

Toward a Pragmatic Ontology of Scientific Concepts

Olin M. Robus

A dissertation

submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2018

Reading Committee:

Andrea Woody, Chair

Arthur Fine

Cass Weller

Joseph Rouse

John Manchak

Program Authorized to Offer Degree:

Philosophy

©Copyright 2018
Olin M. Robus

University of Washington

Abstract

Toward a Pragmatic Ontology of Scientific Concepts

Olin M. Robus

Chair of the Supervisory Committee:
Professor Andrea Woody
Philosophy

I argue that current projects in ‘naturalized metaphysics’ fail to be properly naturalistic, and thereby fail in their stated aim to take one’s metaphysics from science. I argue that naturalism must involve the idea of taking science seriously, and that this can only be spelled out in terms of taking not only the theories of science seriously, but also its practice and its socio-linguistic situatedness seriously as well. This accords with naturalism because not doing so draws an artificial (non-natural) distinction between the epistemic products of science as *essence*, and its socio-linguistic and practical features as *accidents*.

The picture of naturalism which falls out of this is a form of pragmatism. Once this is spelled out, the question becomes, what is the appropriate attitude for the pragmatist/naturalist to have toward ontology? It is not the same as that of traditional metaphysics, since such an attitude (that metaphysical theorizing can make positive epistemic contributions) requires *a priori* commitments which themselves are (1) not subject to empirical review, and (2) are anthropocentric and so potentially distorted. But the pragmatist/naturalist will not go the opposite way, and say that metaphysics is meaningless either, since (true to their pragmatic commitments) ontol-

ogy does things for the scientist and the layperson. It is a mistake to reject it wholesale, since the task of asserting what there is, exposing such an ontology to criticism, both empirical and logical, and revising it, has both small-scale and large-scale consequences. On the small-scale, ontologies are a guide for thinking—scientists, engineers, and lay people can and do use models of existence to navigate occurrent problems in their day-to-day lives. Whether this is posing research hypotheses, developing protocols for generating new materials, or cooling off a cup of coffee, ontology plays a role. Large-scale consequences are more weighty, and more difficult to see. These large-scale consequences have to do with the aims and values which individuals and societies possess, and their interrelation with ontology. This certainly was at the forefront of the earliest ontologies—the Epicureans and Stoics built their moral philosophy on the basis of what ontology they thought was correct. This sort of practice goes on, unabated, today, though with virtually no overt acknowledgement of this important interdependency.


How, for the pragmatist, do we make sense of ontological talk, if we are to eschew traditional metaphysics? Ontology—the naturalist/pragmatist declares—is a tool. Ontology helps us do things, whether it be predict behavior, understand phenomena, blame or forgive someone, and hope or despair about a life after this one (for example). Building an ontology is about building our own concepts, and this, in turn is a negotiation between our beliefs, experiences, and commitments, and the beliefs, experiences, and commitments of with whom we're discursively engaged.

Contents

Copyright	i
Abstract	ii
Table of Contents	iv
Acknowledgments	vii
Introduction	1
0.0 Naturalism & Metaphysics	1
0.1 What is Scientific Metaphysics? What is Metaphysics?	5
0.1.1 Four Kinds of Naturalism	15
1 Putnam, Stein, & Space-Time	22
1.0 Introduction	22
1.1 Time and Special Relativity	23
1.1.1 Putnam, Eternalism and the Relation R	25
1.1.2 Arthur Prior and <i>a priori</i> Scruples	30
1.1.3 Howard Stein and the Equivocality of Theory	33
1.1.4 What Can Settle the Choice Between R and R' ?	36
1.2 Objections	39
1.2.1 Does General Relativity Offer a Better Metaphysical Basis?	40
1.2.2 Symmetry, Interpretation, and Fundamentality	42
1.2.3 Is the <i>a priori</i> Anti-Naturalist?	47
2 Naturalism & the <i>a priori</i>	49
2.0 Introduction	49
2.1 Statements of Rejection	52
2.2 Taking Science Seriously	55
2.3 Why Does Naturalism Reject the <i>a priori</i> ?	57
2.3.1 The Role of the <i>a priori</i>	57
2.3.2 An Epistemological Foundation	62
2.4 Two Principles of Naturalized Metaphysics	63
2.4.1 Science and Conceptual Revolution	65
2.4.2 Naturalism and the ‘Quinean Realization’	69
2.4.3 A Naturalistic Substitute	71

3	Ontological Structuralism	74
3.0	Introduction	74
3.1	A Sophisticated Naturalism	75
3.1.1	Origins of Structural Realism	75
3.1.2	Structure as Metaphysics	78
3.1.3	Mathematical Fixed-Points	86
3.2	Taking Structuralism Seriously	88
3.2.1	Two Naturalistic Failures	89
3.3	Between Math and the World	99
3.3.1	Pythagoreanism or ‘Real’ Structure?	99
3.4	Conclusion	103
4	Pragmatism & Scientific Metaphysics	104
4.0	Introduction	104
4.1	Practice, Concepts, and Pragmatism	108
4.1.1	The Turn to Practice	108
4.1.2	Metalinguistic Negotiation	120
4.1.3	A Pragmatist View of Concepts	123
4.2	Metaphysics in Practice: Astrobiology and Life	125
4.2.1	Definitions of Life, the Problem	126
4.2.2	A Pragmatic Approach	133
4.2.3	Jovian Tornadoes	137
4.3	Conclusion: Why ‘Metaphysics’?	144
	Appendices	146
	Appendix A	146
	Appendix B	150
	Bibliography	151

For Patricia.



Acknowledgments

For better or worse, this dissertation is the product of 9 years of graduate study, and a full 32 years of experience and education. I have been extraordinarily lucky to have many wonderful teachers, friends, mentors, and role-models to help me get to this point. I could not have done this without their guidance and inspiration, and I can only hope that this dissertation is an adequate token of my appreciation.

My first and most profound debt of gratitude is to my family, who at an early age inculcated in me a desire to learn, and taught me an admiration and respect for inquiry. I am the product of my environment and upbringing, and insofar as I have any redeeming qualities, intellectual or otherwise, I owe them to the love, patience, and quiet brilliance of my mother and father.

Another debt of gratitude is owed to the University of Washington Department of Philosophy, and its faculty, staff, and graduate students. My philosophical education truly began in graduate school, and were it not for interactions with my fellow graduate students and departmental faculty, I might have very different beliefs and ideas than I do today. For that I am thankful. In particular I am indebted to Jon Rosenberg, my friend and colleague. I owe much of my philosophical worldview to my many discussions with Jon over the past decade.

I am indebted to the undergraduates in the many courses I have taught in these past years at the University of Washington. Their skepticism, confusion, and enthusiasm have been a source

of frustration, consternation, and inspiration. They say nothing focuses the mind quite like the prospect of death, but the prospect of facing 30 dubious and intelligent undergrads surely comes close. I have learned more philosophy in my discussions with my students than from any talk I've ever attended, or any paper or book I've ever read. To all my students, thank you.

I would also like to thank my dissertation committee for their interest, support, and engagement with this project. Their comments, feedback, and questions have improved my own arguments considerably, and while there will be inevitable points of disagreement, their advice, insight, and wisdom have been most valuable. In particular I must thank my committee chair, Andrea Woody, for her help and guidance in the writing of this dissertation.

For their help reading through drafts of this document, and for their careful editing suggestions, my thanks to Ben Robinson and Matt Robus. Their attention to detail and willingness to endure my prose have improved this work immeasurably. While they have surely caught the most egregious infelicities in the following chapters, whatever errors remain are my own.

Finally, for her patience, for her support, and for not letting me quit, to my wife Patricia I owe my gratitude and love. Thanks for waiting.

Introduction

0.0 Naturalism & Metaphysics

What role should science play in our inquiry and speculation about the world? This question is, I think, the central question to which the philosophical doctrine of naturalism seeks an answer. The prevalence of ‘naturalisms’ in the academy, both in the sciences and humanities, speaks to the importance of this question, especially as science occupies a larger and larger social and cultural position—all the while being subject to political and ideological antipathy.

Naturalism, I take it, consists most broadly in the injunction to ‘take science seriously’. This notion provides a hint at how we should answer the question above: whatever the details, science should play a central and significant role. Beyond this, varieties of naturalism will undoubtedly diverge on whether that central role is exclusive, or admits of fellow adjudicators, and whether the role it plays is to *rule out* speculation about the world, or instead if it just judges what gets *ruled in*.

However, while I see the outlines of these efforts to answer the question as largely well-motivated, attempting to reconcile the role of science and what we know and think about the world has met significant problems. This dissertation will argue that these difficulties are due in large part to the fact that this question has been misunderstood. The consequence of this misunderstanding is that the current ‘standard’ notions of naturalism fail by their own lights (i.e. their positions are undermined by their own self-proclaimed naturalist scruples).

The misunderstanding of the question, and the consequent difficulties in developing a coher-

ent understanding of how to make sense of science's place within a world it is thought to disclose, trace back to a fundamental ambiguity in the question itself, an ambiguity which serves to make a certain, mistaken, answer to the question seem attractive and correct.

The ambiguity is in the use of the term 'science'. There are two alternative understandings that are relevant here. On one hand we might understand it to mean more specifically *scientific knowledge*. This is commonly the way it is used in strong statements of naturalism: "Naturalism is the thesis that scientific knowledge exhausts what one can know about the world." Thus the picture of a naturalized metaphysics based on this answer tends to emphasize scientific theories and their content (together with the implicit model of scientific knowledge as the *having* of this content). On this understanding the connection between metaphysics and science is that the scientific contents are to *replace* the *a priori*/metaphysical contents of metaphysics.

On the other hand, 'science' can be used to refer to *scientific practice*. This sense is the one from which 'the naturalist spirit' most readily springs: the idea that science is valuable as a mode of inquiry is best motivated by thinking about the *activity* of science. The role science plays in our inquiry and speculation about the world is to serve as a *practical model* of inquiry, and is unconcerned with the question of what a theory says, or what its contents are. As a consequence this understanding of science is ill-suited to *replace* the content of *a priori* metaphysics with 'scientific content', and instead (I will argue) casts the notion of a metaphysics 'having' content of any sort into doubt.

It is the thesis of this dissertation that (1) it is this latter use of the term 'science' which should be taken to capture and frame naturalism, (2) current projects in naturalized metaphysics are therefore naturalisms in name only, and (3) that by taking the *practice* of science seriously, we can come to articulate a notion of metaphysical inquiry that is adequately suited to a commitment to take science seriously—warts and all.

In the rest of the introduction I will give a general account of the contemporary notion of naturalized metaphysics and its relations to the wider notions of naturalism and metaphysics. This will involve some discussion of traditional analytic metaphysics, science, and the philosophy of

science. The hope is that by the end of the introduction we'll have a sense of the conceptual terrain surrounding the conjoined notion of *science and metaphysics* and have a better handle on what the term 'metaphysics' is supposed to pick out in the following discussion. I will introduce a general taxonomy of naturalisms which answer the question of science's role in (metaphysical) inquiry. Briefly, these naturalisms are: Logical Positivism, Naïve Naturalism, Sophisticated Naturalism, and Pragmatism. In the following chapters I will then examine these naturalisms in detail (save for Logical Positivism which I discuss in this introduction but pursue no further).

THE PLAN FOR THE DISSERTATION

Naïve naturalism is the view that science can simply *dictate* the correct metaphysical theory, independent from any non-scientific/pre-scientific ingredients. In the first chapter I develop this view, and consider an example of a scientific theory (Special Relativity) which appears to decisively settle a metaphysical question (the reality/determinateness of future events). It will be shown that such a naturalism ultimately fails: it cannot provide a coherent metaphysical account without a residue from the *a priori* ingredients it sought to displace.

In Chapter 2 I develop the argument (usually implicit) for why *a priori* principles are to be avoided by the naturalist in the first place. Ultimately, the reasoning turns on the yoking together of two doctrines: Metaphysical Realism and Epistemological Naturalism. By accepting both the naturalist is committed to (a) there being a fact of the matter about the external world, but also (b) a rejection of the idea that our biologically inherited belief-forming and general cognitive mechanisms are reliable outside of the narrow range of our evolutionary history. The task for the naturalist aiming to overcome the problems faced by the naïve approach (whom I will call 'sophisticated') is then to identify a reliable source for justifying beliefs about the external world which are not vulnerable to the same considerations of contingency and parochialism as our innate concepts. Granted that *a priori* judgments are unreliable outside our narrow evolutionary history and therefore fail *as a priori* judgments, the sophisticated naturalist must show that the claims of science are reliable beyond the narrow range of our evolutionary history.

In Chapter 3 I argue that the sophisticated naturalist also ultimately fails in this task by the lights of naturalism properly construed. I argue that the requisite sophistications commit one to a view of science which undermines the motivations for a naturalized metaphysics as proposed. That is, attempts at replacing traditional metaphysical content by scientific surrogates developed under the banner of naturalism ultimately violate those very naturalist scruples. This is not to say they are unintelligible in their own right—indeed they are perfectly coherent metaphysical theories—but they cannot be justified by the lights of science.

If Chapters 1, 2, and 3 constitute the critical portion of my argument, then the final chapter introduces my positive account of naturalized metaphysics. That is, I advance a positive argument for a properly naturalized metaphysics. Instead of seeking to simply replace the *contents* of metaphysics with more scientific contents, I argue the lessons learned by taking science seriously inform our understanding of what metaphysics *as a practice* itself is supposed to do.¹ To this end, in Chapter 4 I advocate for a form of pragmatism, which attends carefully to the embodied nature of scientific practice, and re-construes questions about ontology as practical questions about how and why to use particular concepts. The role of metaphysics on this view is to evaluate the adequacy of any particular concepts in their application. This account of metaphysics amounts to a repudiation of the typically assumed picture of metaphysics as seeking to put our conceptual scheme into the ‘true’ 1–1 relation with ‘the world’.

To finish with a concrete case, I turn to the field of astrobiology and look at current debates about the definition of ‘life’. Rather than seek to justify necessary and sufficient conditions by which to define it, I instead argue that ‘life’ is not a natural kind. The upshot is that it becomes clear the question ‘what is alive?’ is not a question of identifying properties more clearly, but instead a question of what we should countenance as alive, and that in turn depends on what problems and questions we take to be most salient to us as inquirers. This attitude then ramifies for the practice of astrobiology particularly, and the sciences more generally.

¹That is, to naturalize metaphysics is to change the aims and methods of metaphysics, not merely the sources and content. Science gives us tools for how to think about the world and our epistemic relation to it, it does not simply and merely describe a static, eternal object called ‘the world’.

0.1 What is Scientific Metaphysics? What is Metaphysics?

The term ‘metaphysics’ has been used in many different ways, to serve many different purposes. Famously, the term arose as a bookkeeping device, to describe where Aristotle’s text (the topic of which he called ‘first philosophy’) was to be placed (after *Physics*). It has taken on a life of its own since then, used by bookstores to classify less than reputable claims about crystals and aliens, and used sometimes to describe philosophy generally. Take William James’ summing up of the term: “Metaphysics means only an unusually obstinate attempt to think clearly and consistently.” (James, 1892/1920, p. 461) As the term ‘metaphysics’ is employed currently in academic philosophy, James’ definition clearly casts the net too wide—surely epistemology, ethics, aesthetics, logic, etc. all strive (in an unusually obstinate manner) to think clearly and consistently. But saying that the study of metaphysics is more narrowly circumscribed than James claims leaves the field rather wide open.

One avenue to take in pinning down the content of metaphysics is to return to Aristotle’s notion of ‘first philosophy’. First philosophy, as the name suggests, treats of fundamental, foundational, first, principles; those that ground and guarantee everything else. In fact, immediately following the above quote James hints at narrower description of metaphysics, one that is more identifiable with this notion:

“A geologist’s purposes fall short of understanding Time itself. A mechanist need not know how action and reaction are possible at all... [But] as soon as one’s purpose is the attainment of the maximum possible insight into the world as a whole, the metaphysical puzzles become the most urgent ones of all.”(1892/1920, pp. 461–2)

We may extract two general features of metaphysics from James. First, metaphysics is ‘world-facing’—as a general form of inquiry, metaphysics’ subject matter is the world (as a totality). Second, the *aim* of metaphysical inquiry is achieving maximal insight, in contradistinction to the special sciences of (e.g.) geology or mechanics, whose aims are constrained by more immediate and mundane concerns. So here we have a suggestion as to both the subject matter and aims of

metaphysics—metaphysics’ topic is *the world*, and its aim is *insight* independent of more narrowly circumscribed practical applications.

But while these two features provide a skeleton of what metaphysics is, we should hope to put some meat on these bones. I know of no better way than to poll the experts, and see what common themes we might extract. Schopenhauer defines metaphysics as follows:

By *metaphysics* I understand all so-called knowledge that goes beyond the possibility of experience, and so beyond nature or the given phenomenal appearance of things, in order to give information about that by which, in some sense or other, this experience or nature is conditioned, or in popular language, about that which is hidden behind nature, and renders nature possible. (Schopenhauer, 1819/1966, p. 164)

More recently, E.J. Lowe has offered an account of what metaphysics is or what it may do. Speaking for the tradition, he characterizes metaphysics as “the systematic study of the most fundamental structure of reality.” (Lowe, 1998, p. 2) This phrase and its close cousins can be found in innumerable accounts of metaphysics (and its talk of *structure of reality* matches closely the foundational idea of first philosophy, and *insight* about *the world*), but Lowe recognizes its vagueness, and offers more details of its particular domain of study. Contrasting metaphysical inquiry with the special sciences, he says,

Empirical science at most tells us what *is* the case, not what *must* or *may be* (but happens not to be) the case. Metaphysics deals in *possibilities*. And only if we can delimit the scope of the *possible* can we hope to determine empirically what is *actual*. (ibid., p. 5)

I hold that it is possible to achieve reasonable answers to questions concerning the fundamental structure of reality—questions more fundamental than any that can be competently addressed by empirical science. (ibid., p. 9)

Lowe notably uses the idiom of ‘being’—focusing on questions of existence. We’ll return to this later. Kit Fine (2012) suggests there are five characteristic features which distinguish metaphysics:

The aprioriticity of its methods; the generality of its subject-matter; the transparency or ‘non-opacity’ of its concepts; its eidicity or concern with the nature of things; and its role as a foundation for what there is... Metaphysics is concerned, first and foremost, with the nature of reality. (Fine, 2012, p. 2)

The last philosopher I’ll survey (but by no means for lack of sources!) is Ned Hall: “On a traditional conception, metaphysics aims to answer, in a suitably abstract and fully general manner, two questions: 1. What *is there*? 2. What is it *like*?” (Hall, 2010)

A GENERAL PICTURE OF METAPHYSICS

I want to suggest some themes we might draw out of these accounts, building a clear(er) picture of how we are to understand the term ‘metaphysics’. Schopenhauer draws a distinction between appearances and what lies behind them. This is the classic dualism of appearance and reality, and in these terms we can see (following Schopenhauer) metaphysics as that field of inquiry devoted to studying the reality “hidden behind nature”.

Indeed, Bas van Fraassen (2002), in announcing his rejection of metaphysics, summarizes it succinctly. Van Fraassen’s particular form of empiricist critique involves the following two features: “(a) a rejection of demands for explanation at certain crucial points and (b) a strong dissatisfaction with explanations (even if called for) that proceed by postulation.” (ibid., p. 37) From this we can reconstruct the picture of metaphysics he is rejecting. First, the metaphysician sees the need for offering explanations, and attempting to give them, even in cases where none appear forthcoming. This is just the expression of metaphysics’ inquiry into its special subject matter: seek to understand/explain whatever it is that undergirds appearance. Second, in such cases, rather than permitting the positing of a brute fact, the metaphysician seeks explanations that rely on postulated entities and processes. Notice, since such explanations (in Schopenhauer’s words) “go beyond the possibility of experience”, (1) they *require* postulation, (2) they can only be judged by *a priori* means, and (3) their adequacy is conditioned on satisfying the sought after explanandum.

Lowe construes metaphysics as concerned with the fundamental structure of reality. In particular he emphasizes metaphysics' special concern with *possibility*, so we might summarize him as claiming metaphysics concerns reality's fundamental modal structure. This is consonant with Schopenhauer (and the anti-van Fraassen) insofar as modal or counterfactual relations are precisely those which we typically think serve to render explanations for various phenomena (and even to permit basic predications! (cf. Goodman, 1983)), and such relations are in principle beyond the possibility of experience—their analysis must rely upon something other than observation.

Kit Fine provides a surfeit of features (which he thinks are individually necessary and jointly sufficient) by which to describe metaphysics. I'll focus on just three: its *a priori* character, its eidicity, and its foundational role. (1) As previously mentioned, the *a priori* character of metaphysical investigation appears to be a central feature due straightforwardly to the fact that the supposed subject matter of metaphysics by definition 'goes beyond' what *a posteriori* methods can possibly provide. (2) As for eidicity², Fine describes it as the concern for the nature of things, so to call metaphysics eidictic is to emphasize its concern not for mere accidents of a thing, but for its essence. Again this fits with the picture of investigating the reality behind mere appearances, since in appearances accidents and essences are mixed and the true nature of a thing is obscured. (3) Finally, metaphysics plays a foundational role for any further inquiry. For if we get a good grasp on the totality of the basic bits lying behind the appearances, we seem to have just the sort of epistemic base foundationalism desires.

Hall gives us the simplest account of the metaphysical project. I include it principally because it emphasizes the ontological character of contemporary metaphysics—the concern with 'being' and existence, especially of objects, and in particular with fundamental, individual 'simples' (though this is actually too narrow since the status of simples is itself an item of dispute). Metaphysics, in asking "What is there?" is asking "What is there *really*?"—to be distinguished from what there *appears* to be. And to ask "What is it like?" is to ask after the structure of these simples, to ask after their properties, their essences, and their (modal) interrelations—how they

²Adj: eidictic; In a modern philosophical context this word is, as far as I can tell, original to Fine. It is derived from the Greek 'eidōs' for 'form'/'idea'.

could be.

Summing up, what is our picture of metaphysics? It is the field of inquiry whose special domain is the world(/reality/nature), and whose aim is pure insight (knowledge) about the world (as opposed to more parochial and applied aims of the sciences). The particular task of this inquiry is not to catalog the surface appearances of the world, but—as per its aim—to describe the hidden reality which gives rise to the appearances. More specifically this task is accomplished through the application of *a priori* methods of investigating the unobservable natures (essences) of things. I'll repeat again what I mentioned in the last paragraph: the investigation of the nature of things as it occurs most typically in contemporary metaphysics has a distinctive *ontic* character. That is, it focuses its methods on questions of existence of entities, and seeks to provide an adequate metaphysics by developing an *ontology*—an enumeration of the (kinds of) things in the world. Metaphysics, on this picture, is a sui generis field of study, distinct in its focus (i.e. natures of things) and special in its status (i.e. fundamental or foundational).

SCIENCE AND METAPHYSICS

What has science to do with all this? The ways we might think that the sciences engage with metaphysics spans a broad spectrum. The goal of this section is to provide a grip on what the term 'naturalized metaphysics' is supposed to capture, and it will be helpful to clarify this by starting with a contrast class. The various relations that science is supposed to stand in with respect to metaphysics can be grouped under three general headings.

1: Anti-Naturalism The first is anti-naturalism. On this view, science has *nothing* to contribute to metaphysical inquiry, and indeed disagreements between a metaphysical theory and scientific results are in principle evidence that the scientific theory which is connected to those results³ are mistaken.

This view is not wildly popular in contemporary academic philosophy, but there are notable

³I'm keeping the description purposefully general. The ways that a scientific theory can be implicated in empirical results are complex and numerous. All that's important to say for the view under consideration here is that whatever value science provides, that value falls short of revising or undermining established metaphysical/*a priori* principles.

exceptions. In Chapter 1 I briefly discuss Arthur Prior's (1970) rejection of the arguments for the conventionality of the simultaneity relation in special relativity by choosing to reject special relativity wholesale. While this discussion is incidental to the main theme of that chapter, it's worth reflecting on here momentarily. Prior can be understood as standing his ground by way of a sort of transcendental argument: The very *possibility* of making sense of our experience and talk about the world of experience, and especially the actions and experiences of other persons in that world, *requires* there to be an objective time ordering of events. If, Prior seems to suggest, we give that up, then we give up any sense of a real, sensible, order of experience at all!⁴

Prior, it should be noted, accommodates the stunning empirical success of special relativity with a sort of instrumentalism. Modeling a general anti-naturalism on this example, we might say that for any given scientific theory which appears to assert a fact about the existence or properties of an entity, insofar as that assertion is in conflict with established metaphysical theory, it is to be understood as false and at best an instrumental feature of the theory. Actually, this statement is too weak for a pure version of anti-naturalism, since it permits science to establish metaphysical facts whenever metaphysics is silent. A more rigorous version is simpler to state: No scientific theory or fact can bear on our knowledge of the fundamental structure of reality (which is the province of *a priori* study alone). Thus stated, one would be hard-pressed to find someone espousing this position today,⁵ though some theistic positions might hold something like it.

2: Ontological Naturalism I would venture that the majority of analytic metaphysicians working today can be grouped under the second heading—ontological naturalism.⁶ This is the view that science does have some relevance to the outcomes of metaphysical investigation, but that

⁴While this does seem a bit overstated to my ears, we certainly can't say this position is outside the realm of rationality.

⁵Although a strict empiricist might affirm everything outside of the parentheses since *nothing* can bear on such (impossible) knowledge!

⁶This is also commonly referred to as 'metaphysical naturalism'. (cf. Kornblith, 1994; Forrest, 2000) I avoid this locution for two reasons: first, to emphasize its character as a thesis about what entities do (and couldn't possibly) exist, and second, to distinguish it clearly from 'naturalized metaphysics' which is a common name for the third view of science's relation to metaphysical inquiry.

scientific deliverances and the discoveries of metaphysics are somehow interdependent (to what degree and in what way this interdependency works is a matter of disagreement within ontological naturalism). Lowe, quoted above, gives an admirably clear demonstration of this attitude at work. He distinguishes the special domains of science and metaphysics by metaphysics' study of what is *possible* and science's study of what is *actual*.

While this specific characterization of metaphysics is particular to Lowe, it expresses a common theme in ontological naturalism: One of the principle conditions of adequacy for a metaphysical theory is that it be consistent with the basic requirements of science. What are these basic requirements? Generally, they're understood to be requirements for metaphysical accounts which eschew explanantia that appeal to mentalistic, abstract, or ethereal entities. Hence, the notion of 'natural' employed in the moniker 'ontological naturalism' is that which is to be contrasted with the *supernatural*.⁷

Another expression of this sort of attitude to what entities and kinds are to be allowed into one's ontology is summed up nicely in David Lewis' statement of Humean supervenience:

Humean supervenience ... is the doctrine that all there is to the world is a vast mosaic of local matters of particular fact, just one little thing and then another. (But it is no part of the thesis that these local matters are mental.) We have geometry: a system of external relations of spatiotemporal distance between points. Maybe points of spacetime itself, maybe point-sized bits of matter or aether or fields, maybe both. And at those points we have local qualities: perfectly natural intrinsic properties which need nothing bigger than a point at which to be instantiated. For short: we have an arrangement of qualities. And that is all. (Lewis, 1986b, pp. ix-x)

Notice here that the explanatory project is to decompose everything into mundane, simple, fundamental things. There is nowhere for something so complex as a god, or a mind, or an entelechy.

⁷Failure to distinguish between two uses of 'naturalism' is one of the principal sins of many discussions of naturalism. The distinction is to be made between naturalism as a *negative* thesis: "No positing of the supernatural." and naturalism as a *positive* thesis: "The empirical study of nature is the only correct epistemology." These two theses needn't be mutually entailing.

The success of such a project would be measured by its success to explain and describe *everything* in terms of these simple things.

