundamental physical objects are the relations in which they stand" (*ibid.*, 12).

This chapter will proceed as follows. Section 5.2 will briefly explain GP and how it arises in GR and QM. It is not my aim here to establish that GP holds in these theories, but rather to illustrate and lend some plausibility to the claim that it does. Section 5.3 will turn to metaphysical proposals aimed to account for GP. In 5.3.1 I argue that the proposals of Stachel and Caulton and Butterfield, while promising, are insufficient to the task of accounting for GP. In 5.3.2, I return to GR and draw on the views of Tim Maudlin to motivate *Structural Essentialism.* In section 5.3.4, I offer my own proposal, *Minimal Structural Essentialism*, which aims to provide a structuralist metaphysics adequate to the task of explaining GP. The view is clarified and defended from objections in section 5.4. Finally, in section 5.5 I make explicit the implications of the view for the question of individuality in physics.

5.2 The Case for General Permutability in Physics

The motivation for GP is provided by central topics in GR and QM: the hole argument and quantum statistics, respectively.

5.2.1 GP in GR: The Hole Argument

Models in GR are given by the triple $\mathcal{M} = \langle M, g, T \rangle$ where M is a manifold of points, g is a metric field defined on M and T is a matter field defined on M. The manifold of points M represents all of the points in spacetime (or "events"), the metric field g specifies the spatiotemporal distance relations between spacetime points and the matter field T specifies the distribution of matter and energy across spacetime.

The hole argument concerns a particular transformation of the metric and matter fields (g and T) of a model $\mathcal{M}$.[1] A *hole diffeomorphism* is a smooth transformation of the metric and matter fields that leaves them unchanged outside of some arbitrary region (the "hole"), but changes them within the hole region. The resulting model $d\mathcal{M} = \langle M, d^*g, d^*T\rangle$[2] *apparently* represents a world exactly like ours outside of the hole, but different within it. For example, on an active reading of the hole diffeomorphism, the worldline of a particular galaxy that $\mathcal{M}$ describes as passing through spacetime point p located within the hole no longer does so in the world represented by $d\mathcal{M}$.

Taking the result of a hole diffeomorphism to represent a distinct world is problematic, however. Such diffeomorphic models are not only empirically equivalent to the original model, but also indistinguishable from the perspective of the laws of GR. From the perspective of GR, there seems to be no reason to prefer $\mathcal{M}$ over $d\mathcal{M}$, yet they seem to differ in their metaphysical implications. One way to see the problem concerns indeterminism. If $\mathcal{M}$ and $d\mathcal{M}$ are taken to represent distinct worlds, then specifying the entirety of the region outside of the hole underdetermines the facts inside the hole. If we let the hole be a future region of spacetime, then a complete specification of the past fails to determine the future regardless of our dynamical laws. This is problematic not because of the unpalatability of indeterminism as

---

[1] In Stachel's (1989) reconstruction of Einstein's argument the hole argument is formulated so that only g is changed. On this version of the argument the "hole" must be an empty region of spacetime ($T_h = 0$). The presentation here follows the more general form of the hole argument developed in Earman and Norton (1987).

[2] d^* is the drag-along corresponding to the diffeomorphism d. For example, d^*g is a metric field that results from "dragging" the values of g at each point in the original model $\mathcal{M}$ to its image under d.

such, but because in this case it results solely from an *interpretative choice*.

A popular way to resolve this representational underdetermination is to regard the distinct *models* $\mathcal{M}$ and $d\mathcal{M}$ as representing the same *world* W. Thus, this position—known as "Leibniz Equivalence"—resolves the underdetermination of representations by "collapsing" the two distinct worlds apparently represented by $\mathcal{M}$ and $d\mathcal{M}$ into a single world, redundantly described.

It is a short step from here to GP. What the hole diffeomorphism does is to effectively "swap" points within the hole: manifold point α is now where manifold point β was perviously. As Stachel shows (2002, 235), if we abstract away the differentiability and continuity of spacetime models, we are left with a bare manifold of points and diffeomorphisms become permutations of the points. Thus, if Leibniz Equivalence is the right solution to the hole argument, then GP—its discrete analog—is part of the correct understanding of GR.

5.2.2 GP in QM: Quantum Statistics

Stachel also claims that GP is part of QM in virtue of offering the best explanation for the statistics of indistinguishable particles. While the argument for this claim is not made explicit, there is a popular argument that connects quantum statistics with GP: the case of the "bosonic quantum coins."

Consider two bosons each of which can be in one of two states a or b. Our "classical" intuition is that there are four possible configurations: aa, ab, ba, bb. Assuming all configurations are equiprobable, then the probability of aa would seem to be $1/4$. However, the Bose-Einstein statistics appropriate to bosons tell us that $\text{Prob}(aa) = 1/3$, and this is what is found experimentally. What accounts for the failure of the "classical" intuition?

The answer, so the argument goes, is that the configurations *ab* and *ba* "collapse" into a single state in QM.[3] We would get the statistics wrong if we take the bosonic states *ab* and *ba* to be distinct, so we must view those two possibilities as combining (in some sense) into a single possibility. Thus, there are three rather than four configurations and hence the probability for each of them is 1/3. The reason why they "collapse" is the same as in the hole argument: *representations* related by permutation describe the same physical state of affairs. In that case we reduced the possibilities by taking $\mathcal{M}$ and $d\mathcal{M}$ to represent the same world, and here we analogously take *ab* and *ba* to combine into a single state of affairs ($\frac{1}{\sqrt{2}}(ab+ba)$). In other words, taking on GP justifies the move from classical to quantum statistics. .

There is an important disanalogy with the situation in GR though.[4] The hole argument proceeds by presenting a case of representational underdetermination—between $\mathcal{M}$ and $d\mathcal{M}$—which compels us to "collapse" the two models into one by invoking GP. Yet, in the case of QM, there is no such representational underdetermination. The states *ab* and *ba* which are allegedly "collapsed" are in fact *impossible* for bosons to occupy. Instead there is a single mathematical representation of the state—a ray in Hilbert space—that corresponds to the superposition $\frac{1}{\sqrt{2}}(ab+ba)$ which cannot easily be interpreted as "either *ab* or *ba*."

Yet, all hope is not lost for establishing GP in QM. Adam Caulton and Jeremy Butterfield (2012)(henceforth, CB), argue that the representational underdetermi-

---

[3] This notion of collapse is unrelated to the collapse of the wavefunction posited by some interpretations of QM.

[4] This criticism is raised by Oliver Pooley (2006, 104) in response to Stachel and endorsed by Caulton and Butterfield (2012, 32).

nation present in GR can be found in quantum statistics if a broader understanding of the latter is adopted. Rather than limiting quantum statistics to purely symmetric (for bosons) or antisymmetric (for fermions) states, CB propose that we also allow for "parastatistical" states that are neither symmetric nor antisymmetric. If we allow for the possibility of such parastatistics, then states which were formerly represented by a single ray now correspond to an entire equivalence class of such rays which are each related by permutation. Thus, CB argue that GP must be invoked in QM to remove representational underdetermination just as in GR.

Another possible approach would be to view quantum statistics itself as such a reduction. In the case of the bosonic coins, QM replaces the four "classical" possibilities with the three given by Bose-Einstein statistics. In so doing, one may regard QM *itself* as reducing the two "classical" possibilities ab and ba to the single quantum state $\frac{1}{\sqrt{2}}(ab + ba)$. For this reason, perhaps QM (at least *with* the symmetrization postulate) can be viewed as incorporating GP as part of the statistics it allows.

Far more would need to be said to defend either of these approaches to GP in QM. CB's approach rests on taking parastatistics more seriously than some may wish, while the second approach requires taking seriously the four "classical" possibilities despite their inapplicability to bosons. There may be yet other approaches to GP in QM as well.[5] The claim I wish to defend here is only that the prospects for GP as a feature of QM are good. Permutation invariance, which lies at the heart of these considerations, is a central aspect of QM and strongly suggests GP or a principle

---

[5] Another potential route involves thinking of permutation as placing a constraint on the dynamics of systems (French and Krause 2006). From this perspective, GP might emerge as the principle that motivates the move from a full configuration space to a reduced one.

very much like it.[6]

5.3 Explaining GP: Structuralism, Essentialism and Structural Essentialism

In what follows I will leave behind debates over the status of GP in GR and QM and proceed on the assumption that it is a genuine feature of these theories. These are theories in which representations that differ only with respect to which entities have which properties should be viewed as physically equivalent. The question now is: what accounts for this feature of these theories?

5.3.1 Structuralism

Stachel and CB each argue that GP suggests an increased role for *structure* in one way or another. There are several different projects one may be engaged with in this vicinity. My own goal, as mentioned above, is to find a structuralist metaphysics that grounds GP. I take this to mean answering the following question: what is the nature of the world described by GR and QM such that GP follows from it?

My central criticism of the two proposals considered here is that they are inadequate to the task of providing such a metaphysical grounding of GP. Later I offer my own proposal, *Minimal Structural Essentialism*, as an attempt to provide such a grounding.

Stachel's Structural Individuation

Stachel claims that GP shows that:

---

[6] There is substantial disagreement on how to think of permutation invariance in QM. These considerations are essential to a successful defense of GP in QM. See, Earman (MS); Saunders (2006a) for a sampling of the issues involved.

> The basic building blocks of any model of the universe (for example, the elementary particles and the points of spacetime) are individuated entirely in terms of the relational structures in which they are embedded. (Stachel 2002, 249)

One difficulty in understanding this proposal comes from Stachel's use of the term "model." If this is a thesis only about *models*—i.e., that which is used to represent the world—then it is hard to see how it could be supported by GP, a thesis about the relation between models and the world. If, however, we take the proposal to concern particles and spacetime points *themselves*, the proposal suggests that such objects are structurally individuated. The intuitive idea here seems to be that representations related by permutation attribute the same structure to the world, and hence, if objects are picked out by that structure these representations attribute the same properties to the same objects.

Yet, if we take Stachel's proposal to be that certain physical objects are individuated structurally, further questions remain. For example, what is meant by "individuation" in this context? To *metaphysically* individuate an object x—following Lowe (2007)—is to (metaphysically) *determine* which object x is among members of a certain class of objects. For example, the set $A = \{1, 2, 3\}$ is individuated by its members $1, 2, 3$ in this sense.

Alternatively, Stachel may intend a weaker, epistemic, sense of individuation where to individuate objects of the same kind is to render them *distinguishable*. On this reading, to say elementary particles are structurally individuated is to claim that one can only tell them apart by appeal to the structure they enter into. This reading, however, fails to provide a grounding for GP in the sense articulated above.

In particular, it leaves open the question of what it is about these objects that makes them only distinguishable by appeal to the structures into which they enter.

The most promising version of Stachel's proposal, for our purposes, construes "structural individuation" metaphysically: structure determines which object an object is. However, while this certainly constitutes a claim about metaphysics, Stachel does not provide a specific account of the sense in which structures individuates these physical objects. In lieu of such an account, the view is at best incomplete; it is possible that the correct metaphysical account of GP involves structural individuation, but until such a notion can be spelled out, our task is not complete. Indeed, without further specificity, this view looks utterly paradoxical: structural individuation would entail that structure determines which object is which, but GP is taken to deny there is a fact about such matters.

Caulton and Butterfield's Structuralism

Caulton and Butterfield argue that GP supports a view they call "structuralism" and characterize as follows.

> [Structuralism's] central claim is that individuality is grounded, if at all, only on qualitative properties and relations. (Caulton and Butterfield 2012, 236)

CB's structuralism amounts to a prohibition on non-qualitative accounts of individuality. That is, one cannot explain what makes something the individual object it is by appeal to features that would not be shared by a duplicate. What CB aim to rule out are views like that of Adams (1979), who holds that what grounds the individuality of an object a is its "primitive thisness," the non-qualitative property of being identical to a. Yet, it is misleading to call such a view "structuralism" because, while the structuralist would surely reject non-qualitative accounts of individuality, they would also reject certain views based on *qualitative* properties as well.

As mentioned in section 5.1, structuralism holds that structures are ontologically basic or fundamental and structuralists reject the default "atomistic" metaphysical picture. One particularly clear version of atomistic metaphysics is David Lewis's "Humean Supervenience" account (Lewis 1986, ix-xvi). On this picture, there are only local matters of fact and from these (and the spatiotemporal relations between them) all else emerges. Thus, on this view, relational structures emerge from individual objects with intrinsic properties bearing spatiotemporal relations. This is a view that structuralists deny. They claim that there are relations that are ontologically basic; there is nothing prior to them on which they depend.

My claim here is not that GP requires that Humean supervenience is false, or even that it is false. My claim is simply that structuralists—in virtue of taking the world to be fundamentally structural—must reject Lewis thesis. I happen to think that physics does speak against Humean supervenience, but this is a matter of considerable debate.

CB's structuralism, however, is not sufficiently strong to rule out positions like Lewis's. Their version of structuralism leaves open these possibilities by giving the following ontological options: (1) individuals whose individuality is grounded in qualitative intrinsic properties, (2) individuals whose individuality is grounded in qualitative extrinsic properties, (3) non-individuals. The problem with calling this view "structuralism" is that it leaves open option (1) which most structuralists strongly deny.[7]

Putting to one side whether or not CB's proposal counts as structuralism, how does the proposal fare as an explanation of GP? On one hand, it follows from CB's structuralism that *worlds* related by permutation/diffeomorphism—which agree on all the qualitative facts—are identical ($W = dW$). But on the other hand, their proposal amounts to little more than a bare modal claim[8] rather than a genuine explanation of GP. For individuality to come out as purely qualitative is surely a *desideratum* for a metaphysical grounding of GP, but it does not constitute an adequate explanation in itself. Both this proposal and Stachel's (on the metaphysical reading) aim to explain GP in terms of what grounds individuality (structure in the case of Stachel, and purely qualitative facts in the case of CB). My criticism of these views is not that they are wrong—I agree that qualitative, structural facts should ground the individuality of objects in physics—but rather, that they are insufficient qua *explanations* of GP. They do not provide an account of the nature of reality according to these physical theories from which GP follows. In order to find such an account, we must return to GR to consider another perspective on the hole argument.

---

[7] CB acknowledge this in a footnote (2012, n.3), but argue that it is enough for their purposes to allow for the *possibility* of non-individuals and objects with extrinsic individuating properties.

[8] The term "bare modal claim" is due to Dasgupta (2011). Dasgupta criticizes solutions to the hole argument that proceed by simply denying that W and dW are distinct possible worlds. He calls this a "bare modal claim" and argues that views that deploy such claims have done nothing to alter the commitment to non-qualitative facts about the individuality spacetime points that generate the hole argument. Because of this, these views are committed to surplus structure. Of course, CB explicitly deny such non-qualitative facts play any role in "grounding the individuality" of physical objects, but fail to offer any details of how individuality may be grounded qualitatively. Thus, their claim makes little progress from the bare modal claim $W = dW$.

5.3.2 Essentialism

There is one well known approach that aims to account for the hole argument by appeal to the nature of spacetime points: Tim Maudlin's metric essentialism (Maudlin 1988, 1990). On this view, spacetime points bear their metrical relations essentially, so that it is impossible for the very same point p to enter into different metrical relations with other points. Yet, as an explanation of GP, there is an immediate problem: unlike most commentators, Maudlin *denies* Leibniz Equivalence.

Recall that Leibniz Equivalence is the view that diffeomorphic *models* in GR represent the very same world ($\mathcal{M} \models W$ and $d\mathcal{M} \models W$). Maudlin denies this and takes diffeomorphic models to represent distinct worlds W and dW but avoids the problematic consequences of the hole argument by taking dW to be impossible in light of the essential properties of spacetime points. This approach certainly won't do as an account of GP, but it reveals an interesting relationship illustrated by the following inconsistent triad.

1. **Essentialism**: Spacetime points have certain essential properties (or relations) that make swapping them impossible. ($\neg(\diamond W$ and $\diamond dW$ and $W \neq dW))$
2. **Distinctness**: Diffeomorphic models represent distinct physical states of affairs. ($W \neq dW$)
3. **Possibility**: Diffeomorphic models represent possible physical states of affairs. ($\diamond W$ and $\diamond dW$)

Maudlin chooses to reject Possibility on the basis of Essentialism, but we may elect to reject Distinctness instead. If we do so, then Leibniz equivalence can be combined with essentialism. In fact, this constitutes the beginnings of a metaphysical grounding of GP: the essential properties of the objects in question make permuting *them* impossible, which is why *models* that appear to describe just that difference should be viewed as physically equivalent. That is to say, the *active* reading of the hole diffeomorphism is blocked by the impossibility of a distinct world represented by $d\mathcal{M}$, hence we must adopt the *passive* reading (i.e., GP).

5.3.3 Structural Essentialism

Structural essentialism is a generalization of the view just outlined. It holds that objects (in particular, points and particles) occupy their places in a relational structure essentially. In other words, it is part of the nature of being that particular object that it enters into certain relations with other objects.