Crucially, it is not merely the fact that ontological naturalism forbids reference to non-natural entities in rendering explanations, but also that such explanations take into account, and are amendable in light of the deliverances of science. This works in two ways: First, the metaphysical theories grown out of ontological naturalism are to be ‘inspired by’ scientific theories.⁸ Second, scientific theories are seen as the principal explananda for an adequate metaphysical theory—a mark of success for a metaphysical theory is how well it accounts for and explains salient features of our theories of science. Lewis emphasizes this connection to science (with a characteristic—for ontological naturalism—focus on *physics*): “Most likely, if Humean supervenience is true at all, it is true in more or less the way that present physics would suggest... If physics itself were to teach me that it is false, I wouldn’t grieve.” (Lewis, 1986b, pp. x–xi) Elsewhere Lewis states the thesis of Humean supervenience with an explicit reference to fundamental particle physics—the reason, it seems, is to make clear that the subjects of metaphysics are the entities and properties enumerated by science (and, importantly—see the final clause in the quote below—to pick up the torch when and where science proves inadequate):

[A]ll there is to the world is its point-by-point distribution of local qualitative character. We have a spatiotemporal arrangement of points. At each point various local intrinsic properties may be present, instantiated perhaps by the point itself or perhaps by point-sized bits of matter or of fields that are located there. There may be properties of *mass, charge, quark colour and flavour, field strength*, and the like; and maybe others besides, *if physics as we know it is inadequate to its descriptive task*.

(Lewis, 1986a, p. 14 emphasis added)

One last quote will hopefully solidify this view of the relation between science and metaphysics. Ted Sider, in rejecting the metaphysical view of presentism, takes science to be the

⁸Specifically, the eschewal of the supernatural is taken to be justified by the disposability of such entities in scientific investigations. Science “has no need of such hypotheses.”

central deciding factor in its rejection, and in doing so suggests a general guiding principle for metaphysical practice: “I turn finally to ... the fatal blow to presentism: that it is inconsistent with special relativity... [I]n cases of science versus metaphysics, historically the smart money has been on science.” (Sider, 2001, p. 42) There is a clear deference to science in the ontological naturalist approach, but as we’ll see in the third view, perhaps mere deference is not enough.

3: Naturalized Metaphysics The easiest way to introduce the third category of views is to contrast it with ontological naturalism. Whereas the latter holds that metaphysics should not *conflict* with science, the former—what I’ll call naturalized metaphysics—has a more stringent requirement: that science *dictates* metaphysics. The ontological naturalist leaves open the possibility that metaphysical inquiry is *sui generis* in its domain of study, and the requirement of consistency with science leaves open the idea that metaphysics itself provides conceptual grounding for the possibility of empirical and theoretical science. The ontological naturalist seeks (for example) to explain laws of nature at a more fundamental level than what the scientist needs or cares to (see e.g., Armstrong, 1983). Take as another example the analysis of properties: The metaphysician doesn’t particularly care which properties are actual, let alone their specific distribution. Instead, the metaphysician inquires into the *essence* of **Property** and asks what properties *are*. The particular connection to science that such a method has is that the properties which ontological naturalism takes seriously are those indicated by scientific predicates. Analysis of *any* property must ultimately be cashed out in these scientific predicates. For a particular example, consider Lewis’ analysis of properties as sets, where natural kind terms are those identified with the ‘perfectly natural’ sets. The role science plays here is to provide the criteria for which sets are ‘perfectly natural’. Such investigation and theorizing goes on independently of, and (conceptually) prior to the scientific project of, for example, identifying which objects *in fact* have which properties.

Naturalized metaphysics, on the other hand, holds that treating metaphysics as a wholly autonomous field of inquiry is fundamentally misguided. Metaphysics, for the naturalized meta-

physician, plays a role exactly as far as it serves the needs and uses of the sciences. This view denies that the generalist has a role to play in explicating things like properties or laws. The content of science exhausts the content of what metaphysics might say.

What this amounts to in practice varies according to the particular practitioner of naturalized metaphysics—in the next section I outline four grades of naturalized metaphysician, differentiated by what exactly science licenses us to say about metaphysics. For now it suffices to give a loose characterization, with emphasis on the break between the ontological naturalists and the naturalized metaphysicians.

Return again to Lewis' thesis of Humean supervenience. The basic claim is that the world can be analyzed into local properties instantiated at points, where these points are structured by geometric relations. This constitutes an ontological picture of what there is, and it is a *naturalistic* picture insofar as it purports to explain the phenomena in terms of the existence and natures of purely material/physical entities. That is, this view is consistent with naturalism in the sense of being *non-supernatural*. However, the naturalized metaphysician has a more stringent conception of what it is to be a (properly) naturalistic ontology. The criterion is this: Does physics⁹ make essential reference to the entities in question? Are the entities implied by our best physical theories? Do physicists recognize such things as real? If those entities—be they 'space-time points', 'tropes', 'substances', etc.—play no role in science, then cast them into the flames.

In the particular case of Humean supervenience there has been considerable attention paid to it by naturalized metaphysicians aiming to show that it simply cannot be sustained in the face of our best science. (cf. Ladyman and Ross, 2007; Maudlin, 2007; Humphreys, 2013) While the objections are varied, a common theme in these criticisms is that the thesis of Humean supervenience is scientifically inert.¹⁰ That is, it plays *no* role in scientific theorizing, and in fact (in this particular case at least) the sort of things which scientists make reference to appear to directly undermine such a picture of fundamental ontology.

⁹The transition from the nebulous term 'science' to the specific domain of physics is not a typographical oversight on my part. The naturalized metaphysician often takes for granted (or provide only the most cursory argument for) the idea that physics is the measure of metaphysics.

¹⁰In fact typically something more than this is claimed, namely, that Humean supervenience is demonstrably false.

While naturalized metaphysics appears generally as a united and coherent framework in terms of its negative theses—“reject the entities of a posited metaphysical explanation if they do not play a role in our best science”—the position becomes fractured when one looks to the *positive* claims of naturalized metaphysics. In order to get a better handle on the content of naturalized metaphysics, and to distinguish the particular positions within it which will concern us in the next few chapters, we’ll next turn to classifying the varieties of naturalistic positions with respect to metaphysics.

0.1.1 FOUR KINDS OF NATURALISM

In the previous section we surveyed three philosophical positions with respect to metaphysics: Anti-Naturalism, Ontological Naturalism, and Naturalized Metaphysics. The first two can be grouped together according to their shared acceptance of the epistemic autonomy of *a priori* reasoning. We will call metaphysical inquiry which takes *a priori* principles and intuitions seriously ‘analytic metaphysics’, which is distinguished from naturalized metaphysics (see figure 1).

Naturalized metaphysics then further decomposes into sub-groups. While divisions might be made in different ways, if we distinguish kinds of naturalized metaphysics by what they take science to *license* in terms of ontology, we can group views into four general categories: logical positivism, naïve naturalism, sophisticated naturalism, and pragmatism.

Logical positivism is an edge case for naturalized metaphysics: a position which has a view about metaphysics based on science, namely that there is *no* legitimate metaphysics based on science.¹¹ Logical positivism asserts that *science constrains metaphysical conclusions* to the point that there are no legitimate metaphysical claims. It can still be thought of as a metaphysical position in

¹¹It should be noted, however, that despite the official anti-metaphysical doctrines of logical positivism, in some of its manifestations there are traces of *a priori* metaphysics. For example the project to reduce physical objects to sense data involves (at least *prima facie*) *a priori* metaphysical commitments. My thanks to Cass Weller for making this point to me. Quine in *Posits and Reality* (Quine, 1955) points out that sense data are as much posits as are physical objects or microphysical/theoretical entities.

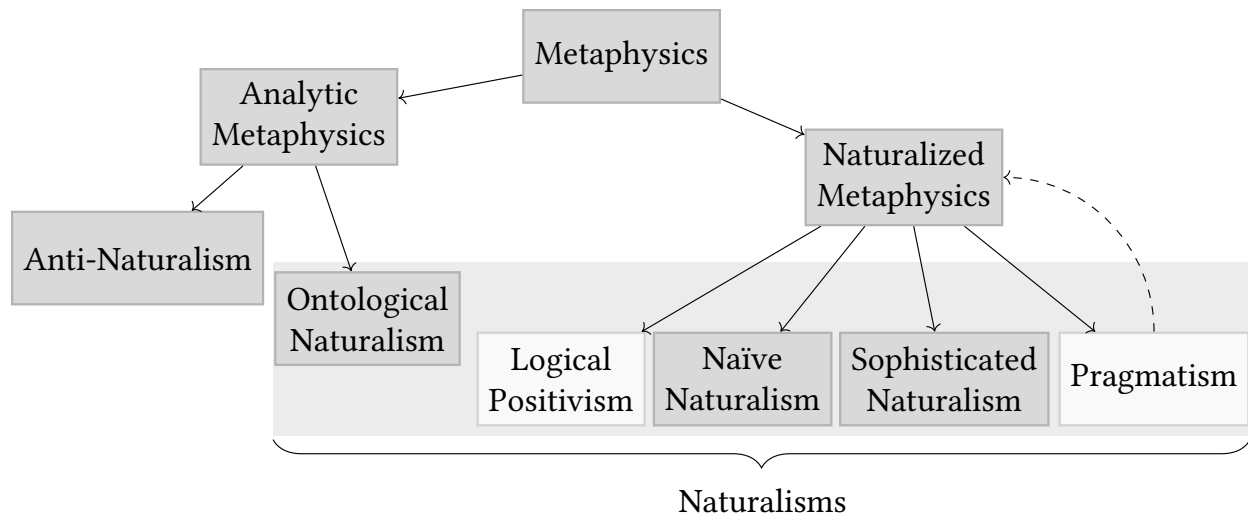


Figure 1: A taxonomy of metaphysical positions.

the (degenerate) sense that it includes a thesis about what inquiry can determine about unseen reality (namely, nothing). I include it here for the sake of completeness and to acknowledge that positivism is a clear case of naturalism (in the sense of a deference to and awareness of science); however, it won't play a further role in the following.

Naïve Naturalism is the position which claims that science can straightforwardly furnish our ontology without *a priori*/analytic distortions. The claim of such naturalists is that all we need to do is (1) look at our best scientific theories, (2) see which entities they quantify over, and (3) take those to be all and only the things which constitute the entities of fundamental ontology. The justification of such a claim is supposed to be derived from the empirical character of scientific experimentation and observation—the sober and methodical processes of science generate trustworthy (though fallible) ontological claims, while the singular intuitions of a philosopher have no similar constraints.

While such a position is an attractive one for anyone dubious of the epistemic legitimacy of *a priori* rationalist techniques, it is not itself free of problems. First, the naïve naturalist tells us to pick our best theories as the ones from which to derive ontology. But the question arises, which

theories are our best? Minimally, this judgement depends on some understanding of *how* one is to rank theories. The problem here is that such a ranking is not given empirically in experience, and will be a function of normative judgements. There is a standard response to this difficulty—‘best’ *is* empirical! We rank according to the degree of empirical confirmation a theory enjoys. For example, we can rank the theory of quantum electrodynamics very high due to the remarkable empirical confirmation it has enjoyed (theoretical predictions of the value of the fine structure constant agree with experimental measurements to within 10 parts-per-billion). However, this response ignores the fact that we have no empirical theory of confirmation—that is, no theory of confirmation that does not rest upon *a priori* assumptions. This is all to say that giving a coherent account of how to rank the best theories, even in terms as supposedly concrete as empirical success, inevitably rely upon non-empirical principles. This isn’t a problem in itself (surely, ever since Hume we’ve learned to make our peace with the vagaries of induction), but does become a sticking point for the naïve naturalist—the entire justification for looking to science rather than analytic philosophy for our ontological conclusions was in order to eschew the untrustworthy, contingent, and flighty intuitions of the human mind.

The second problem for naïve naturalism is that even granting that we can decide by a purely empirical process which theories are ontologically salient, there is no way to begin *drawing out* the ontological commitments of such a theory without some degree of interpretation. Interpretation requires starting points, starting points which themselves are not given in the theory or in experience. This is, really, simply the previous difficulty applied to interpreting a theory rather than, as we did above, applied to *choosing* a theory. The naïve naturalist’s position is that what there is consists of just what our best theories tell us there is.¹² But theories don’t speak for themselves and interpretation is a creative act which itself is subject to all the imaginative biases

¹²To cherry pick an example, Owen Flanagan, on a popular philosophy podcast, explains that metaphysics is not difficult for the naturalist: “If you’re a naturalist ... you’re not going to have much of a problem doing metaphysics, in at least one sense: You’re going to say that the way the world is is the way that our best sciences tell us it is.” (Flanagan, 2012, at 1:10:42) This sentiment is even more prevalent amongst science communicators and popularizers such as Lawrence Krauss, Neil DeGrasse Tyson, Stephan Hawking, etc. Whatever the rhetorical merits (which I think are slim) of this approach in fighting the ‘culture wars’, this sort of characterization is damaging to the naturalist position generally insofar as it encourages a sort of dogmatism about the epistemic certainty of science.

of the interpreter. So again the naïve naturalist is hoist by their own petard—replacing *a priori* metaphysics by reading off ontology from science only masks the *a priori* contributions.

I should note that the difficulties pointed to in the above paragraphs are neither new nor original—I don't claim to be breaking ground by calling them to attention. It has become a commonplace to point out that science relies upon assumptions, that observation is theory-laden, that theories are value-laden, and that science is (normally) done within a framework which imbues theories and empirical results with meaning. Indeed, I doubt any philosopher actually holds the naïve naturalist position, since its deficiencies are so readily apparent. However, is instructive to examine exactly how non-empirical considerations enter into the application of a scientific theory to a metaphysical problem.

The lesson to be learned is straightforward: If the only legitimate metaphysics is one that involves *no a priori* suppositions, then a metaphysics based on scientific theories is as illegitimate as the analytic metaphysics it is meant to replace. Yet, for the naturalist, there does seem to be *something* in the practice of looking to science for our metaphysics which recommends itself as more secure, more trustworthy, more serious than *a priori* speculation which is tested against mere intellectual comfort. The task of the naturalist is now to spell out *what* it is that distinguishes scientifically based metaphysics from analytic metaphysics.

The obvious place to start is by softening the antecedent of the conditional in the paragraph above: “the only legitimate metaphysics is one that involves no *a priori* suppositions”. Revising and weakening this claim so that it admits of scientifically based metaphysics, while still ruling out its analytic brethren, is the task of **Sophisticated Naturalism**. The goal of sophisticated naturalism is to maintain science as the exclusive source for legitimate metaphysical insight, but to explain this legitimacy in a way that does not rely upon overly simple conceptions of science as a pure empirical activity which comes pre-formed from of the simple act of observation. The way this is accomplished is by arguing for a more nuanced account of the foundational starting points on

which to build metaphysical theorizing. The sophisticated naturalist must provide a reason to accept that some privileged class of principles are properly naturalistic, while still ruling out the intuitive flights of fancy which are licensed by analytic metaphysics.

Sophisticated naturalisms reject the claim that *a priori* principles and knowledge are *wholly* untrustworthy. They aim to divide out those parts of knowledge which are admittedly justified *a priori* but are ‘naturalistically hygienic’. Thus, on this view, *a priori* knowledge contributes to metaphysics, but only insofar as those *a priori* principles play the correct role in science. This proviso acts (1) to modify the content of a supposed naturalistic metaphysics, and (2) to identify practical and epistemic constraints on what can be safely (metaphysically) concluded on the basis of our best scientific theories.

We can see this proviso at work in the position of ‘ontic structural realism’. Structural realism has taken the problems of underdetermination and the pessimistic induction¹³ to heart. The lesson structuralists take from these problems is that non-empirical principles play an ineliminable role in making various choices in questions of hypothesis, theory, observation, confirmation, etc. Generally, they accept that we cannot interpret a theory independent of non-empirical assumptions—we bring our pre-theoretical notions with us. The structuralist solution is to identify a way to understand a theory realistically without the usual step of interpretation. To accomplish this they do two things: first, they identify a theory with its mathematical models (i.e. they adopt a semantic view of theories); second, they take the invariants of the theory to be what the theory says are real. Thus, on this view, if theories tell us about what exists, they tell us that it is *structure*, rather than substantial individual things, which, on the classical corpuscular/atomistic theory are what the structure holds with respect to.

This view is plausible due largely to one feature: the widespread agreement that mathematics is an epistemically unimpeachable *a priori* discipline. That is, since math is *a priori*, it is secure in its objectivity and necessity, unlike the more promiscuous speculation of *a priori* disciplines which are really disguised reifications of our parochial conceptual ambit. We can avoid erring in

¹³Briefly, the recognition of the ontological discontinuity in past theory change, and the inference from those past discontinuities to the expectation of similar changes of ontology in the future.

mistaking concepts which *feel* necessary for eternal facts about reality by constraining our vision to just those facts of mathematics. Our fallible intuitions have no sway over the facts of logical and mathematical necessity—or so the argument goes. This mates with the recognition that science, particularly our most successful physical theories, are deeply and essentially mathematical, and empirical discoveries and predictions have largely been presaged by following ‘where the math leads’. Thus, the structuralist concludes, we’re justified in this ontic privileging of mathematical truth on the basis of the ineliminable role it plays in our best science.

Like logical positivism, **Pragmatism** is an edge case for naturalized metaphysics. Unlike the naïve and sophisticated naturalists who presuppose that the methods, aims, and standards of analytic metaphysics are correct, and only seek to replace the ‘content’ of metaphysics, pragmatism, when applied to metaphysics, questions the methods, aims, and standards themselves.¹⁴ That is, a pragmatist approach to metaphysics will ask us to re-conceptualize what the project of metaphysics is, and do so in an idiom which refuses to place abstract notions of ‘Truth’ and ‘Knowledge’ in positions of privilege. The pragmatist approach will help us see that metaphysical naturalisms, in their effort to replace the armchair speculation with scientific theorizing, *largely ignore the lessons of science*—that knowledge (little ‘k’) is hard-won, contingent, messy, and that the fundamental mode of knowing is practical, not propositional (and thus not primarily representational). What metaphysics could be—once we take the lessons of science into account, once we reject the picture of knowledge that places us as spectators, once we repudiate inquiry that takes as its aim the ‘mirroring’ of a static external world—is the project for a pragmatic metaphysics.

Thus, pragmatism as a form of naturalized metaphysics is different in kind from the naïve and sophisticated versions, not aiming to ‘scientize’ metaphysics, but rather to re-conceive the project, its aims, and its value. In figure 1 I’ve represented this feature of pragmatism as a looping back to the category of naturalized metaphysics, transforming the very conception of what such a thing

¹⁴And if the content changes that is incidental, especially since *what it means to be* content will be a function of those methods, aims, and standards.

should be. While I think the recognition that the proper understanding of naturalism is as a form of pragmatism, and while this recognition has a variety of wide ranging effects, some of which can be seen even now in the growing field of ‘philosophy of science in practice’, the positive argument of this dissertation will be limited to looking at how this turn to practice specifically affects questions of ontology—something which I think has been left largely unexplored. In effect, one can see the following project in terms of the recent ‘turn to practice’ within the philosophy of science. While much work in philosophy of science has now taken this turn, questions of metaphysics and ontology have remained within the older ‘logico-theoretical’ tradition exemplified by the realism/anti-realism debates which occupied the second half of the last century. What I aim to explore is how our understanding of ontology should evolve once it too makes this turn to practice, and poses these questions in the context of activity rather than theory.

Chapter 1

Putnam, Stein, & Space-Time

1.0 Introduction

The naturalist believes that science holds the trumps when it comes to inquiry—there is no higher authority in the land regarding questions about the world.¹ With respect to *metaphysical* inquiry, naturalism has been typically understood as the thesis that science should dictate our metaphysics: what there is is what scientists say there is. This claim, of course, requires some justification, especially in light of our (philosophy of science’s) ostensibly anti-metaphysical past.² The idea is that—consistent with naturalism—science enjoys some elevated epistemic status. However, without further elucidation this response is problematic. Simply put, the problem arises when we ask “in virtue of *what* does science enjoy this status?” At the global level the challenge is to articulate an account of this epistemic status which licenses metaphysical conclusions but that does not itself violate the very naturalist scruples which motivated naturalism from the start. I examine this challenge in the next two chapters. Before that, however, there’s a related, local, difficulty with drawing metaphysical conclusions from scientific theories: attempts at ‘reading off’ the metaphysics of a scientific theory inevitably sneak in background assumptions, and

¹The most radical versions of this credo extend the authority of science to cover political, cultural, and ultimately moral questions. This thoroughgoing view of scientific authority is often called ‘scientism’.

²That is, there is a long tradition in the philosophy of science which holds that we should take science seriously, and yet we should entertain no metaphysical claims whatsoever. This tradition shows up in the logical positivists, logical empiricists, instrumentalists, and constructive empiricists.

these assumptions are not themselves typically justifiable in purely ‘natural’ (scientific, empirical) terms. Perhaps this point seems too trite to occupy an entire chapter—such claims of the role non-epistemic values play in theory choice, or the theory-ladenness of observation are a dime a dozen in the literature of the past five decades. However, discussion of concrete examples, especially in the naturalized metaphysics literature, is largely absent. Being explicit about how and when such presuppositions enter will be valuable.

The example I’ll begin with is Einstein’s special theory of relativity. There is a well known argument offered by scientists and philosophers which claims to show that special relativity (SR) *entails* a specific metaphysical hypothesis, namely eternalism. The version of the argument I will focus on is one put forward by Hilary Putnam (1967) since it has the virtue of being explicit in both its argumentative structure, as well as its assertion of this scientific triumph over philosophy. There is a further virtue of the paper—it garnered a response by Howard Stein (1968). These two papers act as foils for the examination of how values/presuppositions intrude when interpreting the theory. This leads naturally to a discussion of what general reasons might be brought to bear in making decisions of interpretation and drawing metaphysical conclusions. I’ll consider several objections to the argument of this chapter, ending on the objection that I have not offered any motivation for why the naturalist *must* reject *a priori* justifications/presuppositions. That objection will be addressed in detail in the next chapter as we ascend to a more global view of naturalism and metaphysics. There, and in the following chapter, we ask precisely why naturalism is motivated to reject the *a priori*, and how naturalized metaphysicians attempt to develop a metaphysics in the absence of the *a priori*.

1.1 Time and Special Relativity

In order to understand how a scientific theory intersects with issues relevant to metaphysical questions we need to first look at the metaphysical theories in question. Presentism is the metaphysical claim that for all events x , x is present to me-now if and only if x is real. Roughly,

everything that is present is real, and everything that is real is present. The first conjunct of this claim isn't terribly controversial—all it asserts is that there are no 'unreal' entities which are present. This is near enough to a tautology to safely ignore. The other direction is where the view of presentism stakes its claim: reality is exhausted by those things which are present. This amounts to a denial that dinosaurs,³ President Lincoln, and the colonization of Mars are real. If (some of) these claims sound paradoxical, presentism has a response: Dinosaurs aren't real *now*, but they *were* real at some past time *t*. Presentism depends essentially on tensed language to make reference to things which are not real but were or will be. Thus, presentism puts a condition on the very notion of reality itself, constraining it to exactly those events which are present.

The dialectical opponent of presentism is eternalism, the thesis that everything—past, present, and future—is real. For the eternalist the ontological status of past, present, and future events is the same. The difference is a relational one, such that events are ordered by 'before', 'same time', and 'after' relations. Thus, the events constituting the existence of dinosaurs, President Lincoln, and the colonization of Mars are related in a way that can be expressed by event statements such as "President Lincoln's Gettysburg Address happens after the K-T extinction event."⁴ and "The colonization of Mars happens (if at all) after the Lincoln-Douglas Debates." Eternalism claims that all tensed language can be translated into this relational language without loss. The inference to be made on the strength of this assertion is that since tensed language is eliminable, there is no need to posit a special ontological status for the present moment—that is, the moment relative to which tensed language gets its truth-value. Parsimony considerations then militate against such a positing, with the consequence that all events in the 'block' universe enjoy equal ontological status.

³All 'entity' and 'thing' language should be explicitly construed as space-time events. To get a better feel for this stipulation, an 'event' can be seen as analogous to a spatial point in Euclidian space. Roughly, they are non-extended locations of a 'happening', where this includes things like an atom existing at a particular time, or a point in space bearing a particular property. *Objects* are then composed of events, such that a persistent object is a continuous series of events that trace a path (world-line) in space-time. When I refer to ostensible objects like Abraham Lincoln or a particular dinosaur, what you should have in mind is an arbitrary event in the relevant world-line. In what follows I will refer to entities named by (e.g.) *a*, *b*, etc. These are always to be understood as referring to events.

⁴'Happening' locutions in these two statements should be understood in their tenseless sense of identifying an event, independent of ascribing it an observer relative time. This is the source of the odd phraseology of the statements.

The debate between presentism and eternalism is an old one, tracing at least as far back as Aristotle's discussing of the sea-fight tomorrow.⁵ Briefly, the question considered is whether there is a determinate truth-value that can be assigned to propositions about future events. On one hand it seems that there must—any proposition has a truth-value (if we accept the law of Excluded-Middle) and if we are unaware of any particular truth-value that is a function of our epistemic position, not an ontological feature of the world. On the other hand, if a proposition about a future event has a determinate truth-value, then it would seem that it has that truth-value *necessarily*.⁶ And if this is the case, then the future is fixed, and necessarily so, which seems to violate our intuition that the future is—in at least some respects—genuinely open. The debate about eternalism and presentism ramifies, affecting views about determinism and freedom, responsibility and action, etc. The fact that this debate has important connections to other fields of philosophy should cause us to pay attention when it is claimed to be *solved*. Hilary Putnam (among others) does precisely this—he claims that special relativity has decisively settled (in the affirmative) the question of whether the future is determinate.

1.1.1 PUTNAM, ETERNALISM AND THE RELATION R

Our first step evaluating Putnam's argument is to translate the notions associated with presentism and eternalism into a vocabulary suitable for analysis in terms of physical theory. Putnam's argumentative strategy is to state as clearly as possible the fundamental commitments of presentism in such a language, and show that those commitments, coupled with SR, entail something quite different from what the presentist has in mind.

The fundamental notion of presentism is that of a well-defined 'present'. It requires there to be something called *the privileged present* which is the collection of all and only the things which are real. Another way to talk about the reality of a thing is in terms of its determinateness.

⁵Aristotle's discussion never explicitly touches on the issue of eternalism and presentism, but it is clear the diverging answers to the difficulty Aristotle identifies depend on one's attitude toward the determinateness of the future.

⁶That is, once a truth value is assigned, it doesn't seem like there's any sense to the idea that the value could *change*.

This is Putnam's terminology, and it is preferable because it can be given a linguistic gloss: An event is determinate just in case any (contingent) proposition about the event has a fixed truth-value.⁷ Thus, the present is the collection of all and only those things which are determinate. One natural requirement to put on this collection is that it defines an equivalence relation R which we'll call the 'determinateness' relation such that Rab holds just in case a is determinate with respect to b . We'll stipulate that 'determinateness' and 'reality' are interchangeable so to say that x is determinate with respect to y is to say that x is real according to y (where 'according to' is a metaphysical condition, not an epistemic one).⁸

Given these conditions, there is a persuasive case to be made that presentism is the layperson's default view about past, future, and present events. Putnam (1967) presumes as much, claiming that the layperson would assent to the following claims (note that these are conditions on the proposed relation R):

1. I-now am determinate.
2. At least one other thing is determinate, and it is possible for this other thing to be in motion relative to me.
3. If a is determinate to b , and b is determinate to c , then a is determinate to c .⁹

(Adapted from *ibid.*, pp. 240–241)

Putnam's argument is that these conditions on which events are determinate—assumptions which are supposed to conform to the metaphysical view of presentism (and thereby lend support to the plausibility of the view)—if conjoined with special relativity, must either be false, or the relation R must be the universal relation (i.e. the relation which holds between *any* two space-time events).

⁷Of course, circularity threatens if one asks after the conditions for when a truth-value becomes fixed, since one attractive answer is to reply in terms of reality. At this point I'll just have to punt, and trust the reader has an adequate intuition of 'reality/determinateness'.

⁸I should stress again that the variables referred to here (a , b , etc.) denote space-time *events*—not temporally and spatially extended objects.

⁹Notice that strictly speaking the conditions here described ((1) reflexivity and (3) transitivity) only require R be a preorder, which is weaker than an equivalence relation. As Putnam expresses these conditions, (2) is about the existence of another 'observer'. While it isn't explicit, the best way to read this condition is as asserting a symmetry condition. The idea is that this other observer takes you-now to be determinate just in case you take the observer to be determinate. The role of the symmetry condition will become relevant later as it does not hold under Stein's interpretation.