In the case of spacetime points, this comes very close to metric essentialism. The geometrical relations between spacetime points *in the world* form a relational structure, and to be a given point p is to have a particular location in that structure. However, structural essentialism differs from metric essentialism in two ways. First, it endorses rather than rejects Leibniz Equivalence and second, it applies to particles in QM as well as spacetime points in GR.

In the case of particles in QM, structural essentialism claims that particles occupy their place in structure essentially, but what are the relevant structures? A hint comes from the emphasis many structuralists place on the phenomenon of entanglement. In particular, the spin-singlet state of two fermions has become the standard example of the sort of irreducible relations taken to support a structuralist metaphysics for QM. Such a state may be written schematically as follows.

$$\Psi = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle) \tag{5.1}$$

According to most interpretations of QM, Ψ is a state in which neither particle has a determinate spin value in any direction yet we know with certainty that a spin measurement of both particles in any given direction will find them to have opposite spin values in that direction. This leads Simon Saunders (2003; 2006a; 2006b) to argue that there is a relation of *having opposite spin from* that obtains in this case, despite the fact that one cannot attribute a spin value to the particles taken individually.

Generalizing from this particular case, the structural essentialist takes the relational structures in QM to be comprised of "state-relations" between particles. Entanglement relations are paradigmatic examples of such state-relations, but other relations are possible as well.[9]

---

[9] Things get more interesting when we consider state-relations between more than two particles. For example, we may write the state of three indistinguishable bosons or fermions as follows (where m, n, s represent orthogonal states).

$$\Phi = \frac{1}{\sqrt{6}}[(|mns\rangle + |nsm\rangle + |smn\rangle) \pm (|nms\rangle + |snm\rangle + |msn\rangle)] \tag{5.2}$$

Φ already exhibits the state-relation of "occupying orthogonal states," but there may be additional structures present in addition. If we take the particles involved to be the electrons in a lithium atom in the ground state, for example, then we may add that two of them are in the lowest-energy state $1s$ (with opposite spins) and one is the next energy state $2s$. Such relations between states are difficult to describe in words, but the point remains that there is a certain relational structure in which such particles are embedded.

If we follow CB and allow for the possibility of paraparticles other than bosons and fermions, then Φ does not capture all of the possible configurations of three indistinguishable particles among the states m, n, s. This allows for further structures of state relations.

5.3.4 Minimal Structural Essentialism

Even from this beginning sketch of structural essentialism, a problem becomes apparent in the cases of both points and particles. Consider spacetime points. The structural essentialist (and the metric essentialist) claim that a point p must bear certain geometrical relations to other spacetime points essentially. Yet, in GR such metrical relations are not independent of the material contents of spacetime (as represented by the matter field T). Thus, if the material contents of spacetime had been (significantly) different, then it follows that p would not have existed. This is a counterintuitive result. Typically we think that our spacetime could have been such that there was a fly at p even if there was no such fly there in the actual world.

Maudlin suggests overcoming this challenge by appealing to Lewisian counterpart theory (see also, Butterfield 1989). According to this reply, while it's true that according to metric essentialism p could not have been occupied by (significantly) different matter than it is in the actual world W, there may be a distinct point p' in a possible world W' which grounds the counterfactual claim that there could have been something different at that same point in spacetime. This is not the place to evaluate Lewisian counterpart theory, but it does seem preferable to preserve the intuition that p itself could have been differently occupied, other things being equal.

In the quantum case things are worse. State-relations typically change as time passes (assuming the Schrödinger picture); an electron e may become entangled with other particles or one of the particles it is entangled with may be detected or absorbed. According to structural essentialism, e would not survive such a change. It no longer enters into the same state-relations in the same way and hence it can not longer be the same object. Again, this is a counterintuitive result. While the reidentification of particles in QM is problematic, surely the same particle can change its state or which particles it is entangled with.

In each case, the problem is that structural essentialism prohibits objects from occurring in distinct structures. But recall that GP only concerns cases in which different representations have the *same* structure. Thus, to ground GP, all that is needed is fixing an object's place in a structure *when that structure is present.* This suggests Minimal Structural Essentialism.[10]

> *Minimal Structural Essentialism (MSE)*: for any relational structure S and any object a embedded in S, a has its place in S essentially whenever S obtains.

MSE claims that objects must occupy the place they do *in a given structure*, but allows that they may also be embedded in *different* structures. In the case of spacetime points in GR, MSE says that p must have its place in S in any world with that geometrical structure between spacetime points. In the case of particles in QM, MSE says that the electron e must have its place in R in any world with that state-relation structure. In cases where the structure of the world is different, the object in question may or may not exist. It is often difficult to say under which conditions the same point or particle is present, and MSE does nothing to resolve this difficulty, it simply allows for the *possibility* of such trans-structural identifications.

---

[10] This view is inspired by the observation of Healey (1995) that defeating the hole argument doesn't require metric essentialism, but rather a weaker view he calls "minimal essence." According to minimal essence, if anything has the location properties spacetime point p does in the actual world, then p does. In other words, there cannot be a point q that usurps p's location properties. As the next section will make clear, MSE is a generalization of this thesis.

In so doing, it avoids the counterintuitive consequences of structural essentialism mentioned above.

5.4 Essentialism and Sufficiency

In order to further clarify MSE, it is worth saying a bit more about the notion of essentialism at work here. A traditional understanding is, roughly, the following.

> *Essentialism*: $\Box\forall x(x = a \rightarrow Pa)$

This is perhaps enough to avoid the standard hole argument[11] in the manner endorsed above: the spacetime point p has its location properties essentially (i.e., it enters into the geometrical relations between other points and material bodies essentially) and so a world (dW) in which p has different location properties is impossible, and therefore, we should view diffeomorphic models as redundant representations. But, there are variants of the hole argument that cause problems for such a view.

Consider, for example, a model in GR in which the manifold point α is replaced by an "alien" point or "mole" μ.[12] This model—call it $m\mathcal{M}$—will also be empirically and nomically equivalent to $\mathcal{M}$ and hence GP would recommend regarding it too as merely a redescription of W. But the argument just sketched fails to deliver this result.[13]

---

[11] I will focus on GR and the hole argument here for ease of presentation, but the points carry over equally well to permutations in QM.

[12] This argument and its name are due to Jonathan Schaffer.

[13] One may wonder whether the idea of replacing a manifold point with a "mole" really makes sense. After all, regardless of what we want to say about spacetime points, *manifold* points are abstracta and hence it is not at all obvious that one can regard them as having identity conditions robust enough to support this idea. If one were willing to claim that manifold points have haecceities, then this idea would make sense, but this does not seem to be a very natural view to adopt. I will simply assume here that there is some framework in which replacement of this sort makes sense.

On an active reading of the "mole argument," one arrives at two distinct possible worlds W and mW that differ only in that spacetime point p has been replaced by a distinct spacetime point m. But the essential properties of p don't have anything to say about worlds in which p does not exist, and hence Essentialism is powerless to rule out the active reading. Thus, it seems that a different condition is needed to rule out worlds like mW.

Sufficiency: $\Box\forall x(Px \rightarrow x = a)$

Sufficiency prevents the "mole argument" by stipulating that anything with the particular location properties of p is (identical to) p in all possible worlds. This way the mole world mW cannot be a distinct possible world and hence we can run the argument just as we did in response to the ordinary hole argument. It is also evident that Sufficiency is strong enough to block the original hole argument by ruling out the world dW as well.

Yet, Sufficiency gives rise to its own problems. In certain symmetrical models of spacetime, distinct points may have all the same location properties. But, according to Sufficiency, any point with the properties of p is identical to p, so it follows that such distinct points are impossible. This point is especially pressing in the case of QM. Even if we limit our application to fermions like those in the spin-singlet state discussed above, such particles will often share *all of their properties and relations* (again, on most interpretations). It would be an unfortunate consequence if MSE required identifying particles sharing all of their properties.

Fortunately, MSE isn't committed to either Essentialism or Sufficiency. Instead, it employs a weaker claim.

Weak Sufficiency: $\Box(\exists x(Px) \rightarrow Pa)$

Weak Sufficiency says that if anything has p's location properties then p does. This blocks the hole argument by denying that points with different location properties can swap places, and it also blocks the mole argument by denying that a mole can take the place of p. MSE endorses Weak Sufficiency by claiming that objects have their properties *whenever a given structure is present.* The claim that a structure S obtains entails the existential antecedent of Weak Sufficiency. If the world (or some bit of it) has a relational structure S then there exist *relata* that enter into those relations. Hence, MSE is simply Weak Sufficiency applied to a number of entities entering into a structure.

5.4.1 A Remaining Problem: Symmetric Structures

One aspect of the problem of symmetry mentioned above remains for MSE though. Unlike Sufficiency, Weak Sufficiency doesn't entail that points with the same properties should be identified, but it does seem to allow swapping them to create a distinct possible world. This may not be too troubling in the case of spacetime where perhaps perfectly symmetric models can be dismissed as unrealistic for our world[14], but in the case of QM there appear to be cases in which we attribute the same properties and relations to indistinguishable particles. If MSE is to adequately account for QM, it must have something to say about such situations.

---

[14] Here are two reasons one might think this is the case. First, highly symmetric spacetime models actually used in physics—such as the Schwarzschild solution—are generally regarded as idealizations and as such, we should not assume that the world is as symmetric as they suggest. Second, if our spacetime includes *us*, then it seems to follow that there will be unique geometrical relations between spacetime points and us.

If we return to the spin-singlet state of two electrons considered above (Equation 5.1), most interpretations provide no unique properties or relations for either electron. Hence, MSE's requirement that each electron have the properties it does in the actual world doesn't rule out a distinct possible world in which the two electrons are swapped. Initially, it may seem that one could reply in the same manner as with spacetime points: if we consider the global structure of state-relations—the relations among (the states of) all quantum systems—then perhaps the symmetry will vanish. MSE claims that quantum systems enter into their state-relations essentially, but these state-relations are never as simple as the two-particle spin singlet state used as an illustration. The proper focus is rather the *global* state-relation structure described by the *universal* wavefunction, and perhaps this will be sufficiently rich to attribute to each quantum system a unique set of relations sufficient to distinguish it from all other particles.[15]

The problem with this initial reply is that symmitrization applies to all wavefunctions of composite systems, even the universal wavefunction. Unlike the spacetime case, there *is* good reason to think the global state structure provided by QM is *perfectly* symmetric (or antisymmetric) because QM itself *requires* as much.

The correct reply to the problem of symmetric structures is to distinguish between the properties and relations that characterize a place in a structure and *that place itself.* Saunder's work on weak discernability is of use here. If we return to the spin-singlet state for the moment, the fact that there is a *irreflexive* relation es-

---

[15] The reference to the "universal wavefunction" here is not meant to suggest that MSE is committed to the Everett interpretation of QM, where that term most often occurs. MSE takes QM to provide a structure of state-relations between particles. If the scope of QM is the entire world, then that structure will be determined by the quantum state of the entire world, whose representation is the universal wavefunction.

tablishes that there are two *places* in the singlet state structure: place A and place B enter into the state-relation of *taking opposite spin values in some direction.* The objects occupying places A and B may share all of their properties and relations, but they occupy distinct places in the relational structure nonetheless, and hence MSE prohibits swapping them by making the occupation of their places essential to them. In other words, distinct places in a structure may be characterized by the very same properties and relations.

The role of weak discernibility in establishing numerical diversity in this context is controversial (see Wüthrich (2009) for one important challenge). My point here is simply that weak discernibility allows us to distinguish *places* in structures from the properties and relations used to characterize them.

Admittedly, this notion of a *place* is in need of further development, but if it is intelligible, then the symmetry problem can be addressed for weakly discernible entities. Two electrons in the singlet state occupy distinct places in the structure of state-relations described by the quantum state of the system they comprise. MSE holds that they occupy these places *essentially* (in the sense explicated above), and hence there is no distinct possible world in which the electrons swap places. This means that models related by such a swap—$\Psi = \frac{1}{\sqrt{2}}(|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)$ and $\Psi' = \frac{1}{\sqrt{2}}(|\downarrow\uparrow\rangle - |\uparrow\downarrow\rangle)$—must be taken to represent *the same world.*[16]

Of course, the symmetry problem remains for systems of absolutely indiscernible entities such as a systems of bosons. It is unlikely, however, that the universal wavefunction will describe a global state structure which is perfectly symmetric

---

[16] Of course, Ψ and Ψ' differ only by a phase change, so there is a sense in which their physical equivalence is trivial in QM. Yet, the present point is that MSE provides a metaphysical grounding for this fact.

(i.e., bosonic) given that we think fermions are included among the fundamental particles of our world.

5.5 Structural Individuality

There is a natural connection between (general) permutability and the issue of individuality in physics. E. J. Lowe (1998), for example, characterizes individuals ("individual objects" in his terms) as having determinate identity-conditions and determinate countability. GP, however, seems to undermine the determinate identity of spacetime points and elementary particles, as Lowe suggests.

> The single electron shell of a neutral helium atom contains precisely two electrons: and yet, apparently, there is no determinate fact of the matter as to the identity of those electrons. This is because the two electrons in the atom's shell exist in a state of so-called 'superposition', or 'quantum entanglement'. Our inability to say which electron is which is not merely due to our ignorance, or inability to 'keep track' of an electron in such circumstances: not even God could say which electron was which, because there is simply no fact of the matter about this. What this means is that an identity statement of the form 'x is the same electron as y' may simply not be either determinately true or determinately false. (*ibid.*, 194)

Thus, according to Lowe, electrons are not genuine individuals ("individual objects") because there is no fact about 'which is which', hence they lack determinate identity. Presumably, the reason for thinking that there is no fact about the matter about which electron is which in a superpositional state has something to do with the applicability of GP; the fact that we can permute the situation and nothing changes shows that there is no fact of the matter about which is which. But it is here that must be very careful to distinguish permutations at the level of *models* from permutations of actual electrons. GP, I have claimed, should be understood as a claim about the *relation* between models and the world; permuting elements of a model results in a distinct model that represents the same world. This means that GP *does not* establish that there is no fact of the matter about which electron is which, and hence, electrons *may* be taken to be individuals after all.

MSE takes advantage of this possibility and argues that elementary particles (and spacetime points) in fact *are* individuals. These objects have determinate countability and identity-conditions in virtue of their place in a relational structure in the manner illustrated in the previous section. For this reason, we may call them *structural individuals*. MSE thus offers a way to implement Stachel's aim of individuating these objects "entirely in terms of the relational structures in which they are embedded."

It is worth noting that structural individuals differ in important ways from traditional individuals. For example, Leibniz's Identity of Individuals (PII) fails if formulated in the usual way: $\forall F[(Fx \iff Fy) \rightarrow x = y]$ where F ranges over properties and x and y are objects . As we've seen, two electrons may share all of their properties while remaining distinct. However, if we understand F as ranging over *places* in structures rather than *properties*, the principle does hold for structural individuals.[17] Thus we arrive at the surprising result that if MSE is the best explanation of GP in QM and GR, then the basic objects of these theories are individuals in Lowe's sense—albeit *structural* individuals for which the PII (traditionally formulated) fails.

5.6 Conclusion

I have argued that MSE provides a metaphysics of objects capable of accounting for GP as it occurs in the two pillars of modern physics, GR and QM. I will conclude by briefly saying something about the sense in which this proposal is *structuralist.* Above it was remarked that structuralism takes the world to be "fundamentally structural." MSE takes objects to be essentially structural: to be a certain point or particle requires occupying a certain place in a relational structure. This bestows on structure an ineliminable role in determining what the world is like fundamentally. While MSE does not go so far as to replace objects with objectless structure(s), it does recognize irreducible structures as part of our fundamental ontology. To this extent, the view embraces the structuralist dictum that the world is "fundamentally structural."

---

[17] Alternatively, we may note that this version of the PII requires *absolute* discernibly for distinctness, and propose instead that it allow for mere weak discernibly as well: $\forall F \forall R([(Fx \iff Fy) \wedge \neg(Rxy \vee Ryx)] \rightarrow x = y)$ where R is an *irreflexive* relation. Either way, there is an important difference in the form of the principle that holds for structural individuals in comparison to the traditional PII.

CHAPTER 6

Conclusion

6.1 Putting the pieces together

My goal in the preceding chapters has been to motivate and sketch a new version of structural realism. The initial motivation I offer is the need to avoid the problems that arise for scientific realists from the history of theory changes in science. Structural realism allows one to avoid these problems by urging commitment to only the structures described by scientific theories. I do not claim that successive theories posit the same structures (they do not), but rather that one can find *continuity* across changes of theories at the structural level that are unavailable to the traditional realist. Worrall (1989) found evidence for such continuity in the preservation of mathematical equations and, indeed, it does seem to be an institutional norm that newer theories can recover the equations of the theories they replace. Yet, Newman's objection makes clear that a commitment to purely "formal" or mathematical structure is not enough to be a scientific realist. Rather, one needs to focus on the *physical* structures built up out of concrete relations that mathematical models represent.[1]

An example of the fruitfulness of the structural realist approach is the case of Newtonian mechanics and special relativity discussed in Chapter 3. The realist

---

[1] It is important to recall that the models provided by our best theories only describe the world approximately. Thus, successive theories attribute different structures to the world, neither of which is perfectly isomorphic to the actual structure of the world.

must offer some account for the enduring success of Newton's theory (or, at least, Newton's laws) and the usual account is that Newton's laws are a limiting case of a suitable relativistic dynamics. However, as Kuhn (1962/1996) emphasized, it is difficult for the realist to make this claim given the different ontologies and ideologies associated with each theory. The structural realist, by contrast, is committed only to the structure of these theories. Hence, for Newtonian mechanics to be a limiting case of relativity is for them to share the same structure at the limit; that is, for their models to be isomorphic. This claim is established by demonstrating that at the classical limit ($c \to \infty$) Minkowski spacetime becomes (isomorphic to) Galilean spacetime.