But interpreting R as the universal relation is flatly inconsistent with the original statement of presentism, and rejecting (1), (2), or (3) is less plausible than rejecting presentism itself. Thus, Putnam urges, we should reject presentism—(all) future events are determinate.¹⁰ Concluding, Putnam says: “the problem of the reality and the determinateness of future events is now solved. Moreover, it is solved by physics and not by philosophy.” (Putnam, 1967, p. 247)

In order to follow Putnam’s argument the only preliminary required is an understanding of the behavior of the simultaneity space of an observer in an inertial frame relative to another observer. First, the simultaneity space¹¹ of an observer at proper time τ_0 is the class of all events reachable by a light-beam sent out from the observer at $\tau_{n<0}$ and received back at the observer at $\tau_{m>0}$ such that $|n| = |m|$.

A mundane, yet analogous, example will help: Imagine sending an RSVP to a friend, A , and you’d like to determine when, exactly, they get the invitation. We first make three simplifying stipulations: (1) It takes no time for your friend to compose a response, (2) they will do so immediately upon receiving the invitation, and (3) the invitation spends no time sitting in the mailbox (i.e. the mail carrier immediately picks it up at both ends). Now, you send the invitation at noon on Monday, and receive A ’s response at noon on Wednesday. From this information we can easily calculate the *total* elapsed time: 48 hours. From symmetry assumptions (specifically, that there is no difference in the average speed of the mail carriers nor is there a difference in distance traveled on the outbound and return routes) we can then calculate that the letter arrived at A at noon Tuesday (that is, exactly $\frac{1}{2}$ the round-trip time). You can now say with confidence that A received notification of your party exactly when you were enjoying your lunch on Tuesday. Now, suppose that you sent out invitations (at noon) in the same manner on the previous Saturday and Sunday to B and C respectively. And further suppose that you received a response from C on Thursday at noon, and from B on Friday at noon. By the same methods described above, you

¹⁰The argument also entails the determinateness of all past events as well. For simplicity in exposition these have been ignored.

¹¹The conditions described here capture ‘Einsteinian simultaneity’. There is a substantial literature on the choice of a simultaneity relation since the relation is, in a well-defined sense, conventional. Nonetheless, the relation described here is widely accepted, and considerations of alternate relations are beyond the scope of this chapter.

would calculate that they too both received their letters at noon Tuesday (a remarkably coordinated mail service!). Thus, the events of A , B , and C receiving invitations, and your eating of lunch on Tuesday all occupy the same simultaneity space—we say they are *simultaneous* with one another (they stand in the ‘simultaneity’ relation S).

For an observer at rest this notion of a simultaneity space is a good surrogate for the idea of the privileged present discussed above: it defines a space with no extension in time, and which captures the totality of events in that space, including the observer. The presentist hypothesis, suitably translated to the language of special relativity, can thus be understood as identifying the determinateness relation with the relation that holds between all and only the events in a simultaneity space, the simultaneity relation S . Relative to an observer, S describes an equivalence class of events which are intuitively the events that are ‘now’. It is a straightforward and seemingly innocent extension to go further and say that such events are also exactly those which are *real*. Also notice that—importantly for the presentist—this class of events is a proper subset of the class of all events (hence it is *not* the class of all events, and so this equivalence relation isn’t the universal relation).

However, because (a) the simultaneity space is defined in terms of the paths of light-rays, and (b) SR stipulates the speed of light is fixed, it follows that the simultaneity spaces of two different observers are not generally coextensive, and in fact are coextensive *only if* the relative velocity of the two observers is zero.¹² These facts about the special theory of relativity will undermine the presentist identification above.

To show this Putnam imagines the following situation: Let me-now be an unaccelerated observer o . Consider a second observer p , moving relative to o with a non-zero velocity (so $S_o \neq S_p$), and who is collocated with o at time t (so $o = p$). It follows from the above fact about simultaneity spaces that not every event in p ’s simultaneity space is in o ’s.¹³ But, by hypothesis o and p

¹²Briefly, this follows (in two dimensional Minkowski spacetime) as a geometrical consequence of the fact that for right triangles the median and angle bisector of the right angle are not in general coextensive (and are only coextensive in the special case where the triangle is isosceles). Translated into a physical interpretation, the median bisector corresponds to an observer’s simultaneity space, yet it is the right angle bisector which is invariant between observers. See Appendix A.

¹³Alternatively, there exists an event $e \neq o$ such that it holds S to o but not p or vice versa.

are collocated and so are ‘in each other’s’ simultaneity space (the ‘observer-events’ which correspond to o and to p ‘happen’ together—they are one and the same). We can now ask between which events R holds on its identification with S and according to the conditions imposed upon it above.

First, o and p are determinate (by (1)), and second, they are determinate to each other (by (2)). From this it follows by the transitivity of the relation R that *every* event in p ’s simultaneity space will stand R to o . Notice what has happened: we identified R with S since it seemed to be a plausible physical candidate for the required metaphysical notion of the ‘present’. Then we found, on the basis of this identification, that $R_o = S_o \cup S_p$, which entails that $R_o \neq S_o$ (assuming, as we will, that no $S_n = \emptyset$). Contradiction. We are presented with a choice: Deny R is an equivalence relation (specifically, by denying transitivity), or find a different physical relation with which we can identify R .

Either choice is a problem for the presentist. On one hand, transitivity seems to be a minimal criterion for a coherent notion of determinateness. The denial that R is transitive amounts to the possibility that two observers could agree that each other were real and determinate, yet part of each’s reality was somehow indeterminate according to the other.¹⁴ This seems like too great a cost to bear. The other choice—maintain R is an equivalence relation—cannot be accepted by the presentist without thereby denying presentism. It can be shown that (on the basis of the above argument) if R is an equivalence relation, the only relation it could be is the universal relation U . (See Appendix B for details on how to extend this argument.) But identifying the relation of determinateness with the relation that holds between every event and every other event is simply to assert the denial of presentism: eternalism.

As we saw above, Putnam concludes that SR has demonstrated that presentism is false, and a substantive metaphysical position (eternalism) has been shown to be the consequence of *physical*

¹⁴Here is where the linguistic gloss comes in particularly handy, since the foregoing amounts to asserting that there is someone else who is real/determinate (indeed, in the example we’re collocated!) and yet there are propositions which have a fixed truth value for one of us, yet have no fixed truth value for the other. Again, it’s important to emphasize here that this is not merely an epistemological wrinkle, determinateness is supposed to be understood in metaphysical terms.

investigation. Here, we are asked to observe, is a clear case where comfortable intuitions about reality and what exists are not only false, but are patently inconsistent!

1.1.2 ARTHUR PRIOR AND *A PRIORI* SCRUPLES

To summarize §1.1.1, what Putnam has shown is that given some natural conditions on the relation R and the assumption that Special Relativity holds, the identification of R with S entailed that either $R = U$ (and so $R \neq S$), or transitivity must fail for R . Putnam concludes that since the latter horn of the dilemma is too unpalatable for anyone to accept, what we're left with is the former: The relation R must be the universal relation.

There is, of course, a third option open to the presentist who wishes to resist this argument. They can simply deny SR. And indeed, some have done so. Arthur Prior is notable in this respect, rejecting SR in order to preserve the obvious and intuitive notion of absolute simultaneity. Prior gives a defense of the notion of the present, aiming to recapture the picture of a distinguished moment which constitutes all and only what is real.¹⁵ He acknowledges a way of thinking which makes the past and future seem like *places* where things happen:

It is tempting to think of the present as a region of the universe in which certain things happen, such as the war in Vietnam, and the past and the future as other regions in which other things happen, such as the battle of Hastings and men going to Mars... [But this picture] doesn't bring out what is so *special* about the present; and to be more specific, it doesn't bring out the way in which the present is *real* and past and future are not. (Prior, 1970, p. 246, emphasis original)

He admits the difficulties with this view, presented primarily by the sorts of scientific considerations rehearsed above. He summarizes the difficulty in the following manner: suppose we're

¹⁵Prior is not directly engaging the argument which we considered above. He considers the more basic question of whether there is an objective time-ordering of space-like separated events compatible with special relativity (that is, a serviceable notion of absolute simultaneity). The above argument for the reality of future events takes the answer to this question ("no") as a premise. Thus, Prior, in rejecting SR on the basis of its lack of an objective time-ordering, would deny the soundness of the above argument.

observing a distant object which is pulsating with light. After seeing one of these flashes of light, we can ask whether the next flash is *happening now*, has *already happened*, or is *yet to happen*. Prior insists that “this is *always* a sensible question.” (Prior, 1970, p. 247, emphasis original) And he goes on to admit this “is what the special theory of relativity appears to deny.” (ibid., p. 248) He sees that for his question to make sense, there must be a notion of absolute simultaneity, a standard against which to check the relation between the two events of seeing pulse n and the production of pulse $n + 1$.

But of course, it is precisely this notion of absolute simultaneity which SR rules out. If we accept that any observer has just as good a claim as any other for making accurate measurements of events in their vicinity, then no such absolute frame can be found. Putnam’s response, embracing eternalism and the denial of a privileged present, is one obvious reaction to this result. Prior, though, insists that we are not *forced* into Putnam’s conclusion.

One possible reaction to this situation, which to my mind is perfectly respectable though it isn’t very fashionable, is to insist that all that physics has shown to be true or likely is that in some cases we can never *know*, we can never *physically find out*, whether something is actually happening or merely has happened or will happen. I’m sure there *are* questions which are perfectly genuine and intelligible questions but which seem to be incapable of being answered. (ibid., p. 248, emphasis original)

Let us be clear about what Prior has committed himself to in this response. He implicitly makes a distinction between two sorts of constraint which special relativity could impose on metaphysical understanding. The first is a metaphysical constraint—the supposed incompatibility of a privileged present with SR reveals how the world *must* be (or minimally, how the world cannot be). The second (and the one Prior prefers) is the idea that SR at best imposes an epistemic constraint—the incompatibility reveals the failure of experimental and empirical methods to discover the true nature of the present. In the framework of the theory, this amounts to insisting there is—in reality—a preferred foliation of space-time which is observationally beyond our reach. No matter, Prior can insist, since *a priori* intuitions about time, the present, reality, etc.

are enough to ground our knowledge that such an absolute reference frame exists, indeed *must* exist. Elsewhere, discussing the same incompatibility between special relativity and our concepts of past, present, and future Prior puts the point more starkly:

Coming back to this allegedly [according to SR] meaningless question as to whether you or I saw the light-flash first, surely what it means is just this: When I was seeing the flash, *had* you already seen, or had you not? In other words, when my seeing it was a *present* fact, had your seeing it become a *past* fact, or had it not? And I just cannot be persuaded that such a question is meaningless—its meaning seems to me perfectly obvious...

So it seems to me that there is a strong case for just digging our heels in here and saying that, **relativity or no relativity, if I say I saw a certain flash before you, and you say you saw it first, one of us is just wrong ...** To put the point another way, we may say that the theory of relativity isn't about *real* space and time, ... the 'time' which enters into the so-called space-time of relativity theory isn't this, but is just part of an artificial framework which the scientists have constructed to link together observed facts in the simplest way possible. (Prior, 1996, pp. 50–51, emphasis original, boldface added)

Prior's position is as clear as it is remarkable from the point of view of the scientist or naturalistically inclined philosopher. The incompatibility between the special theory of relativity and our commonsense notion of what are meaningful questions about my personal time-order and anyone else's time-order certainly demands resolution, but to many of the broadly empiricist persuasion, Prior's adherence to his *a priori* commitments seems mistaken and foolhardy. For someone who is not willing to reject a theory as empirically impressive as special relativity, what are the remaining options? Does SR, in fact, reveal to us the true and fundamental nature of time? We'll see that such revelation is itself contingent on certain choices—choices genuinely left open by experience.

1.1.3 HOWARD STEIN AND THE EQUIVOCALITY OF THEORY

I think it's fair to say that Prior's approach to the applicability of SR to the metaphysical question is distinctly *anti-naturalist* in the sense described in the previous chapter. For this reason we'll move on, looking to the more explicitly science-friendly approaches. Short of rejecting SR, then, has Putnam effectively ended the debate? That is, has Putnam shown that the *only* relation which satisfies the natural conditions on R is the universal relation; that the choice is *forced* on us by the requirements of rationality? No. Howard Stein (1968) shows that Putnam's argument has not, in fact, established that SR logically *requires* determinateness to be identified with the universal relation.

To do this Stein offers an alternative relation, R' , which is a plausible candidate for describing determinateness, yet which is neither the trivial nor universal relation. For Putnam to have established that SR decisively settles the metaphysical question in favor of eternalism, he would have to show that *any* reasonable conditions on a relation meant to capture the notion of determinateness result in the universal relation. If Stein can show there to be an R' meeting some *reasonable conditions*, then Putnam has, effectively failed to show that SR has indeed decided the matter. Below we'll see in detail why and how this is the case, but before that we should get a sense of Stein's relation R' .

R' , to be a plausible candidate, must meet two desiderata: First, it must be consistent with SR; and second, whatever constraints are imposed upon it must plausibly fit with our intuitive sense of the determinateness relation. We'll grant the first—but what constraints should we introduce? Stein suggests the following two: (a) R' is *not* the universal or trivial relation; and (b) R' is a transitive, reflexive relation. (We'll consider below if and why these are indeed plausible.)

Supposing for now that these conditions do no excessive violence to our intuitive sense of the relation, we can now ask if a relation exists that meets these conditions. The answer is indeed yes: it is the relation which holds between any event e and all the events in its 'causal past'—those events from which a light-ray could reach e . Call this relation C . Transitivity and reflexivity hold for C —for any p in e 's causal past, if q is in p 's causal past, q is in e 's causal past, and we'll define

the causal past so that each event is in its own causal past. It is also true that this relation is not (typically)¹⁶ trivial and is indeed not universal. The key feature to this relation satisfying the latter condition is that the light-cone structure of any space-time point is invariant. This means that two collocated observers—no matter their relative velocities—will have coextensive causal pasts. The consequence is that there is no event which holds C to the collocated observers but about which they disagree on its temporal ordering.

From these properties it can be shown that C is in fact *unique*. Any transitive, reflexive, non-trivial/non-universal relation which answers to the natural conditions we placed on R' must be coextensive with C . There is, however, one qualification in all this: symmetry does not hold for C . In fact, C is antisymmetric—if p is in o 's causal past and o is in p 's it follows that o is identical to p . Recall above, in discussing Putnam's approach it was presumed that characterizing the determinateness relation as an equivalence class was a reasonable desideratum. Indeed, one way to characterize Putnam's argument is as trading off the condition that determinateness is always a proper subset of the set of all events in favor of preserving the condition that it is an equivalence class. Stein, on this account, makes a different choice, preserving the former condition and weakening the latter. Notice, however, that there are a variety of ways a relation can fail to be an equivalence relation (six ways in fact) and depending on the desired intension of the relation, some rejections are more catastrophic than others. Previously, in the course of rehearsing Putnam's argument we saw that one option for the presentist was to reject transitivity. This option, it was decided, did rather more violence to our ordinary notions than the alternative, and was rejected. (See footnote 22 below as well.) Rejecting symmetry is less of a cost, though still a significant one for the traditional presentist.

To get a sense of this cost we can ask how R' differs from the intension of the relation R . R 's intensional representation is: “- - - is determinate to ...” Recall, this was intended to capture the person-on-the-street's notion of determinateness, or what is real *now*. Putnam then argued that if

¹⁶We can always arbitrarily define the relation on a pathological space-time so that the relation is trivial and/or universal. For example, any non-empty relation defined on a zero-dimensional space-time manifold consisting of a single event (point) would be both trivial and universal.

we are to preserve some particular virtues of this notion in the face of SR, then it follows that the relation is universal. Taking such a relation to be universally satisfied, the result is the following claim: everything is determinate to everything; everything is real with respect to everything. The sea-fight tomorrow is just as real as my typing this today is just as real as Lincoln's Gettysburg address yesterday (or rather, some 7-odd score years ago).

Is the same intensional reading appropriate for R' ? No. The most natural reading of R' is essentially tensed: “- - - *has become* determinate to ...” If we take the simultaneity space to be a partition of an observer's past and future, one consequence of identifying R' with C is that all events to which e holds R' are in e 's past. R' says that there is no sense to something *becoming* determinate, no fact of the matter (and no sense in asking) about what is *now* determinate with respect to me. The only sense of determinateness we have is when some event ‘shows up’ in our causal past.¹⁷

The cost, then, of adopting ‘Stein-Presentism’¹⁸ is that it also rather decisively undermines traditional presentism. This is due to the fact that traditional presentism *requires* there to be a symmetry in the determinateness relation—to be is to be ‘on equal footing’ with all that (presently) is. Stein's account has—it seems—failed to make good on the presentist notion of the ‘now’. It mischaracterizes the presentist ‘now’ in another way as well: The relation C is such that for any event o , the set of events related by C to o will also be related C -wise to *any* time-like or light-like extension of o into the future! If we accept $R' = C$ then ‘Stein-Presentism’ cannot consistently hold that things pass from indeterminate to determinate, and back to indeterminate

¹⁷To get a clearer idea of Stein's relation, imagine tomorrow we look up at the sky and see a supernova, bright enough to be visible to the naked eye. We might naïvely ask, *when* in the past did that event *happen*? Putnam says that this question makes no sense. What we should say is that it was always determinate irrespective of my epistemic situation. Stein would agree that the question is senseless, but not because the determinateness of the event is a tenseless fact about the universe. Rather, it is without sense because there is *no fact of the matter* about ‘when’ an event becomes determinate, but there is a fact of the matter about when the event has entered my causal past (i.e. just when its light reaches me). We can resist the conclusion that an event s is ‘always’ determinate, by rejecting the notion of *becoming* determinate, and holding that all we can legitimately say (about, for example, the supernova) is that once we have seen it that it *has become* determinate.

¹⁸It has come to my attention that such a view—called ‘light-cone’ presentism—has been defended. The view (attributed by Rickles to Mark Hinchliff (2000)) holds that the null surface of our past-lightcone does indeed define the present moment. Such a view is not *prima facie* attractive however, since it is, as a consequence, committed to the view that the big-bang (or moments shortly thereafter) are happening now (since the cosmic microwave background is light from those times reaching us). This seems to me like a *reductio* of such a view. (Rickles, 2016, p. 219)

again (like traditional presentism does). Instead, once something is determinate, it stays forever determinate. While this is another departure from traditional presentism, it can be seen as a virtue, since it seems to account nicely for the phenomenological asymmetry between future and past events (a persistent puzzle for presentists). Indeed, such a modified theory fits nicely with ‘modified’-presentist accounts like the ‘growing block’, giving them an air of scientific legitimacy.

To take stock: Putnam’s relation is inconsistent with traditional presentism. Stein’s relation seemingly preserves something of presentism, namely, essentially tensed language is required to describe events and the future is open in a sense denied by eternalism. But this relation does so only by abandoning completely the central notion of presentism—the privileged present. I should mention here that my calling it ‘Stein-Presentism’ unfairly saddles Stein with a metaphysical theory of time. He makes no such claim.¹⁹ But let us suppose that a hypothetical naïve naturalist metaphysician wished to be as well informed as possible about what science tells us about the reality of future events. While Putnam has shown one way to derive a relevant metaphysical conclusion from SR, Stein has shown that such a derivation is not unique—there is in principle a second relation which answers well enough to our intuitive notion of reality (becoming, determinateness, etc.) that the temporal metaphysics of SR is not a foregone conclusion. We must ask then, what ought our naturalist metaphysician do?

1.1.4 WHAT CAN SETTLE THE CHOICE BETWEEN R AND R' ?

In order for our hypothetical naturalized metaphysician to ‘let science be her guide’ in matters metaphysical, it appears she must make a choice between the relations R and R' . The possibility space includes at least three options. (1) Accept R as the correct metaphysical relation that holds between all and only determinate space-time events, (2) accept R' , or (3) reject them both.²⁰ A fourth option—accepting both—is ruled out simply for the plain fact that they are extensionally

¹⁹As I was first formulating this chapter and imagining the dialectic between Putnam and Stein as a metaphysical dispute I didn’t imagine anyone (let alone Stein) would seriously defend such a ‘Stein-Presentism’. However, see footnote 18.

²⁰This last option branches into a variety of sub-options, for there can be a multitude of reasons and ways to *reject* any particular claim.

inequivalent. Clearly, the logical legitimacy of ‘accepting’ both claims rests on an articulation of the concept of ‘accepting’. Leaving aside the third option for now, our metaphysician must choose R or R' . In such cases, it is typically understood to be rational to choose the option which has the most evidence in its favor. Our metaphysician, being of the ‘natural’ variety, will impose conditions on what counts as evidence—and this will ostensibly be something to the effect of forbidding non-empirical reasons.

We must then ask what empirical evidence supports one or the other of these relations as being the real, or true, relation. It will help to have a clearer account of what exactly differs between the accounts Putnam and Stein offer. We can do this by first revisiting the person-on-the-street’s notion of the present (which the sought after relation is supposed to formalize). It is supposed to consist of three conditions: (1) The present consists of a proper subset of all possible events and includes something other than me. (2) I am always present, and the present is ‘shared’—if something is present to me, everything that’s present to it is present to me. (3) If something is present to me, I am present to it.

Translated into our discussion of formal relations, the conditions are as follows:

- (1) Determinateness is neither universal nor trivial.²¹
- (2) Determinateness is transitive and reflexive.
- (3) Determinateness is symmetric.

These three conditions form an inconsistent triad and the debate between Putnam and Stein is over which two we should preserve. Putnam chooses to accept (2) and (3) while rejecting (1). Stein, alternatively, rejects the condition of symmetry (3), preserving (1) and (2).²²

²¹Formally: $\exists w, x, y, z (\langle w, x \rangle \notin R \ \& \ \langle y, z \rangle \in R \ \& \ y \neq z)$. Informally: “There are at least two events which are not determinate with respect to one another, and there are at least two distinct events which are determinate with respect to one another.”

²²There is, of course, the third logical possibility of rejecting (2). Since this condition is a conjunction its rejection branches into two alternative options. First, one may reject reflexivity alone; or second, reject transitivity. (I am ignoring the third option of rejecting both since it’s clear how to combine the two.) While there is no logical prohibition on such positions, they appear to be less plausible than the options adopted by Putnam and Stein. The first option entails that there is at least one event which is not real (determinate, present) to *any* event (including itself). This is at the very least both physically and metaphysically unmotivated, and seems to me like the very definition of profligate metaphysics. The second option was already discussed above, and requires a fracturing of reality, which again seems like too great a cost to be worth it.

The difference in their approaches to which condition to reject and which to preserve explains their difference in what metaphysical lesson to draw from SR. What this means is that we should look for the evidentiary status of these conditions. That is, *what is* the evidence for each option, which choice of these two has the best/most evidence in its favor, and what is the epistemic status of said evidence? There are three broad categories under which the evidence could fall. It could be empirical, *a priori*, or theoretical. By empirical I mean something known or justified by direct observation. By *a priori* I mean something known or justified independent of observation or experience. And the third category of theoretical is meant to capture a hybrid sort of evidence, common in scientific reasoning of inferring to things strictly beyond direct experience, yet justified not by pure reason but rather by inductive and abductive inferences.

We can ask, are (1), (2), or (3) justified or known by direct experience? Since all three are universally quantified propositions about relations, it's hard to see what experience could give us direct knowledge of them. Mere justification by experience is only slightly better off—we might suppose direct contact with other things (events) offers some justification for various propositions about determinateness, but more often our judgements about more (temporally and spatially) distant events presuppose these very features, so can't be legitimately justificatory.

How about justification or knowledge by way of theory? This perhaps seems more promising than direct experience, since we certainly have considerable inductive and abductive support for some or all of these conditions. However, it should be clear that theoretical reasoning cannot *decide* the matter, since the theory to be deployed (SR) is precisely what is at issue. That is, deciding the correct interpretation of the determinateness relation cannot depend on SR alone since Putnam and Stein have shown that it accommodates either relation. Neither Putnam nor Stein get the theory *wrong*, yet their arguments yield different interpretations. What differs is not the scientific theory at stake, but rather some non-empirical ingredients that were required to yield a metaphysical interpretation at all.

This brings us to our final option—justification or knowledge of these conditions via non-empirical or *a priori* means. If we are to think that we have any knowledge at all about a relation

between events such as ‘determinateness’ then it seems that knowledge comes by way of *a priori* reasoning and justification.²³ SR does not ‘speak’ until we choose which conditions we wish to preserve, and then its deliverances are a function of those choices. If we imagine two staunch metaphysicians, presentist and eternalist, locked in debate, will SR help their debate?²⁴ They can both have all the facts about SR, but as long as the presentist requires a relation that is not universal, and as long as the eternalist requires a relation that is symmetric, then their disagreement can neither be about the *science* nor can the theory adjudicate a resolution. Of course, it doesn’t follow that there *is no* correct metaphysical fact of the matter, but if there is, the arguments leading to it must contain at least one *a priori* premise.

This is a long way of expressing what Lawrence Sklar summed up in his ‘MIMO’ principle: Metaphysics in, metaphysics out. (L. Sklar, 1992, p. 9) The naïve naturalist pursues metaphysics based on science without proper attention to the significant role that non-empirical features play in theorizing, interpreting, and indeed even making observations and performing experimental interventions. Before looking at how naturalists might sophisticate their metaphysical aspirations in the next chapter, next we’ll consider several objections to the conclusions of the above discussion.

1.2 Objections

There are several obvious responses to the argument offered here. The defender of a naïve naturalism can, first, point out that special relativity is *not* a fundamental theory, and so was never a legitimate candidate for drawing metaphysical conclusions. We ought to, instead, look to our fundamental physical theories, namely, General Relativity and Quantum Field Theory. Second, it might be objected that while my argument does indict *one* scientific theory, I haven’t provided any reason to believe this claim generalizes—perhaps metaphysical conclusions are transparently

²³Note that this sentence is conditional in form—it is always available to simply reject the notion that such a relation exists or is coherent/intelligible.

²⁴Ignoring (see the previous section) the considerable features of traditional presentism which seem to require rejection if one accepts SR. The presentist under consideration here can simply be a ‘Stein-presentist’.

available in the other sciences. Third, and finally, it can be pointed out that *even if* my argument generalizes and in fact *any* attempt to draw metaphysical conclusions from scientific theories will necessarily require the operation of the *a priori*, I haven't offered any reason as to why this is a problem for naturalism.

The first two objections can both be addressed by turning to these more fundamental theories, and lingering briefly to see how the naïve naturalist fares. The final objection will then lead naturally into the next chapter, where we look explicitly at why naturalism rejects the *a priori*.

1.2.1 DOES GENERAL RELATIVITY OFFER A BETTER METAPHYSICAL BASIS?

Special Relativity is not our best physical theory of space-time. That honor lies with General Relativity (GR). The fact that SR offers no definitive metaphysical solutions is not surprising, the objector points out, since we shouldn't expect theories which we *know* to be empirically inadequate to correctly describe the fundamental nature of reality. This line of argumentation is meant to put pressure on the relevance of my conclusion that SR doesn't license metaphysical inferences since its content is metaphysically underdetermined, and any such inference must itself rely on *a priori* contributions. However, this objection presupposes that GR does not fall prey to the same sorts of underdetermination problems. This, however, simply isn't the case—there is at least one example in which GR appears to offer some metaphysical insight but which, upon closer inspection, encounters a host of conceptual troubles. This is in the Substantivalism/Relationalism debate about the nature and reality of space-time as an object (or “container”). In an influential paper John Earman and John Norton (1987) argue that the presence of observationally indistinguishable diffeomorphisms on a manifold presents the substantivalist with a dilemma, one whose only compelling solution is to jettison substantivalism itself.