In Chapter 4, I argue that the best way to be a scientific realist is to be a non-eliminative ontic structural realist (NEOSRist). NEOSR combines ontic structural realism (OSR)—i.e., the claim that there is only structure—and non-eliminativism about individual objects. The key to NEOSR is making sense of the idea that objects are *nothing over and above* structure. I propose to do this by identifying objects with *places* in structures. Unfortunately, it seems that places in structures lack the "completeness" of physical objects; there is no obvious way to compare a place in one structure to a place in another. This would mean that many questions concerning the identity of objects do not have answers. Rather than accept this counterintuitive result, I propose to view objects as positions in *many* structures.

Finally, Chapter 5 argues that structuralism—the view that world is fundamentally structural—can account for a curious feature of modern physics: general permutability. The guiding idea here is that models related by a permutation of items should be taken to represent the same physical situation because they have the same structure. If objects are defined in terms of structures, then models that represent the same structure are equivalent as far as objects are concerned as well. The specific proposal made there is that objects have their places in structures *essentially*. Thus, it is impossible for the same object to occupy a different place in the same structure, as an active reading of permutation would require.[2]

The common thread among these various parts is an emphasis on physical structures over objects, and a metaphysics that locates objects within structures. Yet, questions remain about how to put these various pieces together into a unified picture.

6.2 Conflicting conclusions?

If we focus only on the conclusions of the preceding chapters we get the following: (1) structural realism is the best way to be a scientific realist, (2) structural realism allows the realist to view Newtonian mechanics as a limiting case of relativity, (3) structural realism is best developed as the view that objects are *places* in many structures and (4) minimal structural essentialism offers the best explanation of general permutability. How are these intended to fit together?

The first two chapters (1) and (2) provide motivation for structural realism. In particular, they are intended to convince the scientific realist that structural realism offers the best hope of addressing the problem of theory change. Once we have some reason to favor structural realism, we need to know how the view

---

[2] This is the case even in symmetric structures where distinct places may be characterized by the very same properties and relations. A model in which objects occupying such (absolutely) indistinguishable places are switched must not be taken to represent a distinct physical state of affairs.

should be developed. Structural realism finds itself in the odd position of having many committed advocates while also facing charges of outright incoherence. The conclusion (3) is intended to respond to those critics; structural realism can be made into a coherent view that recognizes both objects and structures. Yet, we may wonder how (3) is compatible with (4). Objects *being* places in structures is not the same as objects *having* their places in structure essentially. In particular, the latter claim seems committed to the "places as offices" view that was rejected in Chapter 4.

However, although Minimal Structural Essentialism is weaker than (3), the two views are compatible. Recall, the sense in which objects have their places in structure essentially is that every possible world in which a structure S obtains, the object a in S must occupy its place in S. While this formulation uses occupancy language, we should not place too much weight on this convenient formulation. The actual modal claim is satisfied if objects simply *are* places. Objects *could* have a full complement of intrinsic and relational properties while satisfying Minimal Structural Essentialism, but they need not. Indeed, it is trivially true that if objects are places in structures then they must "occupy" those places in all possible worlds.

Is there any reason to favor the more sparse "places as objects" view over more traditional views of objects also consistent with Minimal Structural Essentialism? I think there is. Notice that, for objects of the same kind, general permutability ensures that models related by permutation are empirically equivalent. If, however, the objects had unique nonstructural properties, then we would expect that permutation would result in an empirical difference. Indeed, if we adopt the naturalistic principle invoked above (section 4.2.1), we should posit no such properties unless our best science requires them.

In the case of spacetime points, perhaps the story can end here. Each point should be taken to be a place in the geometrical structure of spacetime because science can reveal no difference between them that derives from properties beyond those attributed by the structure they are embedded in. Yet, in the case of the particles of quantum mechanics the case is more complicated. Even though no experiment can decide between models related by a permutation of particles *of the same kind*, there remains a sound scientific basis for the additional properties of elementary particles: the distinctness of different *kinds* of particles.

It is here that, as discussed in section 4.4.1, the proponent of NEOSR must find a suitably structuralist account of the state-independent properties that distinguish among the different kinds of elementary particles. Wigner's approach provides a means for doing just this by viewing particles as group representations. If this approach is successful, naturalism recommends adopting the NEOSR as our metaphysics of objects.

There is another, more general, way to see how general permutability supports NEOSR. Chapter 5 argues for two claims: (1) quantum mechanics and general relativity are (plausibly) generally permutable—that is, models related by permutation are empirically equivalent, (2) such models should be taken to represent the same physical situation. Essentialism accounts for (2): models related by permutation must represent the same situation because they describe the same structure, and objects occupy their position in structure essentially. But, ignoring this claim for the moment, we may wonder why are these models empirically equivalent to begin with. Why would models related by permutation be empirically equivalent? A natural answer is that only the structure—what is preserved under permutation—is empirically relevant. This is another way of noting that all we can know empirically is structure. Then, we can apply the naturalistic argument of Chapter 3 to establish that all there is is structure.

In Chapter 5 I did not go beyond arguing for Minimal Structural Essentialism because I wanted to show the weakest form of the principle needed to ground general permutability. This was done in an effort to make the thesis applicable to a wider assortment of views. Yet, the preceding is intended to show that the arguments of that chapter not only support Minimal Structural Essentialism, but also fit well with the version of NEOSR developed in the preceding chapter.

6.3 Remaining questions

This project is an attempt to develop a structural realist framework to guide our understanding of scientific theories. As such, there remain many questions to be answered in order to develop the program further. What follows are some brief suggestions at directions for future development.

6.3.1 What is the scope of SR?

The brand of structural realism advocated above calls for a reconception of objects as nothing more than places in the structures described by scientific theories. A natural question, then, concerns what theories fall under the scope of this claim. I have focused on theories of fundamental physics, such as quantum mechanics and relativity, but some story is owed about the vast number of scientific theories that are *not fundamental*.

If the foregoing story about fundamental physical theories is true, then there are two basic strategies for dealing with non-fundmental science. First, one may attempt to apply the same structural realist framework to the these non-fundamental sciences. If this approach were successful, we could conceive of the objects in the special sciences—cells, living organisms, people—as places in many structures. Second, one may claim that a more traditional metaphysics emerges when we leave the realm of fundamental physics. On this view, the entities mentioned by non-fundamental science should be thought of in the ordinary, non-strucuturalist way (as "self-subsistent individuals" (Ladyman and Ross 2007)). They have intrinsic and extrinsic properties, they obey the PII and while they *occupy* places in structures they aren't constituted by them.

Both strategies have advantages and costs. There may certainly be a benefit to applying structural realism to sciences other than fundamental physics. For example, Don Ross (2008) argues that structural realism offers important insights in economics and Steven French (2012a) claims that structural realism can be extended to biology. More generally, the functionalist approach in the philosophy of mind might also be viewed as a version of structuralism.

Yet, while it is undoubtably true that certain issues in the special sciences may be served by adopting a structuralist framework, the foregoing discussion does not easily carry over to the special sciences. For example, in Chapter 5 it was argued that structural essentialism explains general permutability by compelling us to view permuted models as redundant representations. Yet, this would seem to suggest that general permutability will occur wherever this version of structuralism is implemented. Thus, one might expect that if we apply this version of structural realism to biology we should expect that biological objects of the same kind also cannot be actively permuted.

This reasoning, however, is mistaken. Minimal Structural Essentialism (MSE) claims only that objects cannot be permuted so as to generate a distinct but *qualitatively identical* world. This means that objects which are (absolutely) discernible can be permuted; when we switch two biological organisms embedded in a common structure the result will be a new structure which is *not* qualitatively identical to the original. MSE only blocks switching two objects when doing so results in a qualitatively identical world.

This means that MSE does not raise a problem for the special sciences so long as it doesn't traffic in (absolutely) indiscernible objects. If fact, we may wonder what force the present version of structural realism would have when applied outside of physics. We may find that, while the first strategy is coherent, it does not provide a particularly new or interesting view of the special sciences.

The second strategy argues that no such revision of our understanding of the special sciences is necessary. On this view, structural realism makes no new claims with respect to the special sciences directly. The interesting question then becomes how the ordinary ontology of the special sciences emerges from a structuralist understanding of fundamental reality.

I have argued that general permutability in quantum mechanics and relativity suggest that the world described by those theories is structuralist: basic objects like particles and points do not exist outside of the structures in which they are embedded. The question we now face is how a more traditional picture of objects and properties can emerge from such a structuralist ontology.

The details of such a story are beyond the scope of the present project, but at the heart of the account must be a notion of increasing independence. As we move from the esoteric objects of physics to ordinary macroscopic objects two things happen: (1) objects have more distinguishing features, (2) connections between objects become weaker. Explaining how these very general features arise as one transitions to the level of description employed by non-fundamental sciences is the main challenge for those who adopt the second strategy.

Whatever the emergence story looks like, it will certainly *not* be a version of traditional atomistic composition. It cannot be the case that the non-fundamental entities are simply collections of fundamental entities given our understanding of the latter. Above it was argued that objects like spacetime points and elementary particles should be regarded as places in structures. If this is so, then they cannot serve as atomic building blocks in the usual sense. The properties of the relational structures in which they are embedded do not supervene on the properties they possess when considered individually, hence any account of their relation to non-fundamental objects must respect this ontological holism.

One structuralist view that discusses non-fundamental objects is that of Ladyman and Ross (2007) discussed in Chapter 1. They propose that non-fundamental objects, like fundamental objects, are real patterns (Dennett 1991). Yet, non-fundamental objects are real patterns which *represent* the fundamental real patterns corresponding to elementary particles and spacetime points. The notion of a representational real pattern does little to resolve the mystery, however. The chief question is how, if not composition, do fundamental objects relate to non-fundamental objects?

The approach I prefer is one in which non-fundamental objects, which may be regarded as individuals in the traditional sense, emerge as part of an "effective theory" that serves certain of our interests. In particular, the structures of state-relations provided by quantum physics give way to a more traditional structure of objects that can be dissociated from their structures as we change the scale of our investigations. Hence, biological entities, for example, emerge from the global quantum structure once we get to the scale at which biological description is relevant.

6.4 Future physics: looking beyond GR and QM

Even if we restrict the scope of the project to theories of fundamental physics, there is a clear limitation to the preceding arguments. It consist of the fact that none of the theories considered represent the current state of the art in physics. In particular, while I have addressed relativity theory and (non-relativistic) quantum mechanics, I have yet to take a stand on quantum field theory (QFT). QFT is important because it is the best candidate for a theory with the sort of virtues scientific realists have in mind when they speak of our "best" or "most successful" theories. It is also a theory of fundamental physics if anything is, and hence should—according to the version of naturalism described in the introduction—be the ultimate source of metaphysical insights. Properly addressing the myriad interpretative issues surrounding QFT would require far more space than is available here, but a few brief comments on its relevance to this project are in order.

6.4.1 The problem of the ontology of QFT

Despite its status as an extremely successful theory, there is widespread disagreement concerning the correct ontology for QFT. Initially, it may seem as though QFT recommends a particle ontology. After all, the Standard Model of *particle physics* exists in the framework QFT provides. In addition, the formalism of QFT contains "creation," "annihilation," and "number" operators that are naturally interpreted as referring to the creation, annihilation and number of *particles*.

There are reasons to resist the particle interpretation, however. Some of these reasons are familiar from non-relativistic QM. For example, particles in QFT cannot be localized to a region of spacetime. Moreover, the possibility of absolutely indiscernible particles discussed in Chapter 5 remain in QFT; on most interpretations, two particles may have all the same properties and relations. QFT also introduces new problems for a particle ontology. One problem is that the occupation number representation is only one among many bases that may be used to represent the state of system in QFT. If another basis is chosen, however, the system will in general not be an eigenstate of occupation number, and hence, the system will be in *a superposition* of distinct occupation numbers. Needless to say, it is not obvious how the particle interpretation can understand superpositions of, say, a 3-electron/3-positron state and a 17-election/17-positron state.

The are further problems with the particle interpretation. The relativistic vacuum state, for example, is one in which the expectation values for various physical quantities are non-zero. This suggests that the vacuum is not a state devoid of interesting physical properties, but it is hard to see how this is possible when, on the particle interpretation, there would be *nothing* present. Another problem with the relativistic vacuum is that according to the Rindler representation—which takes the frame of reference of a uniformly accelerated observer—*there are particles present.* It would certainly be puzzling if, in a relativistic theory, what exists fundamentally depends on which frame of reference we take.

These problems motivate some to adopt a field interpretation of QFT, but problems emerge here as well. Most importantly, QFT describes an *operator-valued* field $\hat{\Phi}(\mathbf{x}, t)$. This is to be contrasted with classical fields which typically consist of the assignment of a particular scalar, vector, or tensor to each point in space(-time). While the operator-valued fields of QFT are formally similar to their classical counterparts (and hence deserve the title "fields"), they seem to represent something quite different *physically.* In particular, classical fields seem to attribute specific properties to each point in space(time), while operators tell us only the *kind* of property under consideration. Thus, a field of operators seems to give us only a distribution of *determinables* rather than a configuration of *determinates* that is required of a proper physical field (Teller 1997).

There are other ontological options for the interoperation of QFT, but they either (1) fail to provide a distinct physical ontology or (2) the ontology they provide requires nothing less than specifying the state of the entire world/universe/multiverse.

6.4.2 The Implications for Structural Realism

The problem of the ontology of QFT may be taken to recommend the kind of structural realism advocated above in at least two ways. First, one can argue that the only ontology compatible with QFT is *structuralist* in the sense that it requires taking the whole to be fundamental. Structures, or networks of relations, are irreducible to the objects—particles or field magnitudes—they are taken to contain.[3] Second, QFT may be taken to support the idea that accurately describing the structure of the world—as opposed to an ontology of individual objects—is what accounts for the success of theory.

Structuralist ontology and QFT

The previous section cast doubt on the traditional ontological options for QFT, particles and fields. Thus, the question remains: what sort of a world does QFT describe? Structural realism provides a promising option. Recall that the ontic structural realist claims that the reason we can only *know* the structure of the world is that all there *is* to the world is structure. From this perfective, the task of determining the ontology of any theory involves locating the relevant physical structure(s) that are described. These structures themselves—understood as physical relations and the places between them—comprise the ontology of the theory. The idea that QFT recommends a structuralist ontology has been advocated by Lyre (2004, 2012), for example.[4]

Structural realism and ontological ambiguity

Independent of the prospects for a structuralist interpretation of QFT, there is another connection between structural realism and problems concerning the ontology

---

[3] Of course, if objects are understood as places in structures, there is a sense in which structures can be reduced to objects. But, this wouldn't be a reduction so much as a redescription of structure in terms of its places.

[4] Lyre emphasizes *group* structures, so it's not clear that the ontological gloss I offer is consistent with his proposal.

of QFT. According to structural realism, the reason our best scientific theories are successful is that they accurately describe (only) the structure of the world. For example, above it was claimed that what accounts for the success of Newtonian mechanics is that it correctly describes the structure of spacetime at the classical limit. From this perspective, earlier theories (like Newton's) contain surplus content in the form of their ontological posits. The objects and properties they are committed to do not exist and, hence, these posits represent a deficiency of the theory.

QFT, however, does not wear its ontology on its sleeve. The successful application of the theory—in its predictions and explanations—involves different ontological pictures in different contexts. As we have seen, there is difficulty in taking either the particle or field interpretation as giving us *the* ontology of QFT. We may take this to recommend antirealism, or to count against the possibility of a "pristine" interpretation (Ruetsche 2011). For the structural realist, however, the ontological quandary of QFT is to be expected. If it is the structure of a theory that accounts for its success, then it is not surprising that a theory could be successful without unequivocal ontological commitments.

If structural realism is correct, the goal of a scientific theory is to capture only the structure of the world. Hence, we should expect science to develop theories which are sanguine about non-structural ontology. QFT provides an example of just such a theory. Particles and fields function as useful metaphors for the application of the theory, but neither should be taken to provide the ontology of the world. Strictly speaking, all that QFT tells us are certain details about the physical structure of the world. I have argued above that if science can only tell us about structure, naturalism recommends positing no more than this. But the argument here is distinct from that of the previous section. There it was claimed that QFT provides a structuralist ontology; here the claim is that QFT is precisely the sort of theory one would expect if science can tell us only about structure.