Earman and Norton offer an argument which is analogous to Putnam's and likewise depends upon non-empirical commitments.²⁵ In Earman and Norton's case they argue that from modest

²⁵I should note, however, that similarly with the case of Stein, interpreting their paper as arguing for a positive metaphysical thesis is going too far. Instead, imagine an enthusiastic naïve metaphysician looking to appropriate Earman and Norton's argument for their own positive thesis.

assumptions about substantivalism, coupled with the mathematical flexibility of the tools used to describe the differential geometry of modern theories of space-time, leads the substantivalist to a dilemma: “Either they [the substantivalist] must reject substantivalism or they must accept a very radical form of indeterminism.” (Earman and Norton, 1987, p. 516) Operating in the naïve mode of naturalistic metaphysics, we maintain that our best scientific theories *dictate* our metaphysics. Thus, we should expect decisive reasons for rejecting substantivalism. Judging from the dilemma presented, it’s clear that we’re expected to think that such a radical form of indeterminism is unacceptable. But notice, in order for GR to offer a purely empirical basis for licensing a metaphysical conclusion it would have to be the case that it would be *irrational* for someone to entertain a belief in a (metaphysically) different account.²⁶ In Earman and Norton’s case specifically the requirement would be to show that substantivalism is incompatible or logically inconsistent with GR—that is, that one could not rationally be a substantivalist and endorse GR. However, this is far stronger than what they do indeed show. At best they do establish that substantivalism is committed to a strong (“very radical”) form of indeterminism about space-time. But is the metaphysical belief in real, yet indeterminate space-time points irrational? Hardly. Are there pragmatic or epistemic reasons for thinking such an interpretation should be avoided? Of course. But recall the supposed force of this objection: it was claimed that analysis of the metaphysical deliverances of GR would *vindicate* the naïve naturalist. Admitting that we have good *practical* reasons for preferring one view over another does nothing to support the contention that our metaphysics may be decisively determined by empirical deliverances of a theory.

None of the foregoing, however, should be construed as a defense of, or argument for, a substantivalist metaphysics of space time. I agree there can be *good* reasons militating against such a view. But as in the SR case, it is a mistake to think that these reasons are there in the ‘empirical content’ of the theory itself.

²⁶Strictly speaking we should strengthen the requirement to the irrationality of conjoined belief in both GR and whatever alternative account is under consideration.

1.2.2 SYMMETRY, INTERPRETATION, AND FUNDAMENTALITY

Neither SR nor GR simply and unproblematically delivers metaphysical conclusions independent of non-empirical choices. Neither theory *dictates* a metaphysics—there is a degree of freedom left open in them which can only be fixed by appeal to something other than empirical observations. However, ‘two cases a universal do not make’—perhaps these failures in space-time theories indicate some pathology special to them. Perhaps, that is, the difficulties arise due to some intrinsic feature of these particular theories—in the case of space-time theories perhaps epistemic considerations might play an undermining role.²⁷ This objection, expanding on the first, holds that while these particular examples don’t seem to successfully carry out the project of naïve naturalism, they aren’t a conclusive argument against it. Instead, we might hold out hope in two alternative fashions. First, we might turn to look at another class of theories, specifically those which deal with the microphysical, e.g. electromagnetism or quantum mechanics. Perhaps here there are metaphysical conclusions which are unavoidable—which are thrust upon us. Indeed, the increasingly bizarre behavior of physical systems on smaller and smaller scales seems ripe for extracting conclusive metaphysical lessons.

The second alternative is to look to future or ‘final’ theories. One might argue that it should be no surprise that our current best physical theories don’t offer ineluctable metaphysical conclusions—we know that they are most likely false. This very falseness is what opens up the interpretational degrees of freedom on which the above arguments rested. This claim, that a final, future physical theory will not leave any room open for interpretational freedom is, while attractive, plagued with difficulties. Though I will respond to objections of the first sort, there is little to say regarding this second kind.²⁸

Returning to the first alternative—arguing that while relativistic space-time theories do not

²⁷That is, perhaps space and time are just too big for us to get the requisite sort of epistemic purchase.

²⁸However, I do have *something* to say about it: First, the idea that a final theory will leave nothing arbitrary or contingent is as much an article of *a priori* faith as any commitment to a privileged present or a substantive space-time. In this sense such positions are self-undermining, since they seek to justify an empirical approach to metaphysics by invoking an *a priori* principle. Second, whatever the prospects of a final theory, the existence of said theory can in no way legitimate or license metaphysical inferences *made now* based upon our current theories, which is precisely the issue under discussion.

conclusively dictate ontology, microphysical theories do offer clear metaphysical lessons independent of any *a priori*, non-empirical, interpretational choice. For example, someone pursuing this objection can point to quantum mechanics: one cannot consistently maintain a local realist ontology if they accept the quantum theory. Thus, the objector points out, metaphysical progress has been made on the basis of a scientific theory!

We should be careful here to distinguish the original aim of a naïve scientific metaphysics, and what the quantum theory purportedly metaphysically licenses. The stated goal of the naturalist metaphysician was to offer a *positive* account of metaphysics based upon our best scientific/empirical investigations. If successful the product would be an account of *what there is* and *what it is like* justified exclusively in terms of the deliverances of science. But note that while undoubtedly a metaphysically salient result from science, the violation of local realism is not itself a positive metaphysical thesis. Indeed, this result shares much in common with the metaphysical result which seems licensed from our above examination of SR. There we concluded that whatever the positive picture of the determinateness of events we may settle on (and which is underdetermined by the theory alone), if SR is even approximately correct then it appears to conclusively *rule out* the traditional picture of presentism which requires a unique, objective foliation of hypersurfaces.

Likewise in this case, the violation of local realism rules out a picture of tiny billiard balls with everywhere-determinate position and momenta and whose interactions are mediated exclusively by contact. But, as the ever-waging debate over the interpretation of quantum mechanics shows, the violation of local realism may be understood²⁹ in more than one way, and which interpretation one adopts is a function of *choice*, which cannot be determined by theory and observation alone.

A particularly apt example of this central role for choice appears in discussions of the interpre-

²⁹It is this criteria of ‘understanding’ which is the bar of adequacy of a naturalistic metaphysics as heretofore described. The naïve naturalized metaphysician is not content to simply say that (for example) quantum mechanics is not a locally real theory and leave it at that. This fact about the theory is something to be *overcome*, to be accommodated and explained. This requirement to offer a singular explanation modeled on the act of description is inherited from traditional analytic metaphysics. Anticipating later chapters, I’ll just comment now that such a requirement does not seem naturalistically motivated, and consequently we should look at what a naturalistic metaphysics becomes in its absence.

tation of physical theories which exhibit symmetries or ‘gauge freedoms’. Indeed the presence of symmetries in both SR and GR are responsible for the respective metaphysical underdetermination found in the discussions above. In the case of SR, the symmetry in question is that between inertial observers, while in GR the gauge freedom of the gravitational potential opens the discussion of substantivalism and indeterminism.

The presence of symmetries in microphysical theories makes metaphysical underdetermination, if anything, more pronounced than in GR and SR. Symmetries appearing in the phase-space representation of a physical system pose an interpretational puzzle which requires some extra-theoretical principles to guide choice. This sort of gauge freedom occurs in classical electromagnetism.³⁰ In the theory of electromagnetism the magnetic field \mathbf{B} can be defined as the curl of the vector potential \mathbf{A} ,

$$\mathbf{B} = \nabla \times \mathbf{A}. \quad (1.1)$$

However, the value of \mathbf{B} is unchanged if one were to add an arbitrary function with zero curl to \mathbf{A} :

$$\mathbf{A}' = \mathbf{A} + \nabla \Lambda. \quad (1.2)$$

Since \mathbf{B} is an observable quantity and \mathbf{A} is not, there is no empirical way to determine the ‘actual’ value of \mathbf{A} . Thus, if we were to ascribe physical reality to \mathbf{A} we would be committed to there being a correct value for the vector potential, one which is in principle unknowable by us. The alternative is to take \mathbf{A} as a mere mathematical convenience, with no connection to any real quantity.

Indeed, this gauge freedom exhibited in specifying the vector potential is a consequence of a more general mathematical characteristic of the integral calculus. To see this, suppose we have an arbitrary function $f(x)$. Integrating over $f(x)$ will give us some new function $F(x)$ related by

³⁰This particular example is ubiquitous in discussions of gauge freedom and the subsequent interpretational issues. The following are a sampling of such discussions: (Baker, 2010; Lange, 2002; Rickles, 2016; Brading, Castellani, and Teh, 2017)

³¹Where Λ is an arbitrary scalar field. The replacement of \mathbf{A} by \mathbf{A}' leaves \mathbf{B} unchanged since the curl of the gradient of a function is zero: $\nabla \times (\nabla \phi) = 0$.

$\frac{dF}{dx} = f(x)$. However, $F(x)$ is not unique, and in fact belongs to a family of continuously many functions which differ by constants C , all of whose first derivatives are $f(x)$. The first derivative of a function does not encode any terms which are constant with respect to the variable of differentiation. As seen in the vector potential case above, this poses several interpretational puzzles when the functions are meant to model a physical system (like the electromagnetic field).³²

The first puzzle is over the metaphysical status of the ‘differential’ relation between a function F and its derivative f if we ascribe some physical meaning to f . Suppose we model some measurable physical quantity by the function f . Suppose also that we *define* this function as the derivative of some other function F . It is typical for realists to assert that f corresponds to, or describes, some physical entity in the world or, simply, that f ‘is real’. Notice, however, that the subsequent status of F is rather nebulous. If we say (as it is defined) that f is the derivative of F , and if we ascribe some physical meaning to f , it might be natural to suppose that we ought to ascribe physical meaning to F as well. Indeed, since in general the derivative measures the *change* in a given function, then we can say that f encodes how F is changing, which implies—at least *prima facie*—that F is real (since one might suppose that to change something requires that thing to exist).

However, it is at least logically possible that no measurable quantity corresponds to F ,³³ and so there is no independent way to observe F (directly or indirectly). The consequence is that we would be committing to add to our ontology an entity whose status as real depends wholly on the ‘proper’ analysis of the relation ‘is the derivative of’. Such an analysis would appear to require appeal to *a priori* principles, if only because it is a piece of mathematical reasoning (which is *a priori* if anything is).

³²In what follows I will offer an abstract example which mirrors the vector potential discussion in many respects.

³³Or indeed, no *determinate* quantity at all! One might notice that the chain of reasoning from the reality of a physically measurable quantity to the reality of its integral closely matches a pattern of reasoning employed by Newton to argue for absolute space (my thanks to Arthur Fine for bringing this point to my attention). That is, from physically measurable (and thus real/absolute) accelerations Newton reasons that velocity must be absolute, and from absolute velocity reasons that there must be an absolute space for this motion to take place in. (For a nice summary of the argument see, Lawrence Sklar, 1974, p. 184) These arguments seem naïve to us now, but it’s worth recognizing their *prima facie* attraction. Nonetheless, recognition of this potentially erroneous train of inference certainly places the burden on those who would argue for the reality of (in the case above) the vector potential (however, the Aharonov-Bohm effect seems to shift this burden in the opposite direction, see fn.34).

The second puzzle, which stems from the first and which has gained the most attention in discussions of ascribing (a correspondence to) reality to various mathematical objects, focuses on the ‘gauge freedom’ which the function F exhibits. The property of f admitting continuously many anti-derivatives of the form $F + C$ is a gauge freedom. Suppose we do think that F corresponds to a real feature of the world. Then we’re in the epistemically unfortunate position of having to accept that there is a fact of the matter about F which is forever foreclosed to us. Alternatively, we might say that all the functions in the family F express one and the same physical situation, in which case F doesn’t correspond to anything real, since it can vary with no change to the physical state of the system.

Here again there is a requirement to make a (conditional) choice—*if* one thinks that f corresponds to an element of reality, and if it is defined as $f = dF$, then either (A) one also takes F to correspond to an element of reality, and so must accept that there are in-principle unknowable facts-of-the-matter about the state of the world (akin to Newton’s absolute space) or (B) deny that F corresponds to anything physically real. Neither option is obviously superior to the other, and regardless, no matter by what metric we assess superiority, such a metric must appeal to something other than the empirical data or the theory under question itself. Accepting the first option (A) is unattractive for the reason of its brazen flouting of Occam’s dictum. And in accepting the second option (B) one will find themselves faced again with the first puzzle, trying to explicate what it means for a physically real quantity to describe the change of something patently unphysical.³⁴

³⁴Regardless of these difficulties, interpretations of the magnetic vector potential face a further ontological challenge: the Aharonov-Bohm effect. While a complete discussion of the effect is beyond the scope of this chapter, a few points are worth mentioning. The Aharonov-Bohm effect purports to demonstrate the reality of the magnetic vector potential A since a charged particle is affected by the current in a solenoid even though the electric (E) and magnetic (B) fields are vanishing in the region of the particle, while the vector potential is non-zero in these regions. Therefore, the reasoning goes, the vector potential *is* a physical quantity with a real effect on the behavior of the electrons. If we are to resist this conclusion and chalk up the effect to the magnetic field, we’re forced to accept that such a field can have an effect where it is not. The reality of the vector potential avoids the specter of action at a distance. However, the vector potential, being unmeasurable, again commits one to an in-principle underdetermination of the actual value of A . A third alternative is to ascribe the effect to the loop integral (holonomy) of the vector potential around the solenoid, but this too has its own difficulties, not least of which is the apparent consequence that the effect is then distinctly non-local. For a more thorough discussion see, e.g., Rickles, 2016; Lange, 2002; Brading, Castellani, and Teh, 2017; Healey, 2007.

In summary: It is far from clear that physical theories other than SR and GR fare better when it comes to offering determinate ontological lessons. Also, symmetries and gauge freedoms appear in theories of the microphysical, and the presence of these seem to require some non-empirical principles to break the metaphysical underdetermination which stands between theory and interpretation. The very act of interpreting such theories would seem to require choices which are not settled by the theory itself, seemingly undermining the naturalist credo that we ought to ‘let science be our guide’.

It would seem that the central difficulty which faces the aspiring naturalized metaphysician is that of *interpretation*. Once one has a theory in hand, the operation of mapping the theory onto entities in the world (i.e. interpretation) appears to be the point at which non-empirical factors enter. One response to this is to argue against the act of interpretation wholesale. This move foreshadows the structuralist approaches to naturalist metaphysics as developed in the third chapter. But before we move on to examine such views there is one last objection to consider.

1.2.3 IS THE *A PRIORI* ANTI-NATURALIST?

In the above discussion of Putnam and Stein’s accounts of the notion of determinateness in the context of special relativity I aimed to show that the choice between their two diverging interpretations would require something more than mere observation or facility with the theory itself. In particular what was required was some non-empirical (and so ostensibly *a priori*) rationale or justification. The consequence was meant to impugn a relatively naïve approach to doing metaphysics based on science—one in which the theory simply dictated the metaphysical picture to be accepted. The implicit reasoning was that if *a priori* principles played an ineliminable role in doing metaphysics, then that metaphysics could not properly claim to be ‘naturalistic’. What I did not do was offer any reason why the naturalist cannot or should not avail themselves to *a priori* knowledge and justification.

It is of course logically possible that one might be a naturalist and still accept some forms of *a priori* justification as legitimate. Or, to strengthen the requirement specifically in terms of

metaphysics—that it is possible for the naturalist to accept that some forms of *a priori* justification license genuine knowledge about fundamental ontology. But while possible, there are reasons for thinking that naturalist metaphysics is committed to denying *a priori* justification such license. It is to these reasons we turn next.

Chapter 2

Naturalism & the *a priori*

2.0 Introduction

Naturalized metaphysics rejects the use of *a priori* principles to ground or justify metaphysical conclusions. What is an example of a ‘metaphysical conclusion’? The claim that the block universe hypothesis is true, or the claim that what exists is structure, not individuals, are both metaphysical conclusions. Their particular attraction for the naturalist, at least *prima facie*, is that their origin and justification is to be found in attention to science, rather than speculation on the basis of intuitions and analysis of commonsense ideas. There are, to be sure, *a priori* arguments for the block universe—these can be arguably dated back to Aristotle’s discussion of tomorrow’s sea-fight. These arguments proceed by analyzing our concepts of time coupled with some principles about propositions. Thus, the traditional defender of the block universe hypothesis might advance an argument like so:

All propositions describe a state of affairs, and any state of affairs either obtains or does not. A proposition is true just in case the state of affairs it describes obtains, and false otherwise. Every proposition, therefore, has exactly one truth-value. Consider, then, the proposition ‘Theseus wins the sea-fight tomorrow.’ Let us suppose that we can analyze away the indexical (‘tomorrow’). In this way we generate a proposition *P*.

P describes a state of affairs, which could either obtain (absolutely, via the elimination of the indexical) or fail to obtain. Thus, *P* is either true or false. But if *P* is either true or false *now*, then (because *P* is stated absolutely) *P* must have that truth value no matter at what point in time it is considered. If *P*'s truth-value is determinate at all points in time, then it must be the case that the state of affairs which *give P* that truth-value must be determinate at all points in time. Thus, the state of affairs which determine *P* are determinate, and so (for us, at time *t*) there are future state of affairs which are fixed. But of course the argument about *P* is perfectly general, so *every* future (and past) state of affairs is determinate—thus, the block universe.

Notice this argument makes no claims which rest upon empirical science, and requires only the conceptual analysis of what a proposition is, when and how such propositions acquire truth-values, and logical extrapolations thereof. This is considerably different than Putnam's argument in the last chapter, where the argument proceeded by way of seeking some relation in accordance with the special theory of relativity itself.

In this chapter we'll examine why the naturalized metaphysician is suspicious of (and ultimately feels the need to reject) arguments like the one above. Briefly, the suspicion arises from two sources. The first source I will label the 'foundations' objection. The argument above is susceptible to responses of the following sort: the opponent of the block-universe hypothesis can quickly identify (often suppressed) premises in the argument which are necessary to render the explicit premises true or plausible, and yet are themselves subject to doubt. Standard objections to the above argument include disputing the account of propositions as bivalent in the way supposed, or arguing against the plausibility of analyzing away temporal indexicals as is required for the argument to work. Disputes of this sort can hit dialectical bedrock quickly, at which point the *rational* basis for either side appears to dissolve into incompatible foundational commitments.

The second source I will call the 'distortions' objection. Whereas the 'foundations' objection is common to many anti-metaphysical traditions (it is, I take it, one of the core objections to *a priori* metaphysics advanced by empiricism), the 'distortions' objection is more particular

to naturalism. It is related to the ‘foundations’ objection, but instead of looking backwards at the source of metaphysical inferences, it looks forward to the content of the metaphysical conclusions themselves. It goes like so: Granting that we all have shared foundations, so that we could advance arguments whose premises were acceptable to everyone, we *still* should not have confidence in the conclusions following from those premises—for a shared hallucination is still a hallucination. The *a priori* starting places upon which we ground our metaphysical conclusions *determine* those conclusions.¹ If it’s possible those starting points are incorrect, then even unanimity should give us no confidence in the conclusions they license. Our *a priori* starting points may *distort* the shape and content of our conclusions—they may systematically mislead us, and foreclose conceptual possibilities.

The present chapter will focus on examining the naturalist motivations for both the rejection of the *a priori*, as well as the continuation of the metaphysical project with ‘naturalistically appropriate’ substitutes. The naïve approach of the last chapter was no improvement over the *a priori* theorizing of traditional metaphysics—we must then ask, what should change in that approach so that a metaphysics based on science will meet the desiderata of the naturalist?

This project is what I will call ‘sophisticated’ naturalized metaphysics, opposed to the naïve naturalism we saw last chapter. There are several intermediate tasks at hand before we can consider a bona fide sophisticated approach in the next chapter. First, in §2.1 we survey some standard examples of naturalized metaphysics rejecting the *a priori* as a suitable *foundation* upon which to build one’s metaphysics, as well as expressing an anxiety about the potentially distorting effects of such a foundation. What these examples leave out, however, is the precisely *naturalistic* reasons why the foundations and distortions worries are relevant.

The remainder of the chapter is therefore devoted to connecting up these anxieties with the conception of naturalism at work for the naturalized metaphysician. §2.2 offers a general account of naturalism, with the aim to characterize the position at once broadly enough to adequately apply to naturalism generally, but precisely enough so that the characterization is more than

¹Indeed, this is just a restatement of the ‘foundations’ objection—different starting points, different conclusions.

merely an empty platitude. It turns out that construing naturalism as ‘taking science seriously’ is sufficient for this task. By taking science seriously (in the sense I elaborate) the naturalist purportedly finds grounds on which to reject the *a priori*.

It will help to get a sense of the thrust of the rejection by seeing how (in §2.3) *a priori* justification is typically employed in metaphysics. Then in §2.4 I’ll offer up two principles which capture the specifically *metaphysical* and *naturalistic* character of naturalized metaphysics, and examine what about these principles ultimately leads to the rejection of *a priori* sources of justification.

2.1 Statements of Rejection

Naturalism claims empiricism as its ancestor, but is not merely a re-labeling of that position. They share a commitment to the idea that justification and knowledge rely upon outward experience generally. But they differ in understanding the character of ‘experience’. Empiricism emphasizes and prioritizes first-person *sense* experience, the inputs and stimulations on our sensory surfaces. Naturalism instead steps beyond our sensorium and takes *science* to be the ultimate bearer of epistemic authority. While the immediate senses and science are alike in their ‘outward facing’ nature, science is not a pure *a posteriori* enterprise. As such, imbuing science with epistemic authority requires naturalism to take on the burden of dealing with the complex, distributed, non-empirical contributions to justification and knowledge which enter into scientific reasoning. Insofar as these contributions don’t exclusively rely upon or consist in direct empirical experience, they involve the *a priori*. However, we typically find naturalists characterizing their project as in opposition to typical metaphysical projects, specifically in opposition to a reliance on *a priori* justification. If the legitimacy of metaphysics (in general) involves recognition of two factors: (a) the presence of a justificatory foundation, and (b) the character of *a priori* knowledge as filling this role, then it’s less than obvious how the naturalist can sustain their *metaphysical* ambitions. In order to see how the naturalist accomplishes this, it will be instructive to see on what grounds they reject the *a priori*.

Ladyman and Ross (2007) devote a large portion of their first chapter to a polemic against *a priori* metaphysics. They call attention to a “tradition which aims at domesticating scientific discoveries so as to render them compatible with intuitive or ‘folk’ pictures of [metaphysics].” (ibid., p. 1) This domestication proceeds by distorting theories and ignoring the latest scientific developments. They argue that this tradition of domestication, which they call ‘neo-scholastic metaphysics’ (see ibid., p. 7 ff.) has reintroduced “esoteric debates about substance, universals, identity, time, properties ... which seem to presuppose that science must be irrelevant to their resolution.”

[These debates] are based on prioritizing armchair intuitions about the nature of the universe over scientific discoveries. Attaching epistemic significance to metaphysical intuitions is anti-naturalist for two reasons. First, it requires ignoring the fact that science, especially physics, has shown us that the universe is very strange to our inherited conception of what it is like. Second, it requires ignoring central implications of evolutionary theory... (ibid., p. 10)

In this quote Ladyman and Ross hint at the two central commitments of naturalized metaphysics which I discuss below: first, a commitment to the idea that there is a universe, and physics ‘shows us’ what it is like; and second, a recognition of the contingency of our evolutionary heritage.

They are not alone in this rejection of the *a priori*, and these two commitments are echoed elsewhere. There is a second dimension of the motivation: a general willingness to interpret scientific theories in a way which asserts the reality of all or some of the theoretical entities posited therein, and to see their results as salient to questions which have historically appeared tractable only by conceptual analysis. Harold Kincaid provides a summation of this general attitude of naturalistic metaphysics:

[It is characterized by first,] an extreme scepticism about metaphysics when it is based on conceptual analysis tested against intuition, and about any alleged *a priori* truths that such intuitions and analyses might yield; and [second,] the belief that scientific

results and scientific methods can be successfully applied to some problems that could be called metaphysical. (Ross, Ladyman, and Kincaid, 2013, p. 3)

The second attitude is an expression of this willingness, a willingness which would seem foreign to many of the positivists of the first half of the 20th century. There is a third sentiment reflecting the naturalist attitude toward metaphysics and the *a priori*—that there is something unconstrained and profligate about permitting the *a priori* to run rampant across our concepts and intuitions. Anjan Chakravartty calls this a ‘slippery slope’ for the naturalist. For example, starting at “the reasonableness of believing in gene transcription” one finds themselves seemingly committed to “the reasonableness of [believing in the] causal relations in virtue of which one is justified in knowing about genes...” (Chakravartty, 2013, p. 28) The naturalist hopes to avoid falling too far down this slope. Why?

[B]ecause one feels that theorizing this far down the slope is simply too far removed from the details of scientific investigation to be of interest to any interpretation of what scientific theories may say about the world... [D]eep metaphysics is too far removed from the details of scientific investigation to yield anything worth having at all... [E]ngaging in metaphysical pursuits too far down the slope is *epistemically impotent*,^[2] and thus a misguided philosophical pursuit. (ibid., p. 29, emphasis mine)

In these quotes we can see the general outlines of the naturalists’ metaphysical revolt: We have empirical reasons to think that nature doesn’t match our expectations of it, and these same empirical reasons might bear their own metaphysical fruit. But if these quotes go some way to understanding *what* the naturalist’s attitude toward *a priori* metaphysics is, they all leave unargued and unexplained *why* this attitude is naturalistically justified beyond the platitudes accusing *a priori* metaphysics of failing to meet the adequacy criteria for scientific knowledge. But

²In fact I think this rather understates the case—the naturalist worries not merely that such pursuits are epistemically impotent, but actively *epistemically harmful*. To borrow an evocative image, allowing *a priori* intuition to be the basis of our beliefs about the fundamental structure of reality is to allow our minds to ‘frictionlessly spin in the void.’ (cf. McDowell, 1994, p. 11) The worry is that there are no external constraints in such a case—it being question begging to insist that ‘reason’ is such a constraint. In such a situation, if we were to think we had bona fide warrant to metaphysical inferences, the picture would be *distorted*.

this, of course, is unilluminating if we're trying to understand what's unscientific about *a priori* reasoning and justification in the first place.

The hostility towards theorizing about the world based on the *a priori* is familiar and unsurprising coming from the naturalist. But familiarity is not the same as a reason, and we need to ask: what's epistemologically worrisome about the philosopher's metaphysical 'speculation' about the world? If our intuitions are *in fact* self-certifying, and if they imply deep requirements on what the world is like, why not avail ourselves of them? To understand the reasoning behind the rejection we'll turn next to spelling out in some detail a general account of 'naturalism' and its relation to science and the *a priori*.

2.2 Taking Science Seriously

What is naturalism? In the introductory chapter I offered up a broad taxonomy of naturalisms. However, the analysis of naturalism there bottomed out at a rejection/suspicion of the *a priori*, without any further examination of how such a rejection was motivated, nor why such a rejection was specifically naturalistic in character. And again in the previous section of this chapter we have seen statements by self-proclaimed naturalists rejecting *a priori* reasoning, but with no clearer or more explicit reasoning as to why such a rejection is licensed.

Naturalized metaphysics rejects appeals to the *a priori* in the service of generating and justifying metaphysical conclusions. This rejection is commonly understood to be related to the scientific character of the naturalist position. This being the case, we should spell out in detail what this connection to science constitutes.

It is uncontroversial to understand naturalism as expressing some sort of general 'friendliness' towards science, but such broad characterization is unilluminating. This is not to say that a *general* characterization will be unhelpful. Indeed, I propose that the appropriate level of analysis at which to begin will take as the correct characterization something arguably as broad as the aforementioned 'friendliness'—*naturalism is the commitment to 'take science seriously'*. There are

two virtues to analyzing naturalism at this distance. First, it is suitably general so that I think no self-identified naturalist would object to this being at least a first-order approximation of their view. Second, and perhaps surprisingly, this characterization is suitably detailed so as to admit of an account of naturalism in concrete terms, and indeed, is detailed enough to allow for at least two distinct understandings of naturalism thus construed.

An analysis of naturalism in terms of ‘taking science seriously’ requires that we have an understanding of what it means to ‘take x seriously’. This is best accomplished by considering two distinctive uses of the phrase. The first sense is that of *believing what x says*. Thus, to take the Bible seriously is to believe that it is a report of the facts, to take eye-witness testimony seriously is to believe that the testimony is descriptively accurate, and—applied to our present concern—to take science seriously is to believe that the propositions of science are true.

The second sense of ‘taking x seriously’ is to recognize the aims, methods, and values of x (whatever it may be) as *important* and *valuable* in their own right. This is the sense of ‘taking seriously’ which applies not to testimony, claims, or propositions, but to persons and things themselves. This is the sense in which you can take (or fail to take) a tradition, ceremony, event, practice, or person or group of people seriously.

On this understanding of naturalism, to take science seriously is to acknowledge the practices which constitute the heterogeneous natural sciences as productively engaged in, and constitutive of, the sorts of inquiry which *matter*. It signifies a commitment to take on the values, aims, and methods of the science as one’s own, to recognize one’s values and aims as instantiated within that practice, and a willingness to embody the norms of inquiry which are exemplified (though not *defined* in any strict sense) by the sciences.

Later (in the following two chapters) I will argue that (1) naturalized metaphysics as typically conceived fails to be appropriately naturalistic in the second sense, and (2) that the second sense is the primary and most important sense in which naturalism should be understood. Setting aside those two claims as promissory notes for now, the aim here is merely to show that ‘taking science seriously’ is not an empty platitude, and in fact admits of a more detailed analysis, while

remaining at a suitably abstract level to apply generally to ‘naturalisms’ as a whole.

The point here is that the naturalist’s rejection of the *a priori* must be linked in some way to this commitment to take science seriously. The precise sense in which they are linked must be in terms of an incompatibility in the way that the *a priori* is deployed by the metaphysician, and what taking science seriously implies about the *a priori* in general.

2.3 Why Does Naturalism Reject the *a priori*?

There is nothing obviously problematic about viewing *a priori* intuitions as genuinely justificatory. Indeed, quite the opposite has been typically presumed throughout the intellectual history of the West. Empiricists like Locke and Berkeley availed themselves of (a limited class of) propositions known *a priori* and saw this as unproblematic. There’s a prima facie burden on the naturalist to explain why helping oneself to *a priori* justification is to be avoided. Further, we should expect that whatever the explanation is, it cannot come by way of non-empirical justification (on pain of contradiction). Thus, the reasoning should proceed by way of empirical reasons, and more specifically for the naturalist, through attention to science.

2.3.1 THE ROLE OF THE *A PRIORI*

The distinction between *a priori* and *a posteriori* propositions is an epistemic distinction.³ To classify a proposition as *a priori* is to assert something about the way in which it is justified and its status as knowledge. As such *a priori* propositions occupy a central role in two distinct modes of

³As has been pointed out elsewhere (see, for example BonJour, 1998) this distinction has often been assimilated (or conflated) with the metaphysical distinction between necessary and contingent propositions, as well as the semantic distinction between analytic and synthetic propositions. While I have nothing of substance to say here about the connections between these three categories, it is important to note that by their very connection to these other notions (whether deserved or not), the distinction between *a priori* and *a posteriori* propositions deserves clarification. In various modes, and for various reasons, the *a priori*, the necessary, and the analytic have all been related either as synonyms, or to explain one or two in terms of the other. For example, one might seek to explain the analyticity of a particular proposition (and the characteristic semantic ‘emptiness’ of such a proposition) in terms of its necessity—it’s true in *all* possible worlds and thus can do no work in helping us discover *which* possible world is actual.

inference.⁴ First, *a priori* propositions may serve as premises in deductive arguments—this is what we might call their ‘classical role’ in rationalist epistemology. Second, they serve as evaluative criteria in assessing competing explanations in inferences to the best explanation (IBE). By virtue of the character of *a priori* propositions—that they are, variously: true by inspection, innately grasped, self-certifying, etc.—they occupy a position of security denied empirical propositions, and so play a singular role in epistemic justification generally.

DEDUCTION

Often, *a priori* propositions are asked to serve the role of epistemic guarantor in inquiry. In this domain, where the explananda outrun experience, *a priori* truths serve as foundations, or ‘fixed points’ in our constellation of belief, from which we can apply the apparatus of inference. The problem of induction, and the difficulty of developing a theory of confirmation in the time since Hume, have lent credence to the idea that no empirical observation, nor any number of such observations, can logically entail a general or universally quantified proposition, and thus can’t serve as indubitable premises from which to derive general truths about the world. Yet, we recognize the logic of deduction to be the most robust instrument of inference possible, and the machinery of deduction sits idle without substantive generalizations on which to act. A deductive system of inference whose only available premises are particulars cannot license inferences beyond those particulars. What inquiry demands are premises consisting of substantive general propositions which can generate significant and surprising consequences not obvious from examination of the original premises. The paradigm example here is mathematics: from ur-premises (called *axioms*), combined with a judicious application of the rules of inference, fruitful consequences can be derived.⁵

The strength of the mathematical example lies precisely in the status which axioms are thought

⁴I am not, however, making any claim that these applications are the only ones to which philosophers have historically put the *a priori* to work. The reason I single out the two I do here is their central role in philosophical and metaphysical inquiry, as well as their role as tools of inference in the sciences.

⁵The classic example here is Euclidean geometry, with its axioms from which all the theorems can be derived. Peano arithmetic is another example of a rich body of theorems which follow from a small set of general propositions.

to enjoy—they are ‘self-evident’ or are capable of being seen to be ‘true by inspection’. That is, the axioms are thought to be known *a priori*, in no need of empirical justification, and immune from even the very possibility of such justification. Axioms are (typically) universal generalizations, and so are suitably fruitful for the mathematician. The trust we can place in the theorems of mathematics rests on the security and simplicity of the axioms. The metaphysician sees non-mathematical axioms—similarly *a priori*, but generalizations about properties, substances, space, time, identity, etc.—as promising the same robust set of premises from which to derive bona fide knowledge about the external world. This is the picture of rationalists like Spinoza and Descartes, who attempt to generate knowledge about fundamental reality from (the appropriately named) first principles.⁶

As an example, consider a paradigm proposition thought to be known *a priori*:

(a) No surface can be red (all over) and green (all over) at the same time.

Notice first that it is a universal statement.⁷ Second, I ask the reader to consult their own intuitive sense of its truth-value, or better yet, to actively test it by entertaining the thought that there exists a counterexample. Can you do it? I’d wager you cannot—it appears obviously true, or true by inspection. What use can such a proposition be put to? Suppose we also know that,

(b) There exists an object whose surface is both red (all over) and green (all over).⁸

From these two premises it immediately follows that the object must exist at two distinct times, which of course implies further propositions about the identity and existence of objects over time (and perhaps facts about the structure of time itself). Conclusions about the perdurance versus endurance of objects, or about identity and qualitative change over time might then be derived. Neither (a) nor (b) could be derived from experience alone, and it is doubtful that there is any experience that could lend unambiguous observational *proof*. The first proposition, (a), being a universal, outstrips empirical support for the standard reasons. But even (b), an existential, is not

⁶It is no mistake that Spinoza seeks to expound his philosophy in the ‘geometrical method’.

⁷Technically, the most natural way to translate this into first order logic is as a negated existential. But under the usual semantics this is logically equivalent to a universal.

⁸Admittedly a more contentious proposition!

amenable to decisive empirical demonstration, since the skeptic about identity across time could always understand the temporally distinct entities to be two distinct objects.

Despite the epistemic appeal of such a deductive scheme of inference, it is typically seen as an overly demanding criteria for most types of knowledge. Proposition (a), while difficult to imagine as false, has perhaps a less secure status than (some of) the axioms of mathematics (like those of set theory or Euclidean geometry); and it is certainly the case that the second proposition hasn't the same intuitive appeal and might reasonably be doubted. It is plain in observing philosophical discourse that interlocutors may (and often do) reasonably disagree about the supposed truth of purported *a priori* propositions. While unanimity is not a guarantee of truth, plurality in views about such propositions is certainly an indicator that *some* ostensibly *a priori* propositions might in fact be false. Thus, there has been a move towards an attitude of fallibilism in regard to even our *a priori* knowledge. (See BonJour, 1998 for a developed account.)

INFERENCE TO THE BEST EXPLANATION

The move to fallibilism lessens the importance of the deductive role for metaphysical inferences, especially when coupled with the idea that first principles must be not only known *a priori* but must be indubitable. Since deduction is only as trustworthy as its premises, if we are fallibilists about the premises it's hard to see what the apparatus of deduction adds.⁹ The embrace of fallibilism about *a priori* knowledge has brought with it the demand for an alternative to deduction which can serve as a licenser of metaphysical inferences (even if it must fall short of *guarantor*).

Inference to the best explanation (IBE) is just such an alternative to deduction¹⁰—a method of inference which seems to straddle the ampliative virtues of induction and the epistemic robustness of deduction. IBE is the inferential scheme which licenses inferences to the truth of whatever

⁹The whole reason that deduction appears to be an attractive tool for inquiry is that validity *guarantees* true premises will produce true conclusions. If our premises might be false, then we have no way to predict the truth value of their consequences purely by examining their logical form. But if this is the case, we'd have to apply looser standards of evaluation in order to assign credences to our premises, and thereby assign credences to their conclusions. But this isn't a significant improvement over other, non-deductive methods of inference—we lose all sense of the original virtue of deduction's epistemic guarantee.

¹⁰This is not to say that IBE is a *justificatory* alternative if this is taken to mean it could justify the same inferences and to the same degree. Rather, it is an alternative in the sense that it is another *form* of justification.

best explains a particular explanandum. That is, as Gilbert Harman puts it, in IBE “one infers, from the premise that a given hypothesis would provide a ‘better’ explanation for the evidence than would any other hypothesis, to the conclusion that the given hypothesis is true.” (Harman, 1965, p. 89) To give a simple example, IBE licenses an inference to the truth of the proposition “Jane left the kitchen light on.” just in case that proposition being true *best explains* the particular phenomenon in question (in this case the kitchen light being left on) and the corresponding explanandum “Why is the kitchen light on?”

By the very nature of this inferential pattern the conclusion does not follow with certainty, but we may have a degree of confidence in the conclusion based on the quality of the explanation. From this, of course, the need then arises to account for the nature of explanation. Daniel Campos explains Harman’s account like so:

[W]e infer hypotheses that explain the evidence better than any of the alternatives. This opens up the possibility of arguing that although we may not be able to determine absolutely what an explanation is, we can at least make relative judgments regarding better or worse explanations. That is, when confronted with two or more alternative explanatory hypotheses for the data, we can make judgments about their relative merit on the basis of some relevant criteria. (Campos, 2011, p. 432)

It is at this point—where we are asked to make judgements of relative merit based on ‘relevant criteria’—where the *a priori* plays a role. In order to evaluate whether Jane’s actions (or lack thereof) *best* explain the explanandum, I need at least one alternative explanation and some criteria which are deemed suitably relevant to the judgement required. What alternative explanations are available to us, as well as our suite of contrastive criteria, depend on what background assumptions we make, and what super-empirical virtues we think are appropriate for evaluating hypotheses (embedded in their surrounding theory).

I might entertain the competing hypothesis that a secret government agency picked my lock, snuck into my house, turned on the light, and then left without a trace. But I rule out this hypothesis on the basis that it is a *worse* explanation than that Jane simply forgot. Why is it worse?

There seem to be two paths we can follow in answering this. The first is an appeal to enumerative induction over past similar cases, in which instances where the light was left on were not caused by the deep-state, but instead by simple forgetfulness. But this route inherits all the ills of the problem of induction, and doesn't seem to get us very far since such a justification would be useless if this was the first observed case of the light being on.¹¹

The second path is to appeal to the fact that the government hypothesis is not *parsimonious*, not suitably *simple* or *elegant*, etc. We appeal to the idea that the government hypothesis has “too many moving parts”, or that it “requires an implausibly vast and elaborate conspiracy”. We appeal to super-empirical virtues which help us weigh competing hypotheses.¹² But these virtues themselves either have the status of *a priori* principles, or have as constituent parts *a priori* principles.

I am not trying to give an exhaustive account of IBE, nor argue in its favor or against it. My point is just to highlight the role which the *a priori* plays in cases of IBE. The worry of the naturalist, then, is about the distorting and misleading effects IBE might have in the hands of the analytic metaphysician who helps herself to purported truths known *a priori* which can then serve to license hypotheses as the best explanation of phenomena in a distinctly naturalistically implausible manner.

2.3.2 AN EPISTEMOLOGICAL FOUNDATION

To sum up the discussion of the *a priori*, its role and utility in traditional metaphysical inquiry is to serve as a foundation to build upon or a yardstick to measure by. Metaphysics (as discussed in the introductory chapter) is specially concerned with proffering explanations about precisely that which cannot be accessed through experience. If we didn't have these fixed points of epistemic certainty (or, more humbly, confidence) we'd be thrust into a thoroughgoing skepticism about the possibility and even coherence of metaphysics. But, (let us suppose) metaphysics *is* possible

¹¹What's more, if we follow Harman (1965) and think that IBE undergirds every case of enumerative induction then we have just reasoned ourselves into a fairly tight circle.

¹²The weighing would become significantly more difficult if the competing hypothesis were to the effect that *I* forgot to turn the light off.

and so, *ex hypothesi*, such fixed points *do* exist. The metaphysician thereby legitimately employs these fixed points in proffering and justifying metaphysical claims.

Two factors contribute to the apparent legitimacy of relying on our *a priori* intuitions to justify purported knowledge about fundamental reality. The first is just the condition above—that without fixed points from which to begin our justificatory practices, we’re left with a skepticism which we have antecedently ruled out. The second factor is simply the recognition that (at least some of) our *a priori* intuitions appear to be plainly true, indeed necessary, and are thereby above epistemological reproach (even if we allow that they fall short of strict indubitability). This combination of a need for epistemological bedrock, as well as the presence of self-certifying propositions which ostensibly fit the bill makes for a clear license to metaphysically theorize.

For the naturalist, these two factors (again: (1) we need epistemological fixed points for metaphysics to be possible, (2) by their very nature we have *a priori* access to such fixed points) are the fundamental source of tension in the naïve metaphysics of Chapter 1, and must be dealt with if one is to provide a coherent naturalized metaphysics.

2.4 Two Principles of Naturalized Metaphysics

Here I’ll offer two principles which broadly characterize the position of naturalized metaphysics. That said, the diversity of positions which identify as ‘naturalisms’, and the diversity of metaphysical positions within that group is great enough that to think that everyone who identifies as a naturalized metaphysician would identify with these principles as stated and with no qualifications would be foolish. Nonetheless, they are broad enough that to deny them would largely disqualify the position from either being considered a naturalism, or a metaphysics. Their conjunction does seem to constitute the basic conditions for a naturalized metaphysics.

1. METAPHYSICAL REALISM: There is an external world, and it is in principle knowable. The aim and measure of success in ontology is proximity to the *true* description of *reality*.
2. EPISTEMOLOGICAL NATURALISM: Our belief forming mechanisms, and the innate beliefs we

possess, must be understood in the fullness of their contingent, diverse, evolutionary context. Human cognition is contingently adapted to reliably track a narrow range of phenomena and should not be trusted outside of that range. As such, human cognitive faculties have no *special* or pre-theoretical claim to correctly anticipating the true description of reality.¹³

The first principle characterizes the way in which a naturalized metaphysics can be coherently understood as a metaphysical project. It is the same principle which analytic metaphysicians accept as a fundamental commitment. Both the analytic and naturalized metaphysician agree on the aim of metaphysics—to describe as accurately as possible the true structure of reality. Both would reject Kantian re-interpretations of the aim as an internal study of the conceptual preconditions of thought and experience. Both would agree that ideal inquiry would approach an end state where our mental model of the world perfectly mirrors reality as it is.

But this is as far as their similarities along this dimension extend. Whereas the analytic metaphysician sees our access to innate and *a priori* truths as fundamentally trustworthy in these pursuits, the naturalized metaphysician is dubious of such purported truths due to the second principle above. There are two ways in which this principle makes contact with naturalized metaphysics.

The first way that epistemological naturalism is relevant to the question of the trustworthiness of *a priori* intuitions has to do with the particular relevance of our best theories of our own biological legacy. Naturalized epistemology emphasizes the *contingent* aspect of our cognitive faculties and the role that *science* (especially the life sciences) has played in this understanding: geology, psychology, evolutionary biology, etc., have contributed to a concrete, embodied account of our evolutionary history. Our primate ancestors evolved in a certain way, dependent upon climatic events, geographical facts, ecological changes, etc. Any of these factors in our evolutionary lineage could have been different, and given such differences we have *no reason* to

¹³Naturalized epistemology is advocated (if not invented) by Quine (1969; 1953), and this sentiment is strongly echoed in Ladyman and Ross's quote above: "Attaching epistemic significance to metaphysical intuitions ... requires ignoring central implications of evolutionary theory..." (Ladyman and Ross, 2007, p. 10)

think our cognitive faculties (let alone our sensory modalities) would have developed in a way resembling anything like our own. If the character of cognition is contingent, we should be wary of purported truths given in that cognition. Science, it appears, tells us reality is far more alien than our mental models might presume and that our cognitive capacities are fragile epistemic things.

Second, as already noted, part of the naturalist spirit is that the suspicion of the *a priori* should be motivated by specifically *scientific* considerations, as opposed to more broadly empiricist or skeptical reasons for distrusting *a priori* knowledge. While these latter reasons perhaps play a part in such motivations, the naturalist considerations are more firmly rooted in episodes from the history of science—most typically episodes in the physical sciences. Purportedly *a priori* propositions about the structure of reality and the furniture of the world have been shown false (or minimally, not unique in their empirical adequacy) by the robust results of our best science.

Thus, naturalized epistemology is motivated not only by a pro-scientific *spirit*, but also by the deliverances of bona fide scientific *results*. Naturalized epistemology represents a skepticism about our contact with reality but, importantly, it is not a totalizing epistemological skepticism. Rather it is a skepticism built on the epistemic success of science. The naturalist must offer concrete episodes in the history of science that are supposed to offer a positive reason to discard previously held beliefs and concepts which had appeared so secure that their denial was unfathomable. The two paradigm cases in the past 100 years are the emergence and development of relativity and quantum theory.

2.4.1 SCIENCE AND CONCEPTUAL REVOLUTION

GENERAL RELATIVITY

Perhaps the most affecting episode of this sort in the past century was the transition from physical theories of space and time to the general theory of relativity (GR). The reason I suggest this was a singular event is not its strange consequences for matter, time, or its theory of gravity. Rather, GR's deepest impact is to be seen in its broad conceptual effects for mathematics, science, and

philosophy. The advent of GR carried with it a shift in attitudes toward abstract and physical geometry, a shift that was historically more radical than its empirical predictions, displacing a tradition which traced back to at least Euclid. Historically, Euclid's *Elements* was not merely a book of mathematical theorems—it was understood to be a true description of the structure of physical space. The relations in real, physical space were understood to behave in conformity with all Euclidean theorems. Not only was it doubtful that alternative *mathematical* systems of geometry could be imagined and constructed, but the idea that *physical* geometry could be other than in accordance with Euclid was unthinkable. The axioms, definitions, and theorems of the *Elements* were so plainly true, and so simply and ubiquitously confirmed in experience and practice, that there could be no doubt as to their unassailable status (however, see fn.15).

This security and universality of the physical reality of Euclidean space was canonized in the physical sciences by Newton, and was similarly placed beyond question in the canons of philosophy by Kant. Kant saw the Euclidean space of Newton as so central to our fundamental conceptual capacities he argued that the very possibility of conceiving of space carried with it the structure of Euclid's geometry—concepts of objects in space are concepts of objects in Euclidean space.¹⁴ The relation between the mathematics of geometry and both our intuitive as well as Newtonian understanding of space was powerful evidence for the fundamentally *a priori* intelligible character of the external world.

But beginning in the early 19th century, mathematical study of the parallel postulate¹⁵ and attempts to prove it from the other axioms led to alternative, non-Euclidean geometries. The confidence in the uniqueness and necessity of Euclid in the abstract domain of mathematics wavered,

¹⁴I should note here, notwithstanding the argument that follows, it's not clear that Kant was wrong if we understand this as a descriptive psychological claim. It may very well be that we, as humans, presuppose a flat geometry when we conceive of objects arranged in space. This of course says nothing about the *actual* structure of physical space. This is, however, just more grist for the naturalist's mill, since the most plausible way to explain why we have an innate disposition to think in a Euclidean manner has to do with the locally flat structure of space at human scales.

¹⁵The investigation of the parallel postulate is an interesting episode in the history of the assessment of *a priori* knowledge in its own right. Euclid famously suspected his 5th postulate was not as secure as the first four, and proved as much as he could without it. Euclid's suspicion is indicative (as evidenced by the eventual discovery of non-Euclidean geometries) of an intuitive fluency with regard to estimating the justification of purported axioms, and should be considered an impressive success in the history of mathematical hunches.

but these alternative geometries were bizarre and fit neither our spatial intuitions, nor our experience. However, the assurance that reality was Euclidean was finally undermined by Einstein's theory, which marshalled the mathematical apparatus of this new geometry to produce a theory more adequate than Newton's.¹⁶

The new physics represented concrete evidence that what we had previously understood as *necessary* was—incredibly—*false*. This transition is not special to this particular revolution, and indeed this conceptual transition is diagnostic of such revolutions.¹⁷ However, this transition is notable because of the magnitude of the conceptual shift. Acceptance of the new physics carried with it a lurking anxiety—if we were so mistaken by relying upon our (seemingly unshakable) intuition, about what else might we be radically misled? This anxiety appears best assuaged by hoping for a more secure source of knowledge about the deep structure of reality than our intuitions. This source is close at hand—it was *science* which helped the Euclidean scales fall from our eyes! Science, the naturalist concludes, must be given pride of place in questions about the external world.

THE QUANTUM THEORY

The second deeply affecting episode of the 20th century was the acceptance of quantum mechanics as an adequate theory of the microphysical. Since Democritus there has been a powerful intuitive appeal to the notion of the material world composed of individual entities whose simple properties and relations determined the character of phenomena at larger scales. These individuals were modeled on the paradigmatic 'things' of experience: spatially and temporally stable, discrete, concrete objects. Fundamental reality was modeled on things such as stones or billiard balls, rather than things like (the phenomenal entity) water—that is, for whatever cognitive reason, humans (or at least philosophers!) have tended to take count-nouns, rather than mass-

¹⁶That is not to say, of course, that there weren't objections and alternatives proposed. As is well known, Euclidean space can be preserved by instead positing appropriate forces to accommodate the predicted observations.

¹⁷Van Fraassen takes this transition to be one of the central puzzles with which we (21st century philosophers of science) must come to grips. (van Fraassen, 2002, p. 64 ff.) It is also a clear example of the Kuhnian phenomena whereby the new paradigm scrupulously reinterprets the content of the previous paradigm so as to make the past beliefs intelligible, but also so the shift appears inevitable.

nouns, as ontologically primitive. Concomitant with this understanding of ontology in terms of concrete individuals is a modeling of the dynamics of reality on the dynamics of those same stones or billiard balls. Thus, just as a stone thrown in the air traces a continuous path through space, the fundamental objects which compose the world trace continuous paths. Indeed, we then (unwittingly allowing the ontological tail to wag the empirical dog) find that we can explain the continuous trajectory of the stone *in terms of* the particles which compose it!

Developments in physical science in the late 19th and early 20th century led to the quantum theory. Once instruments were delicate enough, and once our questions were appropriately focused, the phenomena produced in experiments began to diverge from those expected from classical mechanics. The upshot of these experiments and their theoretical articulation was that somehow we couldn't describe the behavior of the posited fundamental particles if they were modeled in ways which had proven successful (both conceptually and practically) since the early atomists. For example, it appeared that there was no way to describe the particles as having both a continuous trajectory in space *and* to be affected by strictly local influences. On our intuitive notion of objects, not only does the tossed stone have a well-defined position at every moment, but the only possible effects on the stone are those transmitted to the stone through direct contact. The failure of this conjunction of properties (which were taken to be—at least partly—*constitutive* of the concept **object**) marks the crisis in the compatibility of quantum mechanics and our intuitive understanding of the microphysical world.

SHAKY FOUNDATIONS

The effect of these discoveries was again to shake our confidence in what appeared to us as necessary, indubitable, undeniable truths about the nature of the physical world. Notice that the character of the challenge to *a priori* justification which GR and QM pose are somewhat different. The realization that physical geometry wasn't necessarily Euclidean gave us pause, and inductive reason to suspect that other apparent fundamental truths might be similarly wrong. But the non-Euclidean geometries which were developed and applied to relativity were certainly intelligible—

though our three-dimensional thinking struggles to extrapolate from visual models of curved space, the fundamental notion of a curved space doesn't seem to altogether defy conceivability; only expectations.

The parallel realization in quantum mechanics—that we could not consistently describe (let alone *picture* or *imagine*) the behavior of the microphysical world in what seemed like the very concepts required for *coherent thought* about objects—was not succeeded by a coherent replacement. There is no standard, intelligible answer to the question of how to think of the microphysical—every attempt at interpretation trades some basic presuppositions of our conceptual scheme in order to preserve others: Bohm gives up on locality to preserve continuous trajectories, Everett dismisses alethic modality to explain superposition.¹⁸ Whereas GR implies a world alien to our innate sense yet intelligible, QM suggests a world which is—at bottom—truly ineffable, a world which fundamentally outruns the basic mechanisms of thought which were forged in the trees and savannahs of human pre-history. The recognition of our epistemic position with respect to what our best science suggests the world is like adds the final motivation to the naturalized metaphysician's acceptance of epistemic naturalism, and a commitment to an epistemic humility which reaches to the very core of our expectations of reality.

2.4.2 NATURALISM AND THE 'QUINEAN REALIZATION'

This humility about our innate sense of what the world must be like, coupled with the felt need to offer a genuine metaphysical accounting of the world, leads the naturalized metaphysician to what I'll call the 'Quinean Realization'. This is the realization that what we know about our cognitive faculties and their holistic plasticity is fundamentally in tension with the traditional picture of how metaphysical beliefs can be justified. In that picture, *a priori* truths serve as a foundation on which we build the super-structure of metaphysical knowledge. But rejecting the epistemic status of the *a priori* leaves a vacuum in justification—what can serve the role of licensing metaphysical inferences and beliefs in place of these *a priori* fixed points? Ladyman and

¹⁸And Bohr appears to give up entirely.

Ross put the Realization thusly: “individuals are blessed with no epistemological anchor points, neither uninterpreted sense-data, nor reliable hunches about what ‘stands to reason.’” (Ladyman and Ross, 2007, p. 29)

The commitment in epistemological naturalism is tied to a picture of our epistemic faculties which was famously described by Quine: the web-of-belief. In Quine’s web (see fig. 2.1) our conceptual scheme is imagined as impinging on the external world at the periphery, where brute experience constitutes the sole boundary condition on belief. The ‘nodes’ (beliefs) in the web are connected by their logical relations such that changes in beliefs ramify through the web, changing and modifying other beliefs and their interconnections.




Figure 2.1: A schematic illustration of our web of belief.

Concomitant with this picture is the understanding that no node in the web is a fixed point which would furnish a reliable and immovable foundation around which one might orderly arrange belief.¹⁹ This Quinean picture gives up on the notion that we all have at the core of our conceptual scheme access to the same, universal, unshakable knowledge. We, as biological organisms, have those beliefs contingent upon our niche and circumstances. The naturalist suspicion

¹⁹It’s worth quoting the authority here: “Reevaluation of some statements entails reevaluation of others, because of their logical interconnections—the logical laws being in turn simply further statements of the system ... no statement is immune to revision.” (Quine, 1953, pp. 42–43)

of *a priori* knowledge can be seen as the recognition that different arrangements in the web are *possible*, and that different arrangements mean wholesale differences in conceptual scheme, and thereby beliefs formed within that scheme—including those constitutive of our thinking about the world.

The Quinean Realization is just this recognition that there are a plurality of ways our beliefs could be reconciled within the web, with no single method of getting our beliefs to cohere being a more correct way to think of the world than any other (on pain of begging the question). The Quinean Realization also includes the awareness that this situation in epistemology, if granted and left alone, is nothing less than a total and radical skepticism which would undermine not only metaphysics, but with it realist philosophy of science generally. If metaphysics is to be pursued in the absence of the *a priori*, but still with the aim of delivering an (approximately) true description of reality, naturalism must offer a substitute—something from which metaphysical inferences follow, something which may license metaphysical beliefs.