It is worth noting that QFT is not unequivocal when it comes to the structure of the world either. The problem of unitarily inequivalent representations discussed in section 2.7.3 seems to show that there are several candidates for *the* structure of the world according to QFT. Yet, as I emphasized there, the structural realist needn't claim the world has one and only one structure. There are several ways of making sense of the idea that the world has *multiple* structures, even when those structures seem to be incompatible. One option is to appeal to *patterns*—as done by Ladyman and Ross (2007); Dennett (1991); Resnik (1997). Multiple patterns may be found in a single set of data, even when those patterns are incompatible (the "duck-rabbit" provides a simple example). Another way to make sense of multiple structures is by appeal to partial isomorphism (da Costa and French 2003). Two models may each be partially isomorphic with the world without being isomorphic (or even partially isomorphic) with each other. Intuitively, each model may capture different aspects of the structure of the world.

6.5 Summary

The goal of this dissertation has been to present the framework for a defensible version of structural realism. In this final section I have suggested that contemporary physics, and QFT in particular, is likely to provide further evidence for structural realism. This reinforces the central claim of the dissertation that structural realism finds strong support in physics.

REFERENCES

Adams, R. M. (1979). Primitive thisness and primitive identity. *The Journal of Philosophy*, **76**(1), pp. 5–26.

Ainsworth, P. M. (2009). Newman's objection. *The British Journal for the Philosophy of Science*, **60**(1), pp. 135–171.

Alexanderson, G. L. (2006). About the cover: Euler and Königsberg's bridges: a historical view. *Bulletin of the American Mathematical Society*, **43**(4), pp. 567–573.

Bain, J. (2011). Category-theoretic structure and radical ontic structural realism. *Synthese*, pp. 1–15.

Bain, J. and J. D. Norton (2001). What should philosophers of science learn from the history of the electron. *Histories of the electron: The birth of microphysics*, pp. 451–465.

Boyd, R. (1990). Realism, approximate truth, and philosophical method. *Scientific theories*, **14**, pp. 355–391.

Brading, K. and E. Landry (2006). Scientific structuralism: presentation and representation. *Philosophy of Science*, **73**(5), pp. 571–581.

Butterfield, J. (1989). The hole truth. *The British Journal for the Philosophy of Science*, **40**(1), pp. 1–28.

Cameron, R. (2008). Truthmakers and ontological commitment: or how to deal with complex objects and mathematical ontology without getting into trouble. *Philosophical Studies*, **140**, pp. 1–18.

Cartwright, N. (1983). *How the Laws of Physics Lie.* Clarendon, Oxford.

Caulton, A. and J. Butterfield (2012). Symmetries and paraparticles as a motivation for structuralism. *The British Journal for the Philosophy of Science*, **63**(2), pp. 233–285.

Chakravartty, A. (2012). Ontological priority: the conceptual basis of non-eliminative, ontic structural realism. In Dean Rickles, E. L. (ed.) *Structural Realism: Structure, Object, and Causality*, chapter 10, pp. 187–206. Springer.

Chalmers, D., D. Manley, and R. Wasserman (eds.) (2009). *Metametaphysics: New Essays on the Foundations of Ontology.* Clarendon Press, Oxford.

Cruse, P. (2005). Ramsey sentences, structural realism and trivial realization. *Studies In History and Philosophy of Science Part A*, **36**(3), pp. 557 – 576.

da Costa, N. C. and S. French (1990). The model-theoretic approach in the philosophy of science. *Philosophy of Science*, **57**(2), pp. 248–265.

da Costa, N. C. and S. French (2003). *Science and Partial Truth.* Oxford University Press, Oxford.

Dasgupta, S. (2011). The bare necessities. *Philosophical Perspectives*, **25**(1), pp. 115–160.

Demopoulos, W. and M. Friedman (1985). Bertrand Russell's The Analysis of Matter: its historical context and contemporary interest. *Philosophy of Science*, **52**(4), pp. 621–639.

Dennett, D. C. (1991). Real patterns. *The Journal of Philosophy*, **88**(1), pp. 27–51.

Dipert, R. (1997). The mathematical structure of the world: the world as graph. *The Journal of Philosophy*, **94**(7), pp. 329–358.

Earman, J. (MS). Understanding permutation invariance in quantum mechanics.

Earman, J. and A. Fine (1977). Against indeterminacy. *The Journal of Philosophy*, **74**(9), pp. 535–538.

Earman, J. and J. Norton (1987). What price spacetime substantivalism. *British Journal for the Philosophy of Science*, **38**, pp. 515–525.

Esfeld, M. (2004). Quantum entanglement and a metaphysics of relations. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **35**(4), pp. 601–617.

Esfeld, M. and V. Lam (2008). Moderate structural realism about space-time. *Synthese*, **160**, pp. 27–46.

Esfeld, M. and V. Lam (2010). Holism and structural realism. In Vanderbeeken, R. and B. D. Hooghe (eds.) *Worldviews, Science and Us: Studies of Analytical Metaphysics. A Selection of Topics From a Methodological Perspective.*, pp. 10–31. World Scientific Publishing, Singapore.

Esfeld, M. and V. Lam (2011). Ontic structural realism as a metaphysics of objects. In *Scientific structuralism*, pp. 143–159. Springer.

Feyerabend, P. (1962). Explanation, reduction, and empiricism. In Feigl, H. and G. Maxwell (eds.) *Scientific Explanation, Space, and Time*, volume III of *Minnesota Studies in the Philosophy of Science*, pp. 28–97. University of Minnesota Press, Minneapolis.

Field, H. (1973). Theory change and the indeterminacy of reference. *The Journal of Philosophy*, **70**(14), pp. 462–481.

French, A. P. (1968). *Special relativity.* CRC PressI Llc.

French, S. (2003). A modeltheoretic account of representation (or, I don't know much about art... but I know it involves isomorphism). *Philosophy of Science*, **70**(5), pp. 1472–1483.

French, S. (2012a). The resilience of laws and the ephemerality of objects: can a form of structuralism be extended to biology? In *Probabilities, Laws, and Structures*, pp. 187–199. Springer.

French, S. (2012b). 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), pp. 121 – 136.

French, S. and J. Ladyman (2003). Remodelling structural realism: quantum physics and the metaphysics of structure. *Synthese*, **136–141**, pp. 31–56.

French, S. and K. McKenzie (2012). Thinking outside the toolbox: towards a more productive engagement between metaphysics and philosophy of physics. *European Journal of Analytic Philosophy*, **8**(1), pp. 42–59.

Frigg, R. and I. Votsis (2011). Everything you always wanted to know about structural realism but were afraid to ask. *European journal for philosophy of science*, **1**(2), pp. 227–276.

Giere, R. (1988). *Explaining Science: A Cognitive Approach.* The University of Chicago Press, Chicago.

Hacking, I. (1982). Experimentation and scientific realism. *Philosophical Topic*, **13**(1), pp. 71–87.

Hacking, I. (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science.* Cambridge University Press, Cambridge.

Hacking, I. (1985). Do we see through a microscope? In Churchland, P. and C. Hooker (eds.) *Images of Science: Essays on Realism and Empiricism*, chapter 7. The University of Chicago Press, Chicago.

Halvorson, H. (2012). What scientific theories could not be. *Philosophy of Science*, **79**(2), pp. 183–206.

Hardin, C. L. and A. Rosenberg (1982). In defense of convergent realism. *Philosophy of Science*, **49**(4), pp. 604–615.

Healey, R. (1995). Susbtance, modality and spacetime. *Erkenntnis*, **42**(3), pp. 287–316.

Heil, J. (2005). *From an Ontological Point of View.* Clarendon Press.

Horgan, T. and M. Potrč (2000). Blobjectivism and indirect correspondence. *Facta Philosophica*, **2**, pp. 249–270.

Horgan, T. and M. Potrč (2008). *Austere Realism.* The MIT Press, Cambridge, Mass.

Horwich, P. (1998). *Truth.* Oxford: Blackwell, second edition.

Ketland, J. (2004). Empirical adequacy and ramsification. *The British Journal for the Philosophy of Science*, **55**(2), pp. 287–300.

Kitcher, P. (1993). The advancement of science: Science without legend, objectivity without illusions.

Kripke, S. (1980). *Naming and Necessity.* Blackwell Publishing Ltd, Oxford.

Kuhn, T. (1962/1996). *The Structure of Scientific Revolutions.* The University of Chicago Press, Chicago, third edition.

Ladyman, J. (2011). Structural realism versus standard scientific realism: the case of phlogiston and dephlogisticated air. *Synthese*, **180**(2), pp. 87–101.

Ladyman, J. and D. Ross (2007). *Every Thing Must Go: Metaphysics Naturalized.* Oxford University Press, Oxford.