2.4.3 A NATURALISTIC SUBSTITUTE

On the basis of the Quinean Realization, then, the naturalist who remains committed to the possibility of metaphysical knowledge is faced with a challenge. On one hand, taking our cognitive capacities in their biological context suggests that we should doubt the deliverances of those capacities in domains outside our evolutionary niche—this includes metaphysical knowledge. On the other hand, the very scientific tradition which has led us to the above conclusion has within itself suggestions of a metaphysics. Indeed, these are mutually reinforcing since the metaphysics suggested by science appear to be radically different than the assumptions which have characterized much of metaphysics since Democritus.

But while the strangeness of the picture which physics gives us and the picture which our meso-scale conceptual naïveté offers mesh at the level of a plausible psychological story, the naturalist is not free of the fundamental philosophical problem—offering a justification for something which can serve as the fixed-points from which to theorize. Again, it is a caricature, and

a significant one, to imagine that the conclusions of science appeared fully-formed and worked-out from the brute observations we make of the world. As already discussed above and in the introductory chapter, it is a truism that scientific investigation is ‘theory-laden’. There is no such thing as a theory-neutral observation language, free from the linguistic and conceptual trappings of our culture and context. The naturalist would do well to remember that the rejection of the *a priori* which is recommended by naturalized epistemology is thoroughgoing and applies equally to our scientific investigations no matter the empirical connection these enjoy. This, ultimately, is the content of the Quinean Realization—that the conceptual strata on which we operate in both the mundane and grandiose parts of our conceptual lives is less bedrock than shifting sands, and there is no Archimedean point against which we can apply our lever.²⁰

The naturalist cannot consistently argue that the *a priori* principles of the analytic metaphysicians are wrong, and that we just ought to use *their* (the naturalist’s) principles instead. Without further explanation of the origin and justification of those principles, nothing separates them from the *a priori* principles save for the contingent fact that scientists and scientifically inclined philosophers have found them to be more conceptually appealing, more useful for their own interests, etc. A scientific metaphysics cannot proceed from *no* principles at all, but if what is objectionable about the principles employed by the analytic metaphysicians is their *a priori*, and thus unconstrained, character, the naturalist must do better in this regard than simply substitute their own *a priori* principles, no matter how ‘natural’ they are claimed to be.

Since the very possibility of metaphysics requires fixed points from which to start, and the naturalist cannot justify the fixity of such points by asserting their intuitive character, they must offer a suitable substitute upon which to build a metaphysics. They must offer their own fixed points, epistemological anchors which will resist the fluidity of Quine’s web, offer some stability in the tide of interconnected and interdependent belief.

No philosopher currently pursuing a naturalized metaphysics of the sort under consideration endorses so naïve a picture of science as that under consideration in chapter one. They recognize

²⁰Lewis Carroll (1895) helpfully reminds us that even *modus ponens* is ultimately circular. One who does not accept it as a law of reason cannot be *rationally* compelled to accept it.

both the *a priori* component to science, and the need to correct or compensate for its potential distortions.

In the next chapter I will evaluate one of the most popular and successful naturalistic metaphysical programs—(Ontic) Structuralism. Structuralism can be understood as a direct response to the problems of the naïve approach to a scientific metaphysics. This sophisticated naturalism is motivated by the sorts of concerns introduced above. Structuralism seeks to replace the content of a traditional *a priori* approach to metaphysics by identifying fixed points which are licensed by science. From these fixed points are supposed to follow metaphysical conclusions which enjoy the epistemological bona fides of science, avoiding the distorting effects of anthropocentric *a priori* biases.

Chapter 3

Ontological Structuralism

3.0 Introduction

The criteria for finding a suitable fixed point upon which to base naturalistic metaphysics are two: first is the negative constraint that the justification for such a point cannot be that it merely *seems*, *feels*, or *appears* true, or that it cannot be *imagined* to be false. As is clear from the emphasis, such appeals baldly rest on the cognitive faculties which evolutionary biology (etc.) have given us a reason to distrust. The second criterion is a positive constraint requiring that the appropriate source of justification for whatever fixed points we might find satisfies some naturalistically acceptable conditions. We should expect the sources of justification—our fixed points—to come from taking science seriously. The task for the naturalist then, is to articulate fixed points which meet such requirements. One might try to identify a subset of the *a priori* which not merely *feels* true, but more importantly has the appropriate naturalistic justification. This is the strategy of ontic structuralism.

3.1 A Sophisticated Naturalism

3.1.1 ORIGINS OF STRUCTURAL REALISM

Ontic structuralism finds its genesis in debates about scientific realism. A realistic interpretation of a theory involves taking its central non-observational terms to be genuine referring terms. An approximately true theory can be understood as imperfectly referring to (and hence approximately describing) the real entities in the world along with their behavior. The principal anti-realist argument against this idea is to point out that over the history of science there is little ontological continuity across theory change, and often none at all. The humors, phlogiston, caloric, and aether of past theories were not carried over (even approximately) to the theories replacing those for which they played central roles. According to realism, these theories—being the best theories available—committed us to belief in such entities, entities which we now have very good reason to believe do not exist. Thus, in historical cases the realist approach to theory interpretation gives the wrong answer. The anti-realist then projects forward, pointing out that our epistemic position with respect to our current theories is no different than in those historical cases, and we thus have no principled reason to think our theories are (approximately) the final, correct, descriptions of reality save for a blinkered historical chauvinism. Indeed, we have very good inductive reason to think our sciences will ultimately undergo the same sort of theory change as the historical cases. This anti-realist argument is called the ‘pessimistic-induction’ and has become a yardstick against which to measure the success and plausibility of competing realist theories.

The pessimistic-induction against realism leverages the discontinuity in putative ontologies between successive theories. Structural realism was born out of attempts to articulate the ways in which there *is* genuine continuity between theories whose ontologies are *prima facie* incompatible. Introduced by John Worrall (though Worrall acknowledges historical antecedents which could be considered versions of structuralism), structural realism refocuses attention to the *structure* of a theory rather than the interpretation of its central theoretical terms. (Worrall, 1989)

While the putative entities of a theory are often jettisoned in its successor, it is rare for the mathematical structure of the theory to be likewise removed. Indeed, one of the procedural norms in theory revision and the acceptance of a successor theory is that the structure of the theory to be replaced is preserved, often as a ‘limiting case’ in the new theory.

A standard example is the relation between the Galilean and Lorentz transformations. The relation of two inertial frames in Newtonian mechanics only depends on their relative velocity:¹

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

whereas the relation in relativistic mechanics depends additionally on the Lorentz factor:

$$x' = \gamma(x - vt) \tag{3.2}$$

where

$$\gamma = \frac{1}{\sqrt{1 - \frac{v^2}{c^2}}}. \tag{3.3}$$

When the difference between the speed of light (c) and the measured velocity (v) is large, the contribution of the Lorentz factor is small and approaches unity as we allow c to approach infinity. We say that ‘in the limit’ as the speed of light grows arbitrarily large the Lorentz transformation approximates or ‘reduces to’ the Galilean transformation.

The relevance to realism of this limiting behavior in the equations of a theory and its successor is twofold. First, in the example above this relation and its dependence on the relative difference between measured velocities and the speed of light provides a psychological/historical explanation of why our scientific forbears believed the former equation accurately described reality. The

¹What is ‘structural’ about equations (3.1) and (3.2)? These equations identify a class of models under certain coordinate transformations. (In the Galilean case the identity condition is straightforward: identify models which preserve angles between unaccelerated worldlines.) The claim that structure is preserved across the change to relativity amounts to the claim (in this particular case) that the class of models identified under coordinate transformations in the relativistic case preserves the class of models to be identified (and this class is described precisely by the Lorentz transformations), save for the modifying term γ —and that the classes are strictly equivalent (isomorphic) as γ vanishes.

speed of light, on terrestrial scales of space and time is—for all practical purposes—infinite and so in the absence of sufficiently delicate measurements the Galilean transformation describes the world to within reasonable experimental error.

Second, and more relevant to the discussion here, is the mathematical/structural features equation (1) and (2) share. Undoubtedly (the structuralist argues), the reason that past scientists were so convinced of the correctness of equation (1) was that they *were in fact* describing the true relation between inertial reference frames, and it was only due to our contingent limitations as creatures adapted to the savannahs where distances were small and times were large compared to c that we didn't write down the complete equation but only an approximation. Having written down the Galilean transformations we were getting the world *right* save for the missing term, a term which vanishes in anthropic domains, and one which we have since filled in after overcoming our primate limitations.

This stands in significant contrast to the picture of this theoretical change according to the traditional realist. To successfully *interpret* each of the theories in which equations (1) and (2) are embedded requires distinct physical ontologies. In the Newtonian case, absolute space and time are posited, which in turn suggests a picture of space as a substantial *thing*. In relativity, not only is absolute space not required, it is ontologically idle!²

The connection between the problem of ontological discontinuity for standard scientific realism, and the naturalized metaphysician's project to remove *a priori* distortions from metaphysics should be unsurprising. The pessimistic induction has teeth as long as the purported referring terms of successive theories are extensionally inequivalent. Structuralism sees this difficulty arising from the very attempt to interpret the central terms of a theory *as* referring. For the structuralist it is this act of interpretation itself which introduces the distorting effects of our *a priori* prejudices about what the world must be like. That ontological discontinuity occurs is simply further evidence that allowing our *a priori* intuitions to enter into our understanding of reality leads to an unconstrained, confused, and ultimately false picture.

²This idleness of an absolute space is manifested in the consequence that if such a preferred frame existed it would be in-principle undetectable.

3.1.2 STRUCTURE AS METAPHYSICS

How, then, is structuralism supposed to help remove the *a priori* intuitions from scientific interpretation, and thus from metaphysical inference? The key is in the structuralist account of *interpretation*. Structuralism was motivated by the recognition that attempts to interpret scientific theories realistically inevitably presuppose some metaphysical principles which have snuck in from our pre-theoretic folk concepts about the world of “medium-sized dry goods.” The quandaries of quantum mechanics offer a nice example. Attempts to interpret the theory ‘realistically’³ famously appear paradoxical and somehow unintelligible. Structuralism suggests this is because typical attempts to interpret the phenomena presuppose a model whose domain consists of ‘objects-bearing-properties-in-space-and-time’ that would be recognizable to Aristotle or even *Homo neanderthalensis*. Naturalism points out that such a presupposition is by no means essential to the theory, and in fact we have every scientific reason to doubt this presupposition!

The structuralist insight is that the *structure* of a theory is often perfectly coherent and “well-behaved,” even if there are serious difficulties *picturing*⁴ the phenomena described. In the case of quantum mechanics the mathematical structure of the theory is consistent and well-understood. The structuralist suggests we simply refrain from attempting *any* interpretation (thereby avoiding its distorting effects), and instead infer from the adequacy and coherence of the mathematical structure to its positive ontological status. Instead of attempting to grasp the idea of a physical particle passing through both slits, or the idea of a physical particle which neither has a particular property nor lacks it, nor both has and lacks the property, we should understand that what exists is the structural relations into which we’ve been trying to fit our *a priori* notions of what the world must be like.

In order to offer a convincing account of an ontology which is fundamentally structural/relational, ontological structuralism provides two arguments: first, a negative argument against the exis-


³As used in this particular context, ‘realism’ is something of a term of art. It will suffice here to understand this as meaning an interpretation which construes the theory as describing concrete objects whose spatial arrangements continuously evolve in time.

⁴As in, having an image in our heads of entities (objects or events), bearing properties and relations, evolving in space and time, etc.

tence of *individuals*⁵ in our fundamental physical theories. Second, they offer a positive argument as to why we're justified in asserting the reality of the structure which remains and what such an assertion means.

AGAINST INDIVIDUALS

The strategy of the first argument is to show that some natural conditions which we would expect to hold for individuals at the fundamental level are not, in fact, borne out in our best theories. One example of this argument focuses on the surprising results of quantum statistics, a feature known as 'permutation invariance'. Consider the statistics of permutations on particles⁶ if our model of reality consists of classical microscopic individual objects. Construct a system of two (indistinguishable) particles and two boxes.⁷ Classically, and intuitively, the system can be in one of four states, that is, there are four unique ways the particles can be distributed in the boxes (here labeled A and B):



States (1) and (2) are distinguishable, since we will posit that the boxes themselves are labeled. States (3) and (4) however, are *not* distinguishable—by hypothesis the two particles are identical in all respects, save for their spatiotemporal location. By looking in the boxes one could not tell if the particles had been switched. Nevertheless, on this classical picture, there is still a matter of

⁵An 'individual' in this context is an ontological simple. This could be, depending on the account, atoms of substance with no proper parts, or space-time points instantiating natural properties, etc. The common notion in object ontologies is the idea that these simples are uniquely identifiable and distinguishable (if only up to spatiotemporal coordinates). It is this notion of individuals which is the target for structuralism's central critique since its centrality to object ontologies makes it the lynchpin in attacks on such views.

⁶The use of particle should be understood as agnostic between the classical notion of an individual object, and the structuralist idea of *relata* parasitic on the underlying relation. The term is unfortunate since it clearly elicits the former idea—but this is just further evidence that our intuitive, folk metaphysics pervades even the deepest parts of the physical sciences.

⁷This example is adapted from Ladyman and Ross, 2007, p. 133. All their caveats about the simplifying assumptions and unrealistic nature of this example apply here also (see their footnote 6, p.133).


fact about the locations of each particle. This is indicated above by representing the particles as black and white. These properties could be thought of as intrinsically hidden identity properties, or ‘haecceities’ (the essential ‘thisness’ of each individual object).

Suppose we contrive a mechanism by which we can randomly generate these states (perhaps by using a ‘Plinko’ board). If our process is genuinely random then we would naturally assign each state outcome the probability $1/4$. While a full half of the possible outcome states would be indistinguishable, we can have indirect evidence of their distribution. All we need to do is look at the statistical outcomes of this experiment. Assuming each state is equiprobable we can sample the relative frequencies of states (1) and (2). We would find that each has a relative frequency of $1/4$. Thus, on the basis of the probability calculus (and observation) we can conclude that states (3) and (4) jointly occur the other $1/2$ times. While we cannot know their individual frequencies, on our supposition that all states are equiprobable we should assign them each $1/4$.

How does this situation change when we look at the same system but with quantum particles? We’ll start this time with the measured frequencies and work our way back to the physical picture. Again, we posit that the particles themselves are indistinguishable, and we arrange the system so that all states are equally likely. Run the experiment a number of times and then examine the relative frequencies of the two states that are in principle distinguishable (in this case (α) and (β) which correspond to states (1) and (2) in the first set up). What we’ll find is that the relative frequencies⁸ for states (α) and (β) are each $1/3$! What is going on here? That is, how can we describe the behavior of the states of the system? Without rejecting our hypothesis that each state is equally likely there is exactly one consistent description of the system: There are three (not four) possible states of the system.⁹ Call this third state (γ) . We say that the system is *invariant* under permutations of the constitutive particles, that is, the permutation (swapping) of particles in (γ) does not result in a new physical state. The picture looks like this:

⁸The particles described here are behaving in accordance with the Bose-Einstein statistics.

⁹Since the total probability space must sum to 1, two states measured with probability $1/3$ entail exactly one further state such that all outcomes are equal.



In the classical system states (3) and (4) represented two ontologically distinct though empirically indistinguishable outcomes, and this was supported by measurement of states (1) and (2). If we are to accept that measurement of the two distinguishable states (in this case (α) and (β)) is evidence for the ontological fact of the matter in the empirically opaque circumstances of the classical example, this should hold for the quantum system as well. It follows that there is no corresponding distinction between a state in which one particle is in A and the other in B and the reversed permutation. There is exactly one way in which the particles can separately occupy the boxes, and this arrangement occurs $1/3$ of the time.

In the quantum case it appears as if ‘particle’ is behaving as a mass noun, rather than a count noun as we might naïvely expect. Instead of randomly assigning particles to the boxes, imagine a certain quantity of a fluid (say, water) is poured into the boxes. The stream of water is deflected by means of a paddle which randomly directs the stream to A, B, or doesn’t direct the water, but splits the stream evenly into both boxes. In this case, there are exactly three possible outcome states, corresponding to (α) , (β) , and (γ) . Prima facie this fact itself seems to militate against the view that particles in our physical theories are individuals, since there is no sense in which referents of mass nouns possess individuating identity conditions above and beyond their constitutive properties.¹⁰

To recap the argument: If particles are genuine individuals then there should be some individuating characteristic which each uniquely possesses independent of the presence and state of any other particles. This ‘individuater’ makes it the case that there is a metaphysical fact of the matter in distinguishing two states in which the particles are interchanged. But this has measurable statistical effects. These effects are not borne out in actual measurements. So, particles are *not* genuine individuals.

¹⁰A fishbowl full of water is not treated as a collection of separable, discrete, individuals. Compare with a fishbowl filled with stones.

But if particles are not individuals, what are they? If the negative portion of the structuralist argument is successful, they must offer a coherent replacement ontology, else the very notion of science offering a metaphysical picture of the world seems to evaporate. This sets up the dialectic nicely, since now the structuralist is in a position to offer a structure ontology as such a successor.

THE FUNDAMENTALITY OF STRUCTURE

The argument that structure *exists* rests on the notion of an ‘invariant’ under a certain set of transformations, and how this plays a role in physical theories. Schematically, the argument is that invariants of a transformation are objective, and thus real.¹¹ This is in distinction to things which change under transformations and so are ‘illusory’. This principle is then applied to our best science, in which transformations and invariants play a large role, thereby offering a criteria for existence. Finally it is argued that what is invariant in these cases is structure—the mathematical/nomological relations that hold in the theory.

Invariants, very generally, are those things in a system which remain the same when then system undergoes some sort of change. For example, the shape of my water bottle is invariant under rotations through the z-axis, though rotations through either the x- or y-axes do not preserve shape. We speak of such invariance as a symmetry of the object. Likewise, the shape of my coffee mug is not rotationally invariant through any axis, though it is translationally invariant—its shape does not change if I change its location in space. In the case of my water bottle, someone observing it from a position rotated from my vantage by an arbitrary angle (about the z-axis) will describe the water bottle exactly the same way, modulo the angular difference. What this means is that if the person were rotated into my position they wouldn’t have to change their description. This is not generally the case for the mug. Suppose I’m viewing the mug such that its handle is protruding straight out to my right. Consider the description of someone who is viewing the mug from my left by a rotation about the mug of 90 degrees. Rotating them into my position, their description of the mug would no longer match their view of the mug itself.

¹¹I should note that this is, of course, all too quick. I give a fuller gloss of the structuralist approach below (page 83).

At first these examples might appear as counterexamples to the notion that only invariants of an object (or system) are objective and thus real. Surely the shape of my mug, and the mug itself, is objective! We shouldn't take these perceptual examples too seriously, and analogies can only be pushed so far. The question (in this case) is not whether the object exists or has an objective status, but rather about the completeness of the description of the object—how we represent it. A description of the mug based only on a perspective from one angle around the z-axis doesn't capture truly *general* features of the object, since each description will mention a little bit of handle, or a lot of handle, or no handle at all. Russell makes a similar point in *The Problems of Philosophy* (1912) when discussing our perception of a coin, or a desk. In our phenomenological description of the coin its shape is a specific flattened oval, whose characteristics vary over a number of transformations. Likewise, the desk phenomenologically described is, from one perspective a trapezoid, from another a parallelogram, from a third a complex irregular shape. In these cases where the descriptions are not invariant, it is not that the desk is not real, but that any one of the descriptions does not characterize an objective feature of the desk.¹²

Notice, however, that in the above discussion, the idea that what is invariant is by that token real is at best assumed, with no argument. What is the argument? It goes like so: Consider an object. Suppose that two different people (or the same at different times) describe the object from different vantages. For example, one close, the other distant. Or one above it, the other below, or each rotated about the object in different ways. Now, we compare their descriptions. Consider the things about which their descriptions disagree. Perhaps color. They can't both be correct, and they might both be mistaken, so we shouldn't put too much credence in the idea that the object has, as an essential characteristic, a determinate color. But now consider the things they do agree upon—whatever property reappears across a range of different perspectives, whatever doesn't change no matter how you look at a thing, those properties seem to have a better claim on being objectively *in* the object itself (so to speak). Of course, with only two observers

¹²One is reminded of the parable of the blind men and the elephant—each concludes things about the nature of elephants on the basis of their particular contact with a part of the elephant. These conclusions, of course, are distortions of the correct characterization.

there is a risk of accidental agreement, some noise obscuring the signal. So, take the limit of continuous transformations for infinite observers—this idealized examination should hopefully bring out exactly those things which are objective about the entity.¹³ Finally, if such a property is invariant across every perspective and transformation, then it doesn't seem like it can be an artifact of the observers themselves—it must be real.

The above argument is at best idealized in the case of everyday objects.¹⁴ The examples lean heavily on the notions of perception and description, both of which are complicated, messy, and ill-defined at the level of the manifest image. The case, it is hoped, gets stronger when one starts looking at characterizing the entities in theoretical physics.

By applying this notion (identifying existents with invariants) to physical theories, we can both clarify and make more persuasive this argument, but also argue that it is, in fact, structure which is invariant under transformations of the theories. To illustrate this we can look at the simple example of gravity.

Consider a gravitational field. This field can be characterized as a potential field, assigning to each point in space a number, which is the value of the field at that point. Massive particles accelerate in this field by following the (negative) gradient of the field, which describes how the value of the field changes throughout space. Since these accelerations on particles can be described in terms of the force on a particle, we can understand the gradient of the field as describing the (gravitational) force a particle feels.

Now consider an object which moves from point p_0 to p_1 . Can we determine the value of the gravitational potential at either or both of those points? Acceleration is an observable, and if we know the mass of the particle we can write down a function describing the forces on the particle from p_0 to p_1 . This function can be integrated over the interval p_0 to p_1 to find the value of the

¹³Perhaps I'm overstating a simple point—the very notion of objectivity is that anyone, from any perspective, will agree on the property in question.

¹⁴However, even with the commonsense objects of perception we can get a good sense of how *relations* and *structure* come to play a natural role in descriptions which have the desired property of invariance. Presented with the problem of offering up a description of the mug which is invariant under the relevant transformations, we might ambitiously set out to list the position of every atom in the mug relative to every other atom (in the mug). This would in effect be a description of the internal structure of the mug, which would be invariant under any transformation.

potential between the start and end points. However, there is a remaining constant term left over from the integration—what is the physical meaning of this constant? Suppose after evaluating the definite integral we have the following expression:

$$U(p_0 - p_1) = (10 + C) - (2 + C). \quad (3.4)$$

The constant of integration will cancel out, resulting in a relative potential difference of 8. But what is the value of the field at each point?

We can think of this case in terms of two observers, Hyde and Lois. Hyde finds himself near p_0 . Most of his life has been spent here, and he's used to thinking about the behaviors of objects relative to p_0 . So it's natural for him to assign the field a potential of zero at the point p_0 .

$$U_H(p_0) = 0. \quad (3.5)$$

Lois, however, lives near p_1 , and like Hyde, has grown accustomed to thinking of objects relative to her surroundings. So she assigns to the field a potential of zero at p_1

$$U_L(p_1) = 0. \quad (3.6)$$

Hyde and Lois both observe an object moving in the field, traveling from p_0 to p_1 . Hyde describes this as the object moving from a potential of 0 to a potential of -8 . Lois alternatively describes the motion in terms of the object falling from a potential of 8 down to a zero potential. They fundamentally disagree on the values of the field at any given point. However, *they will agree on all observable facts about the falling object.*¹⁵

¹⁵As an example, they will agree on the velocity of the object at the point p_1 . (We will assume they are at rest with respect to one another.) To see this consider how they each describe the total energy of the system at an initial time and at a final time:

$$E_{initial} = E_{final} = mgh + \frac{1}{2}mv^2. \quad (3.7)$$

Let $m = 1$ and $g = 1$. Hyde will describe the energy of the system initially as: $E_h = 0 + 0 = 0$. Lois has: $E_l = 8 + 0 = 8$. For the final state of the system we have Hyde: $0 = -8 + \frac{1}{2}v^2$, so $v_f = 4$. For Lois: $8 = 0 + \frac{1}{2}v^2$, so $v_f = 4$.

We can compound this by considering continuously many observers each assigning to p_0 a real number. Each would be able to derive exactly the same behavior as either Hyde or Lois. We can say that the value of the potential at any given point in the field is *not invariant* under a certain class of transformations. Were we to insist that such points do in fact possess some value or other, then it is in-principle unobservable and undetectable. Assigning a value to a point in the potential field is a gauge freedom of the theory. However, the relation *between* points in the potential field is invariant—every observer will agree on the relative difference in the field between points.

The values of the potential, it appears, are parasitic on the relations between those potentials—they do not have fixed and determinate values, but rather are defined in terms of the relations in which they participate. This puts pressure on views that see spacetime points to be individuals which possess determinate properties. The alternative, as the structuralist recommends, is to take the pattern of relations which fixes the structure of the field as real. Returning to the example of permutation invariance it turns out, similarly, that we can better understand why there is only one state (γ) in terms of relations. If we take the relation as the fundamental entity then there is no mystery why the statistics work out the way they do: there exists only one state (γ) rather than two simply because there is exactly one (unique) undirected graph connecting two vertices.

3.1.3 MATHEMATICAL FIXED-POINTS

Structuralism must satisfy the original naturalistic desideratum: offer naturalistically acceptable epistemological fixed points from which to ground metaphysical inference—epistemic foundations which are not contingent on our particular evolutionary legacy and are secure from the distortions of our conceptual and material ancestry. The structuralist's focus on mathematical structure accomplishes this goal in two ways. First, as seen above, it offers a way to positively avoid the *a priori* distortions which are part and parcel with attempts to 'interpret' a theory by constructing a model in a domain of concrete, property-bearing, individuals. By focusing on the mathematical structure of the theory, there's no temptation to ask what the math is 'about'. Sec-

ond, it offers naturalistically acceptable fixed-points in the (*a priori*) truths of mathematics and logic.¹⁶

This account of sophisticating naturalized metaphysics doesn't actually eschew *all a priori* knowledge and justification. Instead it identifies a privileged subset of the *a priori* which meets the naturalist criterion. It develops a more fine-grained account of the sources and justification for *a priori* knowledge, and holds that the truths of math and logic pass the naturalistic test.¹⁷ This approach thus holds that there are, in fact, fixed nodes deep at the core of our web-of-belief upon which we can build a naturalistic metaphysics.

It is important to recognize that this approach amounts to a de facto repudiation of Quine's proclamation: "no statement is immune to revision [even] the logical law[s]..." (Quine, 1953, p. 43) Structuralism responds to the Quinean Realization by simply denying it—it is *not* the case that the possibility space of conceptual schemes is infinitely plastic and mutable. Any constellation of belief, no matter its contingent history, will share a common, necessary core of mathematical truths, and thus metaphysical conclusions which issue from this core are naturalistically warranted.

Before moving on we should dwell momentarily on the lynchpin of the structuralist account of naturalized metaphysics, and its sophistication over naïve naturalism. The fundamental complaint with analytic metaphysics, and what ultimately sank naïve naturalism, was their treatment (explicit or not) of *a priori* intuitions and principles as epistemically licensing metaphysical conclusions. Trust in such intuitions—the complaint went—was neither motivated nor encouraged by what we know (through the sciences) about the world and ourselves. Structuralism avoids the complaint, not by removing all reliance on propositions which appear *a priori* true, but by a

¹⁶On this view math is a good *a priori* fixed-point because of its good deductive behavior—it lends itself to theorems etc. This rigor is what imbues it with legitimacy, or so the structuralist claims, but see Mark Wilson's comment below.

¹⁷Notice that this could only be the case if the truths of mathematics and logic satisfied the constraints of epistemological naturalism. To do so appears to require that their status is independent of our psychological/cognitive evolutionary legacy. That is, that our mathematical knowledge is independent of our being the sorts of creatures we are. We will return to this below.