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

Laudan, L. (1981). A confutation of convergent realism. *Philosophy of Science*, **48**(1), pp. 19–49.

Laudan, L. (1984a). Explaining the success of science: beyond epistemic realism and relativism. *Science and reality*, pp. 83–105.

Laudan, L. (1984b). Realism without the real. *Philosophy of Science*, **51**(1), pp. 156–162.

Lawvere, W. and S. Shanuel (1997). *Conceptual Mathematics*. Cambridge University Press, Cambridge.

Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, **61**(4), pp. 343–377.

Lewis, D. (1986). *Philosophical Papers*, volume II. Oxford University Press.

Lowe, E. (1998). Entity, identity and unity. *Erkenntnis*, **48**, pp. 191–208.

Lowe, E. (2007). Sortals and the individuation of objects. *Mind and Language*, **22**(5), pp. 514–533.

Lyre, H. (2004). Holism and structuralism in U(1) gauge theory. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, **35**(4), pp. 643 – 670.

Magnus, P. and C. Callender (2004). Realist ennui and the base rate fallacy. *Philosophy of Science*, **71**(3), pp. 320–338.

Maudlin, T. (1988). The essence of space-time. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*, **Volume Two: Symposia and Invited Papers**, pp. 82–91.

Maudlin, T. (1990). Substances and space-time: what Aristotle would have said to Einstein. *Studies In History and Philosophy of Science Part A*, **21**(4), pp. 531–561.

Maudlin, T. (2012). *Philosophy of Physics: Space and Time*. Princeton University Press.

Maxwell, G. (1970a). Structural realism and the meaning of theoretical terms. In Winokur, S. and M. Radner (eds.) *Analyses of Theories and Methods in Physics and Psychology*, volume 4 of *Minnesota Studies in the Philosophy of Science*, pp. 181–192. University of Minnesota Press, Minneapolis.

Maxwell, G. (1970b). Theories, perception and structural realism. In Colodny, R. (ed.) *The Nature and Function of Scientific Theories*, volume 4 of *University of Pittsburg Series in the Philosophy of Science*, pp. 3–34. University of Pittsburg Press, Pittsburg.

Maxwell, G. (1972). Scientific methodology and the causal theory of peception. In Feigl, H., H. Sellars, and K. Lehrer (eds.) *New Readings in Philosophical Analysis*, pp. 148–177. Appleton-Century Crofts, New York.

Melia, J. and J. Saatsi (2006). Ramseyfication and theoretical content. *The British Journal for the Philosophy of Science*, **57**(3), pp. 561–585.

Morganti, M. (2004). On the preferability of epistemic structural realism. *Synthese*, **142**, pp. 81–107.

Musgrave, A. (1985). The ultimate argument for scientific realism. In Churchland, P. and C. Hooker (eds.) *Images of Science: Essays on Realism and Empiricism.* The University of Chicago Press, Chicago.

Newman, M. H. A. (1928). Mr. Russell's "Causal theory of perception". *Mind*, **37**(146), pp. 137–148.

Okun, L. B. (1989). The concept of mass. *Physics today*, **42**(6), pp. 31–36.

Paul, L. (2006). Coincidence as overlap. *Noûs*, **40**(4), pp. 623–659.

Paul, L. A. (2002). Logical parts. *Noûs*, **36**(4), pp. 578–596.

Pooley, O. (2006). Points, particles, and structural realism. In Rickles, D., S. French, and J. Saatsi (eds.) *The Structural Foundations of Quantum Gravity*, pp. 83–120. Oxford University Press, Oxford.

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

Psillos, S. (2006). The structure, the whole structure, and nothing but the structure? *Philosophy of Science*, **73**(5), pp. 560–570.

Putnam, H. (1966). What theories are not. In Ernest Nagel, P. S. and A. Tarski (eds.) *Logic, Methodology and Philosophy of Science Proceeding of the 1960 International Congress*, volume 44 of *Studies in Logic and the Foundations of Mathematics*, pp. 240 – 251. Elsevier.

Putnam, H. (1975a). *Mathematics, Matter and Method.* Cambridge University Press, Cambridge.

Putnam, H. (1975b). *Philosophical Papers*, volume 2: Mind, Language and Reality. Cambridge University Press, Cambridge.

Quine, W. V. (1948). On what there is. *The Review of Metaphysics*, **2**(5), pp. 21–38.

Quine, W. V. (1951). Ontology and ideology. *Philosophical Studies*, **2**(1), pp. 11–15.

Quine, W. V. (1968). Ontological relativity. *the Journal of Philosophy*, **65**(7), pp. 185–212.

Quine, W. V. (1981). *Theories and Things.* Harvard University Press.

Resnik, M. D. (1997). *Mathematics as a Science of Patterns.* Oxford University Press, Oxford.

Roberts, B. W. (2011). Group structural realism. *The British Journal for the Philosophy of Science*, **62**(1), pp. 47–69.

Ross, D. (2008). Ontic structural realism and economics. *Philosophy of Science*, **75**(5), pp. 732–743.

Ruetsche, L. (2011). *Interpreting Quantum Theories.* OUP Oxford.

Russell, B. (1911). On the relations of universals and particulars. *Proceedings of the Aristotelian Society*, **12**, pp. 1–24.

Russell, B. (1927/1992). *The Analysis of Matter.* Routledge, New York.

Sandin, T. (1991). In defense of relativistic mass. *American Journal of Physics*, **59**, p. 1032.

Saunders, S. (1993). To what physics corresponds. In French, S. and H. Kamminga (eds.) *Correspondence, invariance and heuristics*, volume 148 of *Boston Studies in the Philosophy and History of Science*, pp. 295–325. Springer Netherlands.

Saunders, S. (2003). Physics and Leibniz's principles. In Brading, K. and E. Castellani (eds.) *Symmetries in physics: Philosophical reflections.* Cambridge University Press.

Saunders, S. (2006a). Are quantum particles objects? *Analysis*, **66**(1), pp. 52–63.

Saunders, S. (2006b). On the explanation for quantum statistics. *Studies In History and Philosophy of Science Part B*, **37**(1), pp. 192 – 211.

Schaffer, J. (2010). Monism: The priority of the whole. *Philosophical Review*, **119**(1), pp. 31–76.

Shapiro, S. (1997). *Philosophy of Mathematics: structure and ontology.* Oxford University Press, Oxford.

Stachel, J. (1989). Einstein's search for general covariance, 1912–1915. In Howard, D. and J. Stachel (eds.) *Einstein and the History of General Relativity*, pp. 63–100.

Stachel, J. (2002). The relations between things versus the things between relations: the deeper meaning of the hole argument. In Malament, D. B. (ed.) *Reading Natural Philosophy/ Essays in the History and Philosophy of Science and Mathematics*, pp. 231–266. Open Court, Chicago and LaSalle, IL.

Strawson, P. F. (1950). On referring. *Mind*, **59**(235), pp. 320–344.

Suppe, F. (1989). *The Semantic View of Theories and Scientific Realism.* University of Illinois Press, Urbana and Chicago.

Suppes, P. (2002). *Representation and Invariance of Scientific Structures.* CSLI Publications, Stanford.

Tegmark, M. (2008). The mathematical universe. *Foundations of Physics*, **38**, pp. 101–150.

Teller, P. (1997). *An Interpretive Introduction to Quantum Field Theory.* Princeton series in physics. Princeton University Press.

van Fraassen, B. (1980). *The Scientific Image.* Oxford University Press, Oxford.

van Fraassen, B. (2008). *Scientific Representation.* Oxford University Press, Oxford.

Wallace, D. (2010). Decoherence and ontology. In Saunders, S., J. Barrett, A. Kent, and D. Wallace (eds.) *Many Worlds? Everett, Quantum Theory, and Reality*, pp. 53–72. Oxford University Press, Oxford.

Wigner, E. (1939). On unitary representations of the inhomogeneous Lorentz group. *Annals of Mathematics*, pp. 149–204.

Wikipedia (2013). Seven Bridges of Königsberg — Wikipedia, The Free Encyclopedia. [Online; accessed 3-September-2013].

Worrall, J. (1989). Structural realism: the best of both worlds? *Dialectica*, **43**(1-2), pp. 99–124.

Worrall, J. (2007). Miracles and models: why reports of the death of structural realism may be exaggerated. *Royal Institute of Philosophy Supplements*, **61**, pp. 125–154.

Wüthrich, C. (2009). Challenging the spacetime structuralist. *Philosophy of Science*, **76**(5), pp. 1039–1051.

Zahar, E. G. (2004). Ramseyfication and structural realism. *Theoria*, **19**(1), pp. 5–30.