What this entails is that the statements of math and logic are necessary—independent of whatever contingent conceptual scheme a possible being might have. Among the various theories of mathematical truth, it would seem that Platonism about mathematics appears to be the one most clearly accounts for this feature. While not problematic in itself, the merging of naturalism and Platonism is at least surprising. Advocates of structuralism appear to be aware of this: "One distinct, and very interesting, possibility is that ... the traditional gulf between Platonistic realism about mathematics and naturalistic realism about physics will shrink or even vanish." (Ladyman and Ross, 2007, p. 237)

principled distinction between *a priori* intuitions which are *not* naturalistically justified (perhaps, for example, the intuition that no object can be colored red and green at the same time), and those that *are* so justified (for example, the intuition that if it is true that A&B, then it is true that A).¹⁸ The burden for justifying a sophisticated naturalized metaphysics now rests on this distinction, and whether such a distinction can indeed adhere to naturalistic scruples. We must now ask, is this distinction borne out in taking science seriously?

3.2 Taking Structuralism Seriously

Structuralism as a naturalistic metaphysics depends for its coherency on being able to show that the metaphysical conclusions it arrives at pay due attention to science. But due attention to science does not only involve understanding the formal apparatus of a theory, but also attention to how scientific knowledge is generated, and what the scientific attitude, and scientific practices, teach us about inquiry.

Structuralism specifically claims that the world as disclosed by science is a structure, where this means mathematical structure—i.e. the relations that hold between entities are (ontologically) fundamental. But for mathematical structure to be the ontological conclusion of naturalistic inquiry, it has to be explained what warrants it. It is the omnipresence of mathematics in scientific theorizing coupled with the seeming fundamentality and unshakeableness of mathematical truth and reasoning which makes structuralism look so attractive.

We should ask—is structuralism indeed as successful in packaging a metaphysics for science as it might appear? And perhaps more importantly, does the view proceed by way of a spirit allied with the general practices and attitudes of the scientific process? That is, does it take science seriously, warts and all?

There are, I contend, two reasons why structuralism fails in this bid. Both focus on the nature of mathematical representation in the sciences, and what taking science seriously will lead us

¹⁸These two (ostensibly necessary) propositions are to be distinguished in that the latter is a *logical* truth, whereas the former is not. Drawing the distinction in this way is merely for an example—there may be other ways to divide up the naturalistically acceptable and unacceptable propositions.

to conclude about the role of mathematics. Not only does structuralism fail to account for the diverse and subtle ways in which mathematics is employed non-representationally in practice, but it assumes features of our cognitive relation to mathematics which are entirely unmotivated by attention to science, and indeed presupposes a robust *a priori* metaphysical commitment to the nature of mathematics, one that is at odds with at least one scientific account of our cognitive relation to mathematics. Nowhere, I claim, is this *a priori* commitment borne out by attention to science, and only comes from the lurking non-naturalistic commitment to naturalistic metaphysics itself.

In the next chapter I suggest that proper attention to science—taking it seriously—results in a patently deflationary attitude about ‘realism’ with respect to the world, and we ask—what does this modified naturalism recommend with respect to metaphysics? Science teaches us to suspect our *a priori* intuitions, and our sense of how things are or should be. But it also teaches us that those intuitions are inescapable and that it is folly to seek a view from nowhere. Instead, the lesson to take is that the questions we put to the world will return answers which depend themselves on the very context and nature of those questions. A naturalistic investigation into metaphysics—into the nature of the world—is equal parts careful attention to the products of our empirical queries, and careful deliberation and reflection on which questions are the right ones for us? Which questions are worth pursuing? Which concepts are best to articulate those questions? Which concepts should we use? I develop a positive account of just this sort of metaphysical project, and offer up a concrete example of how it may advance our metaphysical and scientific project. In the remainder of this chapter we turn to look at the failure of structuralism as a naturalistic metaphysical outlook.

3.2.1 TWO NATURALISTIC FAILURES

MATH AND METAPHOR

If structuralism is to be naturalistically acceptable it must be the case that mathematical concepts are *independent* of the contingent conceptual capacities which are accidental to us as hu-

mans. If instead we found that mathematical concepts *were* so dependent, then there is nothing to distinguish—in principle—between the *a priori* intuitions of the analytic metaphysician and those of the naturalist-structuralist. That is, naturalized epistemology would then have no principled way to identify those intuitions which are corrupted by our anthropocentric situation, and those which are independent of our contingent biology.¹⁹ Structuralists presuppose that mathematical concepts are independent in this manner—that they float free from the environmental and ancestral moorings of our wider conceptual-linguistic capacities. If we were to meet intelligent aliens—the structuralist invites us to imagine—our concepts relating to color, texture, solidity, mass, etc. (that is, all our concepts related to, or derivative of our sensory modalities) might not have any commonality with the concepts of the aliens. But our mathematical concepts would be translatable—the expectation is that since no sensory modalities are required to engage with and manipulate mathematical concepts, and since the only thing required is a ‘rational’ capacity, the aliens would also be able to recognize these concepts and our representations of them, providing a common conceptual ground for communication.

This idea is not new, and indeed has been put into practice in some attempts to communicate with extraterrestrials. In 1974 Carl Sagan and Frank Drake designed and transmitted an FM radio signal 1679 digits long from the Arecibo radio telescope. Why this number? They relied on whatever received the signal to recognize that 1679 has the unique prime factorization of 23 and 73, and thereby assemble the signal into a 23 by 73 array of ‘picture elements’ (i.e. pixels), the pixels colored by the binary pattern.²⁰ Additionally, the authors of the Arecibo signal authored a *Scientific American* article some 20 years later: “How could we be sure that a particular radio signal was deliberately sent by an intelligent being? It is easy to design a message that is unambiguously artificial. The first 30 prime numbers, for example, would be difficult to ascribe to some natural astrophysical phenomenon.” (F. Drake and C. Sagan, 1997)

Much older ideas about such communication also leveraged mathematics—Carl Friedrich Gauss

¹⁹Or, rather, there would be no reason to suppose there was such a thing as the latter.

²⁰For more information on the content of the message and its hoped for interpretation see: Brau, 2017 and links therein.

is credited (perhaps erroneously, see citation) with the idea of carving a geometric demonstration of the Pythagorean theorem into the Siberian forest so as to be visible to inhabitants of the Moon or even Mars, the thought being that such a display would signal to the aliens the presence of intelligent life. (Wikipedia(b), 2017) Indeed, there are a variety of suggestions for formal, quasi-formal, and mathematical languages through which we could signal our presence and communicate with extraterrestrials. (Wikipedia(a), 2017) These are all motivated by the idea that mathematics is the *lingua franca* of the cosmos. Notice, for example, in the Arecibo signal all the mathematical expectations which are implicit. First, we must expect that the binary structure of ‘0’ and ‘1’ will be recovered from a signal at 2380 MHz frequency modulated by 10 Hz.²¹ Second, the signal must be recognized as a *number* of bits. Third, this number must be recognized as semi-prime (i.e. uniquely factored into exactly two primes), and the fundamental theorem of arithmetic must be known and employed to recover the two prime factors. Fourth, these numbers must be conceived of as cardinal numbers (as contrasted with ordinals) for the interpretation of the numbers as representing the rows and columns of a 2-dimensional array.²²

Let’s be explicit about what structuralism requires: Structuralism consists in the claim that what exists is structure, or that mathematical structure correctly represents reality. This claim is motivated by noticing that the alternative ontological account—what exists are fundamental ‘individuals’, objects which compose the world according to their intrinsic properties—relies unjustifiably on concepts which trace their origins to the contingent vagaries of our time and place of evolution. These concepts are a result of our parochial ancestry, and so cannot be trusted to represent the world faithfully. Implicit in the move from objects to structure is the claim that mathematical concepts are *not* so contingent upon our evolutionary history as human animals. But this is a strong claim, and should, according to the naturalist, be supported by attention to

²¹Which is on the order of a difference of 1 part per billion—for comparison typical FM radio broadcasts modulate their carrier wave by ~ 75 kHz, which is roughly 1 part per thousand.

²²Perhaps I am being unduly hard on Drake and Sagan. The Arecibo signal was more of a symbolic public relations opportunity than an earnest attempt to contact alien life. For one, the Arecibo antenna is not optimized to transmit on frequencies which are thought to best stand out against cosmic noise. Nevertheless, Drake, Sagan, and many others see the basic premise of transmitting in some fashion which signals our mathematical competencies as fundamentally sound.

what our best science indicates about the nature and origin of mathematical concepts. If we find there is *no* positive reason to think mathematical concepts have this special objective status, then structuralism's claim to be genuinely naturalistic is undermined.

That science does, or could ever, vindicate the claim that mathematical concepts are independent of our cognitive situatedness is doubtful. Attention to cognitive science suggests there is evidence that mathematical concepts *are not* independent in the way structuralism requires. George Lakoff and Rafael Núñez (2000) argue that our fundamental mathematical concepts rely upon conceptual metaphors linked to our embodied experience. This is an extension of the general project of cognitive linguistics which studies conceptual metaphor and how abstract and higher-order concepts are parasitic on our familiarity and conceptual competence with mundane experience. For example, the abstract relational concept **affection** is modeled in terms of warmth and coolness: "She *warmed* up to me." "He gave me an *icy* stare." (Lakoff and Núñez, 2000, p. 41) Lakoff and Núñez give a variety of further examples: **importance** in terms of physical size,²³ **similarity** in terms of physical proximity.²⁴ (ibid., p. 41)

Similar conceptual metaphors exist for mathematical concepts. For example, the concepts of categorical logic, sets, and basic arithmetic all rely on an 'image schema', specifically, the *Container schema*. As Lakoff and Núñez explain it:

The Container schema has three parts: an Interior, a Boundary, and an Exterior. This structure forms a gestalt, in the sense that the parts make no sense without the whole.

[...]

To get schemas for the concepts **In** and **Out**, more must be added to the Container schema. The concept **In** requires that the Interior of the Container Schema be "profiled"—that is, highlighted or activated in some way over the Exterior and Boundary. In addition, a figure/ground distinction must be added.

²³"This is a *big* problem." "Don't sweat the *small* stuff."

²⁴"These colors are very *close*." "Over the years, our opinions have *diverged*."

[...]

Image schemas have a special cognitive function: They are both *perceptual* and *conceptual* in nature. As such, they provide a bridge between language and reasoning on the one hand and vision of the other. Image schemas can fit visual perception, as when we see the milk as being *in* the glass. They can also be imposed on visual scenes, as when we see the bees swarming *in* the garden, where there is no physical container that the bees are in.

[C]omplex image schemas like **In** have built-in spatial “logics” by virtue of their image-schematic structures. (Lakoff and Núñez, 2000, pp. 30–31)

The Container schema provides a linkage between our visual capacities and abstract reasoning. Specifically, in terms of the spatial “logics” referred to above, the Container schema provides the framework for supporting inferences of the following type: “If x is in A , and A is in B , then x is in B .”²⁵ That is, if the milk is in the pitcher, and the pitcher is in the refrigerator, then the milk is in the refrigerator.

We can see how categorical logic can be constructed from the Container schema by producing a mapping between concepts in the ‘source domain’ (here, Container) and the ‘target domain’ (Category). See table 3.1.²⁶

Table 3.1: Categories are Containers

Containers		Categories
Bounded regions in space	→	Categories
Objects inside the bounded regions	→	Category members
One bounded region inside another	→	A subcategory of a larger category

In turn, this mapping supports the basic inferences of categorical logic. Listed below (Table 3.2) are the maps for *Excluded Middle* and *Modus Ponens*. As Lakoff and Núñez note, these

²⁵Note, this is, in general, *not* a valid inference in set theory. Set theory adds further structure to the Container schema.

²⁶Tables 3.1, 3.2, and figure 3.2.1 are adapted from (Lakoff and Núñez, 2000, pp. 43–45).

Table 3.2: Categorical Inferences

Containers		Categories
<i>Excluded Middle</i>		<i>Excluded Middle</i>
Every object y is either in Container A or out of Container A .	→	Every entity y is either in category A or out of category A .
<i>Modus Ponens</i>		<i>Modus Ponens</i>
Given two Containers A and B and an object x , if B is in A and x is in B , then x is in A .	→	Given two categories B and A and an entity x , if B is in A and x is in B , then x is in A .

mappings seem obvious, and perhaps trivial, on inspection. However, the crucial point is that the inferential structure of the abstract conceptual domain of categories is supported by basic, *embodied*, spatial structure.

Abstract concepts and our basic ability to evaluate and manipulate representations which employ them depend on these conceptual linkages. Venn diagrams (fig. 3.2.1)²⁷ are conceptually transparent because of the manifest spatial structure of the Container schema. Lakoff and Núñez argue that from the point of view of cognitive science “spatial logic is primary and the abstract logic of categories is secondarily derived from it via conceptual metaphor.” (Lakoff and Núñez, 2000, p. 45)




Figure 3.1: Spatial Logic of *Excluded Middle* and *Modus Ponens*

Lakoff and Núñez offer up four ‘grounding metaphors’ for basic arithmetic, following the above pattern of mapping patterns and operations from an embodied source domain to an abstract mathematical domain. (ibid., p. 50 ff.) The four metaphors are ‘Arithmetic is Object Collection’,

²⁷Technically, this is an Euler diagram.

‘Arithmetic is Object Construction’, ‘Numbers are Measuring Sticks’, and ‘Arithmetic is Motion Along a Path’. Each of these metaphors employs concrete physical domains with which we are practically acquainted to ground reasoning in the abstract domain of arithmetic.

While Lakoff and Núñez’s project is fascinating and ambitious, a full exploration is beyond the scope of this dissertation. The important feature for our concerns here is the explicit connection they make between embodied, physical, and concrete phenomena and our mathematical-cognitive framework. Abstract mathematical concepts are not plucked *sui generis* from the transcendental ether—they are metaphorical extensions of the manifest structures accessible to our physical bodies. If this conclusion is correct, then mathematical concepts are vulnerable to the very same Naturalized Epistemological concerns about the situatedness and contingency of our concepts to which the intuitions of the *a priori* metaphysicians are exposed.

It’s worth noting here what I am not claiming—I am not arguing that this account of mathematical cognition is correct and unassailable. However, insofar as Lakoff and Núñez’s project is a legitimate project within cognitive science and linguistics, we as naturalists have a *prima facie* duty to take their claims seriously.²⁸ The force of Lakoff and Núñez (2000) is to problematize an uncritical assumption of many naturalists: that mathematical concepts are wholly distinct from our ‘natural’ cognitive abilities. It also lays bare the uncomfortable position of the naturalist who wishes to preserve this distinction.

Recognizing the embodied and metaphorical aspect of mathematics throws the non-naturalistic character of the structuralist’s assumption into stark relief. Simply, there is no positive reason to think that mathematical concepts are exempt from the dislocating consequences of the Quinean Realization save for an *a priori* bias. The conclusion we should draw from this is that naturalized metaphysics—as exemplified in the ontic structuralist programme—fails in its naturalistic ambitions and succumbs to the very sort of *a priori* metaphysics is sought to displace. As we’ll see in the next chapter, the fact that mathematics appears to be *necessary* has a much to do with our *choice* (explicit or not) of concepts, as it does with the requirements of an external reality. This

²⁸NB: This, emphatically, *does not* mean simply believing their claims.

is just to say, alien life will share our mathematics, or diverge widely from it, in part based on what we will countenance *as* 'life' at all. We may decide that mathematics which matches ours is a precondition for counting something as (intelligent) life, but that reveals a conceptual choice on our part, not one of nature's joints. Lakoff and Núñez agree:

It is no accident that our branches of mathematics are linked in the way they are. Those conceptual connections ... express ideas that matter to us. The way the branches of mathematics are interrelated is a consequence of what is important to us in our everyday lives and how we conceptualize those concerns." (Lakoff and Núñez, 2000, p. 451)

TOOL OR TRUTH: MATHEMATICS IN PRACTICE

But it is not simply that ontic structuralism fails to take what science *says* about mathematics seriously. They also mischaracterize a significant portion of what scientists *do* with respect to mathematical practice. In order to maintain that attention to science tells us that *structure* is real, the structuralist requires that the central and primitive role of mathematics in scientific practice is *representation*. Attention to scientific practice instead suggests that mathematics is deployed in a wide and heterogenous array of tasks and roles, and it is only an *a priori* bias which singles out mathematics' representational role as special or privileged.

In thinking about the role that mathematics plays in scientific practice, let us begin with a mundane analogy, the practice of a woodworker. There are many tools available to the woodworker in her shop. Some tools are general purpose, designed to accomplish a variety of tasks, while others are more narrowly specialized. Drawing up plans for building a particular piece of furniture (say, a chair) is a crucial aspect of the woodworker's craft. Her skill is drafting plans for the chair ahead of time will give her a greater degree of control and precision as she proceeds with the build. The tools she uses (pencil, paper, straightedge, compass) might be thought—by an observer—to be specialized for the activity of drafting. But the observer would quickly become disabused of this supposition—these tools are employed broadly, in different ways in different

contexts. Indeed, some of these contexts are arguably more fundamental to the woodworker's craft than their deployment in drafting. For example, the pencil is used to mark wood-stock for cutting and for recording measurements. Paper might be used to check the tightness of a joint, or put down on a work surface to protect it from paint or glue. Straightedges serve to scrape excess glue from a joint. And one of the most valuable functions of the compass is to invariantly transfer length measurements.

What is the point of this vignette? Just as it would be a mistake to conclude that the pencil, paper, straightedge, and compass—because of their visible and important role in the drafting process—are fundamentally tools *for* drafting (to the exclusion of other uses), it is a mistake to see how mathematical tools are most visibly deployed (i.e. representing) and conclude that mathematical tools are *for* representing. To be sure, using mathematics to model physical phenomena is undoubtedly a valuable and powerful use of mathematics. But the historical attention of the theoretical side of scientific practice has instilled in many philosophers a lopsided view of the nature of the contributions which mathematics make.

The non-representational roles which mathematics play can be subtle, first for the reason that they are typically not highlighted as *the* things mathematics is supposed to be doing for a scientific theory. But also for the reason that there is a fine distinction to be made between an equation (for example) standing in a representing role, and it serving some other purpose. The examples are varied and sundry: physical measurements, both spatially and temporally require mathematical tools, as does the collection of data and the construction, adjustment, and manipulation of instruments. More abstractly, a mathematized theory which the theorist understands as representing the dynamics of a particular system, can be taken by the experimentalist and used to calculate (for example) the diameter of the aperture in their detector, or the voltage needed for the detector to function. The numbers, variables, algebraic operations, etc. in these cases have as much to do with representing their purported subject-matter as using a pencil to mark a cut-point on a board has to do with drawing plans for the finished chair—related, sure, but different.

The other way that we can see mathematics not on a representational footing is in the limi-

tations of mathematical methods themselves. It is often unproblematically accepted that mathematical equations are transparent upon inspection and that any difficulty occurs only in terms of computational power. This is, however, far from the reality of applying mathematics in attempts to (for example) model phenomena. Mark Wilson (2013, p. 166 ff.) gives a nice example of how this can arise in seemingly simple cases such as a bead moving along a wire. Given a differential equation which describes the bead's velocity as a function of its position, we can ask: what is the curve the bead traces through space-time? Ostensibly this is *represented* by the differential equation itself, yet our methods for extracting such a curve will (in all but a tiny class of such equations) be approximate. And such approximation is not the exception but the rule. 'Line-fitting' is a general case of this situation, where there is no tractable function which we might write down that fits the data as we have measured them. Either a function which *is* transparent is substituted, or in some cases a series expansion can approximate the complex behavior of the data, but only in a limited domain of inputs.

These cases would come as a surprise to anyone who was familiar with science only through the writings of much of the 20th century philosophy of science. Wilson diagnoses the problem like so:²⁹

[T]he most powerful inferential schemes utilized within applied mathematics often prove erratic in their performance: they sometimes work well and sometimes work badly, without displaying evident marks to distinguish the cases. Allied woes are rarely evident within the well-behaved inferential patterns of first order logic, *whose atypically cooperative behavior has bewitched many philosophers* into overlooking the computational logjams that substantially shaped the topography of working science. (Wilson, 2013, p. 165, emphasis added)

In many of these cases the mathematics doesn't serve a representational role because it *can't* serve such a role. Yet line-fitting, series expansions, Euler's method all play a part in the skillful practice of science. These methods, tricks, approximations are wielded as *instruments* for gaining

²⁹Recall the reasoning as to why mathematics was seen as a good fixed-point for the structuralist project (§3.1.3).

control over a complex and poorly understood domain. To deny these uses of mathematics play a role at least as fundamental as representation can only be sustained by begging the question against them from the outset.

3.3 Between Math and the World

3.3.1 PYTHAGOREANISM OR ‘REAL’ STRUCTURE?

At this point the structuralist might be expected to object to this line of reasoning. They point out that I have failed to adequately take into account the distinction between ‘mathematical’ and ‘physical’ structure. The criticisms I’ve leveled above focus on the non-naturalistic presuppositions which must get taken up on the way to an ontic structuralism. However, my criticisms take aim specifically at *mathematical* structure, and I presuppose that the only sort of structure the ontic structuralist could be speaking of was abstract and Platonistic in nature. But, the structuralist replies, they’re interested in *physical* structure—the structure of the perfectly real, concrete world we see before us. They can even allow that ontologizing mathematical structure *does* violate our naturalist scruples, but this simply fails to connect with their proposed metaphysical project. No doubt mathematical structure is central and crucial in *representing* the complex physical structures, but we should keep the distinction clear in our minds.

But while this appears to be a promising course to steer at first, on closer examination it looks like there is no non-question-begging distinction to be drawn which neatly divides the two notions of structure. This difficulty is what Steven French calls the ‘collapse problem’.

[I]f ... a ‘purely’ structural description [of the world is accepted], then it may become hard to discern any difference between the physical world and the mathematical world. Indeed, given the mathematization of science, and physics in particular, the structural description of the physical world may appear to be entirely mathematical... (French, 2014, p. 192)

The problem the structuralist faces, articulated here by French, is that once descriptions of a physical system are stripped of any intrinsic ‘objectual’ features, and we’re left with nothing but the structure, there seems to be nothing left that distinguishes the physical system from a mere mathematical structure.

For a simple example, consider an attempt to carefully describe the rotational behavior of a wooden block carved into the shape of an equilateral triangle. We find that there are three rotations which have a special status—any combination of these rotations will return the block to the same orientation as one of these three rotations. Additionally, one of these rotations doesn’t change orientation at all (rotation by 2π radians), and for any of these rotations, there is a rotation which will ‘undo’ the change the former made. This description of the rotational behavior of our block of wood can be formalized,³⁰ and any mathematician will recognize it as the cyclic group C_3 .³¹ Now, if we (somewhat implausibly) maintain that what is *real* about the block is the structural relations which the composition of rotations encapsulate, then *all we’re left with* is the group C_3 .

In the present case we can block this implausible conclusion by pointing out that our wooden block has wide range of other properties (both intrinsic and relational) which clearly make it something apart from the mathematical structure C_3 . But in the cases where structuralists are primarily concerned—the unobservable particles and forces of fundamental physics—there seems to be nothing analogous to a description of a particular knot in the grain, or the particular heft, color, or texture our block enjoys. At the level of fundamental physics there is no recourse to the sorts of everyday, commonsense, features of the physical world which distinguish it from a mathematical description. We can understand it when our physics teacher says that a ballistic projectile follows a parabola—but no one would mistake the rock for a parabola.

For the objection under consideration—which aims to push back on my critique that structuralism uncritically presupposes some special justificatory status for mathematics by distinguish-

³⁰Which I won’t do here, although I hope it’s clear how one could do so.

³¹In the description of the rotating block above, strictly speaking, I left out one requirement for its rotations to constitute a group, associativity: $a \circ (b \circ c) = (a \circ b) \circ c$. Interestingly, in physical instantiations of a group (like our block) associativity is trivial.

ing between physical and mathematical structure—the collapse problem poses a considerable challenge. One must find some feature of physical structure that earns it its status as physical, and yet it not just more individuals disguised.

Ladyman and Ross face these troubles then they offer up their distinction between “formal and material modes of discourse.” (Ladyman and Ross, 2007, p. 119) They insist that when they utter their ontic credo “What exists are (‘real’) ‘patterns’.^[32] ‘Real patterns’ should be understood in the material mode.” They contrast this with talk about ‘structures’, “understood as mathematical models ... that elicit thinking in the formal mode.” (ibid., p. 119) Yet *what* differentiates the real pattern from the structure is left unexplained—the authors trusting we’ll all know it when we see it. Later they appear to forget their own distinction, speculating about the diminishing gulf between Platonism and their brand of philosophy of science (see footnote 17). And there is good reason to embrace the elision, as it deals neatly with two difficulties which a structuralism which sought to sustain the distinction had trouble with.

First, the identification licenses the application of the concept of isomorphism between mathematical and physical structures in a strict (i.e. non-metaphorical) sense. One of primary pieces of evidence which structuralists point to in making their argument are the existence of isomorphisms between physical systems and mathematical structures. The mapping of the rotations of the wooden block to the cyclic group C_3 is one such example. This example also provides us with the obvious difficulty claiming there is an ‘isomorphism between the block and C_3 .’ An isomorphism, in its strict—mathematical—sense is a mapping between two mathematical structures, and so cannot hold between a mathematical structure and a physical object. If the structuralist wants to preserve a distinction between mathematical structure and physical systems, they face the choice of either settling for a notion of isomorphism which is merely metaphorical, or inventing their own notion of isomorphism which explains the relation.

The second advantage the strict identification between math and the world confers to the structuralist is that it avoids the suspect conceptual distinction between physical and mathemat-

³²See Ladyman and Ross, 2007, pp. 119, 228, 233.

ical structural entirely. It is, as evidenced by heroic efforts on the part of structuralists (see, French, 2014, pp. 192–230), difficult to make coherent even the notion of physical structure over and apart from mathematical structure, let alone offer criteria for distinguishing the two. Physical structure appears to stand as a metaphysical/ontological mystery—not objects, yet physical, not abstract (mathematical), yet purely relational. It seems to require the status of an unanalyzable ontological primitive. But this hardly seems like a move licensed by attention to the relevant scientific theories and practices. Scientists do not posit some primitive substance called ‘physical structure’ in order to carry on with their experimental and theoretical practices. Indeed, the structures which they deal in (at least on the theoretical side) are explicitly mathematical, and they are forthright in their employment of *mathematics* to help *model* physical systems. Taking physical and mathematical structure to be one-and-the-same seems, in this regard, to carry with it the fewest ‘unnatural’ attachments carried in from outside the sciences.

Even French, after devoting a chapter attempting to articulate the distinction between physical and mathematical structure, ends by saying:

Perhaps then we simply have to accept that the distinction between the mathematical and the physical has, at the very least, become blurred ... or that it cannot be drawn at all... Perhaps there is no answer to the ‘problem of collapse’. Perhaps we should follow Ladyman in dismissing this as a pseudo-problem... (ibid., p. 230)

The pseudo problem referred to here by French is what the collapse thesis denies—a distinction between ‘mathematical’ and ‘physical’ structure. Taking these statements at face-value it appears that the structuralist must take mathematical structure to be the focus of their ontologizing. Thus the structuralist can find no force to the objection distinguishing between physical and mathematical structure.

3.4 Conclusion

Ontic structuralism fails to take science seriously, and cannot coherently maintain their naturalist credentials. But it is too hasty to write off structuralism as an unqualified naturalistic failure—there are aspects of structuralism worth salvaging and lessons worth learning. A naturalistic metaphysics properly so called should:

1. Be scrupulous in the commitment to take scientific practice (in addition to theory) seriously.
2. Harbor a suspicion of metaphysical conclusions arrived at by appeal to *a priori* principles.
3. Recognize the concomitant (naturalistic) acknowledgement that *a priori* appeals are *unavoidable*, and this recognition should constrain and guide a naturalistic metaphysics.
4. Focus on articulating the concepts required for scientific theory—where articulating such concepts involves an openness to radically different ways of thinking of the world. Structuralism is to be commended for its vision of upending our normal conceptual schema.

A naturalistic metaphysics asks what the scientific attitude should be to questions of metaphysics.³³ The response is that science has shown us nature answers the questions we ask, but that our questions are fundamentally human-reflexive, and are diagnostic of what we take to be important. This in turn offers a view to the self-conscious activity of metaphysical theorizing. The next chapter advances a new approach to naturalized metaphysics, one which takes the lessons learned from taking science seriously, and modifies the enterprise of metaphysics itself.

³³NB: This is *not* equivalent to asking what a scientist or scientists think the attitude toward metaphysics should be or is. In this respect scientists are as unqualified to give an opinion on the basis of their status as any layperson. The *naturalistic* study of science takes a critical attitude towards the concepts of science itself, and does not invest scientists with any *prima facie* authority in this critical attitude.

Chapter 4

Pragmatism & Scientific Metaphysics

4.0 Introduction

This chapter is intended to make good on my claim that a *properly* naturalistic metaphysics requires a deeper revision of our conceptions about the practice of metaphysics than what the so-called naturalized metaphysicians have aimed at. They are, I argued, betraying the spirit of naturalism itself in seeking merely to replace the *contents* of traditional metaphysics with contents drawn from science. When Ladyman and Ross condemn neoscholastic metaphysics for ‘domesticating science’, what they offer in contrast is a project aimed at discovering the way the world *really* is by careful attention to our best scientific theories. But what the naturalist—who takes science seriously, warts and all—objects to in both these projects is the assumption that science in any way licenses the move from engaged, practical, empirical inquiry, to describing the world in terms of a fundamental, atemporal, intelligible *essence*. This move, from the practical, engaged, activity of science, to a static, eternal, ‘crystalized’ picture of the world is, from the point of view of the naturalist, just one more *a priori* supposition which we should suspect.

But the naturalist is neither a dogmatic quietist, nor positivist, nor instrumentalist regarding metaphysical questions. Metaphysical inquiry is *valuable*, both as a prelude to scientific inquiry, and as part of its conclusion.¹ But, this value has been mistaken—it is not epistemic in the clas-

¹Indeed, stating it in these terms only highlights the artificial division between metaphysical inquiry and scientific

sic sense we might expect. That is, metaphysics is not about producing “deep, worldly truths” (Thomasson, 2017a, p. 364) or “discovering especially deep or fundamental facts about the world.” (Thomasson, 2017b, p. 101) Rather, the naturalist endorses the idea that “metaphysics fundamentally involves normative conceptual work.” (Thomasson, 2017a, p. 364) Perhaps the importance and significance of this claim is not obvious at first glance, but as I will argue, this idea signals a break from traditional *a priori* metaphysics in a manner far more profound than the distinction to which naturalist metaphysicians have typically clung in the past.

In this chapter my aim is to articulate a picture of metaphysics done in a way consistent with the naturalist attitude of taking science seriously—one which *springs out of* such an attitude. My central criticism of the so-called naturalistic metaphysics of the naïve and sophisticated variety considered in Chapters 1 and 3, respectively, is that they allowed themselves an unserious, idealized picture of science, one that has little connection to what scientists *do*.

In the naïve case, science was idealized to consist of a pure empirical product, formulated in some ‘natural’ observation language, and standing independent from any interpretive, theoretic, or pre-theoretic contributions by the humans who discover, formulate, or employ such a language. Thus (on this view) ontology can be unambiguously and decisively extracted from theory.

In the sophisticated case, science is admitted to be linked to the interpretive, theoretic, and pre-theoretic contents of the scientists who wield it. However, a more subtle idealization is imposed, formulating science in terms of two separable parts—the first being that portion of science shot-through with anthropocentric, meso-scale, folk conceptual trappings, while the second is the part standing independent of these human fetters. It is this latter part which licenses extracting what ontology one can—with the qualification that such an extraction is no longer trivial, and is perhaps less far-reaching than in the naïve case.²

inquiry. There is just one thing, inquiry.

²To be absolutely clear, for the ontic structuralist view which represented the sophisticated approach in chapter 3, this second part is *mathematical structure*. This is not, however the only sophisticated view. An alternative which I do not address in this dissertation is the naturalized metaphysical view of Anjan Chakravartty, which finds this

What is the alternative? That is, beside a wholesale repudiation and abandonment of metaphysics as a domain of inquiry, what other options are left for the naturalist? There is at least one alternative, which I elucidate here. The central motivating factor in this account is—as the reader might expect—the admonition to take science seriously. The first task, then, is to spell out what it means to take science seriously. This requires attention to the embodied, concrete, human-scale and social-scale activities which comprise scientific *practice*. This ‘turn to practice’ has been a centerpiece of much of late 20th and early 21st century philosophy of science, but has been rather neglected in the specific area of naturalized metaphysics. Taking the turn to practice seriously shows how focused attention on the practical and embodied activities which comprise science should reorient our conception of (1) the products of scientific inquiry, and (2) the nature of scientific knowledge. This reorientation ramifies throughout the naturalized metaphysical project.

What sort of metaphysics does such a view license? Attention to the turn to practice suggests that the naturalist depart from traditional analytic metaphysics along a wholly different dimension than what naïve and sophisticated naturalism attempted. This departure distinguishes itself in terms of re-conceiving the practice of metaphysical inquiry. This re-conception can then be examined in two parts. I begin by looking to the deflationary approach of Amie Thomasson’s ‘easy ontology’, which looks to recast metaphysical debates in terms of meta-linguistic conceptual negotiation. But, this view does not go far enough, and so we’ll transition to the second part. Here, I will look to the American Pragmatist tradition to draw lessons about what role metaphysical inquiry and debate might play within the scientific context, as well as more broadly across society.

The view I develop here occupies an intersection between the notions of *concepts*, *science*, and *authority*. We can contrast this view (which one might call ‘naturalist-pragmatist’) with naïve and sophisticated naturalism in these terms: For the latter views the project of naturalized meta-

second part in the *empirical content* of a theory. (see, e.g. Chakravartty, 2013) Where the ontic structuralist errs by idealizing mathematics beyond what is naturalistically licensed, Chakravartty’s view errs by helping itself to the Experiential Given—the idea that (to put it crudely) experience is already carved up into propositionally structured pieces, ready-made for our cognition. (If one were adamant on parallelizing the failures of the structuralists’ and this view, we could recast the former’s failure in terms of helping itself to the *a priori* Given.)

physics consists of taking *science* to be the ultimate *authority* for evaluating the adequacy of our *concepts* in correctly matching (that is, mirroring) the world. The central doctrines of pragmatism highlight three myths active in this statement: First, pragmatism denies there is such a thing as ultimate authority. Second, pragmatism denies there is any sense to the idea that concepts correctly mirror some external, fixed, reality. And third, pragmatism denies there are general, universal criteria for evaluating concepts, and that conceptual evaluation is fundamentally a contextual, concrete, practice-based task.

A naturalist-pragmatist approach suggests an alternative understanding of the metaphysical project. It consists in tracing how the practice of *science* gives rise to new *concepts* (or revisions and rearrangements of old ones), and how these concepts then exert *authority*³ within our practical activities (but here especially science) to shape future practice, valuations of various products and aims, etc. But it is not merely a descriptive project—along with tracing this connection between science, concepts, and authority, metaphysics is the domain in which normative critique of the sciences and their conceptual store is possible, especially with respect to weighing other competing aims and values in the wider social and political contexts in which these practices occur.

Finally, I'll consider a concrete example. I'll do this primarily by focusing attention on a particular area of contemporary research, one that I see as potentially benefitting greatly from this approach. This is the research field of astrobiology: “the study of the origins, evolution, distribution, and future of life in the universe.” (NASA, 2017) Astrobiology is a nascent research area in its own right, and is a fascinating case study for the philosophy of science for a variety of reasons. Here, however, I will focus exclusively on one of astrobiology's foundational concepts, the concept **Life**. There is no consensus on what constitutes a living system. Indeed, as we'll see, many discussions which grapple with the definition of life (typically in the context of questions about the origins of life) are deeply confused. It is not hard to see that this confusion stems in part from the fact that these questions find themselves in close proximity to questions and debates

³This isn't quite right—of course concepts can't *themselves* exert authority. Rather they are seen *as* authoritative, in the sense of framing, grounding, making possible certain vocabularies, etc.

more typical of ontological discourse. I'll argue that the seeming intractability of this question is best understood as a problem about how our concepts are and should be applied. Life, I'll argue, is not a natural kind, and the decision to countenance something (some system, object, process, etc.) as living will be a decision we'll have to make (insofar as we see that such a concept is needed).

4.1 Practice, Concepts, and Pragmatism

4.1.1 THE TURN TO PRACTICE

We begin our account of naturalist-pragmatist metaphysics by looking to the late 20th Century movement in the philosophy of science called the 'turn to practice'. Under this banner renewed attention has been placed on the broader contexts of science, society, and the material conditions of how scientific institutions work to generate knowledge. This is in distinction to what might be called the 'logico-theoretical' view of science. The 'logico-theoretical' view holds that science is best understood, and best exemplified, by the formal theories which scientists produce, and that, fundamentally, genuine philosophical questions orbit this notion of *theory*.

The relevance of the turn to practice to the present project is primarily due to the fact that naturalistic metaphysics has largely been left behind by the wider philosophy of science community in this respect. Accounts of metaphysics drawn from, inspired by, or grounded in science are still firmly rooted in the logico-theoretical view. Even when their exponents invoke language which appears to endorse the view that science is constituted by practice⁴ the content of their metaphysical accounts are almost invariably in terms of what a theory represents, and how well it does so. This should not be surprising, since there's a considerable legacy within metaphysics and logic connecting *theory* and *ontology*.

⁴For example: "Thus science is, according to us, demarcated from non-science solely by institutional norms: requirements for rigorous peer review before claims may be deposited in 'serious' registers of scientific belief, requirements governing representational rigour with respect to both theoretical claims and accounts of observations and experiments, and so on." (Ladyman and Ross, 2007, p. 28)

In this respect the project of ontology has been—for the past half-century—connected with Quine’s dictum: “To be is to be the value of a bound variable.” While this statement is suitably nebulous to provide home for numerous understandings, on the received view the general thrust of the dictum is to connect ontology with first order quantificational logic. What exists, according to this standard picture, are just the things in whatever domain models the theory in question. Clearly, if we suppose that what science *essentially* consists in is theory, then it is a small step to submit such a theory to this sort of analysis. We take the sentences of the theory and find an interpretation of those sentences. The objects in the domain of such a model are then what—according to the theory—there is. This trivializes the extraction of ontology from a scientific theory, and what *interesting* work remains is in the process by which we decide which theory is best supported by the evidence, and in the long run, is true.⁵

However, if we shift away from the idea that theory is what is central to understanding science, this picture of metaphysics and the extraction of ontology from said theory no longer looks attractive, let alone coherent. Thus, the turn to practice is fresh ground for a scientific metaphysics to explore. Insofar as taking science seriously leads to the practice-turn, a naturalist-pragmatist metaphysics owes a picture of metaphysics which is likewise revised in light of the turn to practice. And what reason do we have to accept that the proper attitude of naturalism will lead us to the turn to practice? Doing otherwise would require that we impose artificial assumptions and prescriptions about what parts of science we should attend to, and which we should ignore. The naturalist is committed to taking seriously the full gamut of scientific activity, and so takes seriously its practical aspects.

The turn to practice reorients the study of science—emphasizing (to put it rather crudely) what

⁵Ontic Structuralism changes this formula, but only by degrees—the idea (as discussed in the previous chapter) is to take the *structure* of the model (in a simple first-order example this would translate to the set of relations defined over the domain) as fundamental, and the objects of the domain as unreal, or only derivative of the structure. Note though, nothing about the structuralist account here questions the assumption that the locus of analysis is *theory*. Ladyman and Ross offer modifications of Quine’s dictum: in their (2007) book they proclaim that “To be is to be a real pattern.” (Ladyman and Ross, 2007, p. 226) where a real pattern is to be cashed out in terms stable structures. In their (2013) article they amend this, and offer a new thesis: “the world is the totality of non-redundant statistics (not of things).” (Ladyman and Ross, 2013, p. 111) Again, here the ontology of non-redundant statistics is to be drawn from analysis of theory.

scientists *do* rather than what they *say*.⁶ As Soler et al. (2014a, p. 9) puts it, “analysts of science not only should consider the contemplative-representational-ontological dimension of science, but should also examine the transformative-technical-pragmatic dimension of science, with its material, somatic, skilful and utilitarian aspects.”

It is beyond the scope of this chapter to trace the history and etiology of the turn to practice,⁷ and additionally, it is not any one doctrine or thesis. This makes it difficult to succinctly state *what* the turn to practice *is*. Nevertheless there are several common features which I will highlight as being of particular interest and importance here. In their introduction to a collected volume focusing on the turn to practice⁸ the editors offer up a collection of ‘shifts’ which characterize the practice turn. Of the six they discuss I will focus on just two:⁹ the shift from science-as-contemplation to science-as-transformation (of the world), (p. 22) and the shift from scientific products to processes. (p. 21)

The first shift (from *contemplation* to *transformation*) concerns (among other things) the value which we place in science. Under the logico-theoretical view the value of science is to be found in the propositional and factive character of its products. Science is worthwhile, under this view, primarily because it delivers to us true (or as rigorously justified as we presume possible) propositions which we can then hold in contemplation, or in Rorty’s analogy,¹⁰ with which we can bring our minds to more perfectly mirror the world. The turn to practice denies this as the primary value of scientific practice itself, and even as the value of those parts of science typically taken to be most straightforwardly representational. “[T]heories not only are representations that must

⁶To put a finer point on it, the turn to practice does not ignore what scientists say, but instead recognizes that *saying* is fundamentally a kind of *doing* and so is to be understood itself in the wider context of ‘doings’.

⁷Soler et al. (2014a) gives a comprehensive accounting of this history.

⁸Léna Soler et al., eds. (2014b). *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. Routledge

⁹Not for lack of interest or importance of the remaining four, but solely for the constraints of economy and relevance to the present concerns.

¹⁰“The picture which holds traditional philosophy captive is that of the mind as a great mirror, containing various representations—some accurate, some not—and capable of being studied by pure, nonempirical methods.” (R. Rorty, 1979, p. 12)

be assessed for their empirical adequacy, but they also are, perhaps primarily, material means to achieve cognitive aims.” (Soler et al., 2014a, p. 23) Representations, on this view, are pieces of technology, no different in that respect from a microscope or table saw.

The recognition that science is connected inextricably to the development and application of the methods and means to actively transform the world, both in the material technology it devises and promulgates, as well in the conceptual and calculational apparatus of its theories, has in turn forced a reckoning in the role that science plays in the arenas of politics and culture. Without the reassuringly abstract picture of science as impartial, disinterested, and detached from the more mundane exigencies of everyday life, it becomes clear that the question of *what* scientists choose to do, are incentivized to do, are prohibited from doing, and refuse to do (etc.) demands reflection on the values and aims of a wider population.

The second of these shifts (from *product* to *process*) emphasizes the “the dynamic unfolding of actual scientific activities and ... the actual historical process through which particular achievements come to acquire the status of knowledge.” (2014, p. 21) This shift can be understood in two (not exclusive) senses. First, this shift tracks a recognition of the messy, complicated, ambiguous, and fundamentally open-textured nature of the human activities placed within scientific practices. In this sense it marks a shift in *emphasis* or *priority*—that these activities play a central role in the corpus of scientific knowledge, and perhaps reveal more about the nature of that knowledge than what is possible by any examination of theory abstracted from this context. The second sense marks this shift as one in the *intelligibility* of science. That is, according to this shift, in the absence of an awareness and appreciation of the embodied, ‘somatic’, discursive, and social aspects of science, the study of science as purely abstract theoretical enterprise is at best a sort of ‘false consciousness’.

What I’d like to emphasize about this shift is the call to look at the local, practical, detailed features which constitute the activity of science. It is not enough, according to the practice-turn,

to merely theorize about what science is like, how it justifies, or even formulates its claims—one should *look* to see what actually happens.¹¹ What one will find when they look and see is a rich and tangled nest of tacit skills, educated guesses, and (don't forget!) endless roadblocks, confusions, caveats, and failures.

Take, for example, the famously 'simple' experiments performed by Galileo to study the motion of objects under uniform acceleration, and with which he purportedly discovered his law of fall. Study of this experiment, especially study of the practical arrangement and operation of the various parts helps to show how non-trivial it is to measure something as humble as rectilinear motion. This simple example highlights the fact that many typical discussions of Galileo and the inclined plane treat the concrete experiment as either trivial, ancillary, confabulated, or even simply a rhetorical fiction. On the other hand, those who have undertaken to perform these same experiments have something to teach us about the delicacy of reliably intervening in the world. Indeed, assumptions made about (for example) the veracity of Galileo's claims of precision, reveal a misapprehension of the actual physical skills required and acquired by the experimentalists. These misapprehensions in turn belie a distortion of our view of the world that affects the metaphysical estimations one might make.

Galileo, in the *Discourses* describes the setup as follows:

SALV. [...]

A piece of wooden moulding or scantling, about 12 cubits long, half a cubit wide, and three finger-breadths thick, was taken; on its edge was cut a channel a little more than one finger in breadth; having made this groove very straight, smooth, and polished, and having lined it with parchment, also as smooth and polished as possible, we rolled along it a hard, smooth, and very round bronze ball. Having placed this board in a sloping position, by lifting one end some one or two cubits above the other, we rolled the ball, as I was just saying, along the channel, noting, in a manner

¹¹Wittgenstein's exhortation (during his discussion of family-resemblances) comes to mind: "Don't say: 'They *must* have something in common, or they would not be called "games"' — but *look and see* whether there is anything common to all... To repeat: don't think, but look!" (Wittgenstein, 2009, §66)

presently to be described, the time required to make the descent. We repeated this experiment more than once in order to measure the time with an accuracy such that the deviation between two observations never exceeded one-tenth of a pulse-beat. Having performed this operation and having assured ourselves of its reliability, we now rolled the ball only one-quarter the length of the channel; and having measured the time of its descent, we found it precisely one-half of the former. Next we tried other distances, comparing the time for the whole length with that for the half, or with that for two-thirds, or three-fourths, or indeed for any fraction; in such experiments, repeated a full hundred times, we always found that the spaces traversed were to each other as the squares of the times, and this was true for all inclinations of the plane, i.e., of the channel, along which we rolled the ball. We also observed that the times of descent, for various inclinations of the plane, bore to one another precisely that ratio which, as we shall see later, the Author had predicted and demonstrated for them.

For the measurement of time, we employed a large vessel of water placed in an elevated position; to the bottom of this vessel was soldered a pipe of small diameter giving a thin jet of water, which we collected in a small glass during the time of each descent, whether for the whole length of the channel or for a part of its length; the water thus collected was weighed, after each descent, on a very accurate balance; the differences and ratios of these weights gave us the differences and ratios of the times, and this with such accuracy that although the operation was repeated many, many times, there was no appreciable discrepancy in the results.

(Galilei, 1914, pp. 178–179)

How has this description of the experimental setup been received in accounts of Galileo's discovery of the law of fall? In the book *Great Experiments in Physics* (Shamos, 2012), Galileo's law of fall is given an entire chapter. However, this book—focused as it is on *Experiments*—devotes 16 pages to Galileo's (albeit ingenious) abstract geometric constructions, and *one* page to the

description I've included above, with virtually no commentary on the latter. Clearly, for the author, the actual description of the physical apparatus, and even the fact that it was employed to collect data, are at best ancillary.

Perhaps this attitude toward Galileo's experiments can be explained by the general estimation of the actual worth of those experiments in past decades. For example consider Albert Einstein's assessment of Galileo's experiments:

The experimental methods at Galileo's disposal were so imperfect that only the bold-est speculation could possibly bridge the gaps between empirical data. (For example, there existed no means to measure times shorter than a second.) (foreword, Galilei, 1967, p. xix)

And Einstein is not alone, here is Alexandre Koyré's impression of Galileo's description:

A bronze ball rolling in a "smooth and polished" wooden groove! A vessel of water with a small hole through which it runs out and which one collects in a small glass in order to weigh it afterwards and thus measure the times of descent (the Roman water-clock, that of Ctesebius, had been already a much better instrument): what an accumulation of sources of error and inexactitude!

It is obvious that the Galilean experiments are completely worthless: the very perfection of their results is a rigorous proof of their incorrection." (Koyré, 1953, p. 224)

Koyré dismisses out of hand the evidential value of Galileo's experimentation. Instead, Koyré assumes that the law of Galileo's was derived through pure intellect, and the description of the experiment was a harmless fiction to convince the reader its testability in principle. Koyré's (and Einstein's) chief complaint in Galileo's written account is that of time-keeping. He finds it incredible that Galileo could achieve the accuracy he claimed (on the order of 0.1 sec). Why does he think so? Examining the time-keeping mechanisms available to Galileo convinces Koyré that none were reliable clocks to the requisite precision. He does not, however, attempt the

experiments himself, and is satisfied by the abstract analysis of time-keeping machines. He makes, on my view, several unjustified assumptions, but principle among them is the assumption that theoretical accuracy of time-keeping mechanisms is a reliable surrogate for the estimation of an experimenter's accuracy. On examination this assumption is simply not borne out, and our discovery of this error comes from several researchers who have set out to *actually* perform Galileo's experiment.

Koyré himself points to several of Galileo's contemporaries (specifically, Marin Mersenne and Giovanni Battista Riccioli) who sought to repeat his findings. However, neither Mersenne nor Riccioli attempted to repeat Galileo's experiment as he describes it, and Koyré comments approvingly on the experimental superiority of their methods. Regardless of the concrete empirical improvements which these methods might have offered (they both conducted experiments on freely falling bodies), these hardly count as evidence for the impossibility of the sorts of measurement and precision claimed by Galileo.

For this we can look to more recent attempts to reconstruct Galileo's experiment, and to see if such precision is possible. The first is a report from Paul D. Sherman (1974), wherein Sherman along with students in his physics class set out to perform Galileo's experiment as close to the original description as feasible. Their question was one of possibility—were Galileo's experimental methods sufficient to (1) justify his law of fall, and (2) achieve his claimed precision, or is Koyré's incredulity justified?

There are several fascinating things to extract from Sherman's report. First is to note the non-triviality of several aspects of the construction and intention of the inclined plane itself. As Galileo reports, the groove was made "very straight, smooth, and polished," and lined with likewise smooth and polished parchment. The question to ask, then, is whether the bronze ball used in the experiment was rolled along the edges of the groove, or wholly within it. A groove 'a finger's breadth' in width poses some interpretive difficulty, first because there is no standardization of such a measure, and so no good estimation of how to convert it into modern units.¹²

¹²However, in Galileo's case we can actually describe a concrete process for which we could determine the size of the groove quite accurately (supposing that Galileo's used *his* finger for the measurement). Fortuitously, Galileo's

Second, a groove of such a diameter (probably on the order of 1.5–2.5 cm) falls within a range of plausible sizes for bronze balls such that the question of rolling within the groove or on its edges is wholly unclear. Sherman notes that a previous attempt at recreating Galileo’s experiment by Thomas Settle (1961) assumed the latter, rolling the ball along the edges. Sherman, in contrast reasons that Galileo would not have noted the effort to ‘smooth and polish’ the groove if this had been the case. Thus, he chooses a ball fit to roll inside the groove.¹³ Sherman’s analysis of the text is not decisive, however, since the easiest way to polish and smooth the *edges* of a groove is to simply polish and smooth the *entire* groove.

Nevertheless, for the present purposes it is sufficient to note that this particular question, and the consequent interpretive and practical task of determining which arrangement to use, is almost invisible to the reader who assesses the experiment abstractly in contemplation, but impresses itself at once upon the experimenter since a choice *must be made*.

Besides the non-triviality of the actual construction of the inclined plane and the accompanying choices,¹⁴ the second feature to extract from Sherman is the need for *skillful* execution of the measurement processes required.

[The] question is one of coordination. It is obviously necessary to start the water at the moment the ball is released and stop it at the end of the run. The distances along the groove were marked as increasing from the bottom to the top of the slope. At the bottom end we placed a tin bucket to the ball made a sharp “clang” at the end of its run and it was fairly easy for the person at the burette [which served as the water-clock] to coordinate with this. Much more difficult was coordinating the timer

right middle-finger was liberated from his remains when they were exhumed and reburied in 1737. These finger bones are on public display in Florence. Using the techniques of modern forensic anthropology, we could establish with a high degree of confidence the width of his index finger (which is most likely the unit measure). That said, determining the precise width of the groove is probably not that important.

¹³Notably, both Settle and Sherman report precisions and accuracy within the range described by Galileo, so the choice appears unimportant. This unimportance could only have been determined, however, by conducting experiments in both configurations.

¹⁴One further experimental question—which I’ll only note in passing—is whether Galileo operated the device alone, or with assistance. Sherman offers textual evidence that he had at least one assistant, but again, *such a choice* could very well affect the outcome of the experiment. Indeed this choice has downstream consequences for the second feature which follows.

with the ball's release. It is clear from Galileo's own description that he experienced similar difficulties which is why he states "we repeated the experiment more than once ..." and "having assured ourselves of its reliability .." and "the operation was repeated many, many times." *Practice is needed.* (Sherman, 1974, p. 346, emphasis added)

Sherman and his students were able to coordinate their individual tasks to a sufficiently precise degree that their timing error was around 0.1 sec. What is also notable is that this error shrinks as the number of trials increases, strongly indicating that the skill in coordinating the water clock and the ball was improving with practice. Settle reports similar error in timing (it is notable that Settle experimented alone).

The historian of science Stillman Drake conducted his own version of Galileo's experiment, (S. Drake, 1975) aided by a (then) recent discovery by Drake of a page from Galileo's unpublished notes. These notes contain recorded data from what was ostensibly several trials on the inclined plane, as well as calculations fitting the $d \propto t^2$ law. The details of Drake's ingenious detective work are extraneous to our concern here, and we'll focus instead on Drake's proposal for Galileo's method of time-keeping. The note (named *f. 107v*) records distances traveled over 8 equal measures of time. Drake's key concern is how Galileo could have reliably apportioned these measures, especially with sufficient accuracy given the recognized paucity of accurate time-keepers. Drake points out that lack of accurate time-keeping mechanisms in no way entails the inability to keep accurate time, especially over short spans. The particular evidence Drake offers is musical time-keeping—humans have a keen sense of rhythm and meter: incongruities in musical pace on the order of (Drake estimates) $1/50$ th of a second are noticeable. Exploiting our rhythmic talent, Drake speculates on a method for dividing out equal measures of time like Galileo:

Place a grooved inclined plane about $6\frac{1}{2}$ feet long at an angle of about 1.7 degrees. Fit a stop at the higher end, against which a steel ball can be held by resting a finger on it lightly. Now sing a simple march such as "Onward, Christian Soldiers" at a tempo of

about two notes per second,^[15] very crisply. When the tempo is established, release the ball at some note and mark with chalk the position of the ball at other notes (half-second^[16] intervals). In three or four runs eight positions can be marked; put a rubber band around the plane at each mark... Then, making many repeated runs, adjust the rubber bands so that the audible bump made by the ball in passing each one always coincides exactly with a note of the march. When the inclined plane has thus become a kind of metronome, measure the distance in millimeters from the resting position of the ball to the bottom of each rubber band. (S. Drake, 1975, p. 101)

Whether this method was what Galileo *in fact* used is uncertain (though Drake make a persuasive case). What is important in Drake's method, and those of Sherman and Settle is that they demonstrate that Einstein and Koyré's presumption is *false*. These latter two are correct to say that no mechanism existed which could reliably measure fractions of a second. However, they neglect two important features of Galileo's experiment. First, what Galileo measures are *relative* spans of time¹⁷ and so measuring relative quantities (such as water) are sufficient for the desired accuracy, independent of any periodicity (or lack of it) in the time-keeping device. Second, by coupling the motion of the ball to an audible signal (as Drake hypothesizes), one can discern irregularities in periodicity (and thereby tune the distances) far more accurately than available time-keeping mechanisms would allow.

What lesson should we draw from these considerations? Einstein and Koyré suppose that Galileo's law must have been discovered by the operation of his intellect and abstract reasoning. Indeed, throughout the *Discourses* he demonstrates his considerable creativity and genius. However, Einstein and Koyré come to have a distorted view of what Galileo was *actually able to do*

¹⁵That Drake includes this comment describing the tempo in seconds defeats the purpose of this method, but it can be removed without affecting the experiment. Drake included it, no doubt, to give the reader a sense of the tempo, but of course any tempo (within reason) would produce correct results. In fact, the reason he is giving quantitative measurements at all is because he is describing a physical setup which will recover *the same data* that appear in *f.107v*.

¹⁶See fn. 15.

¹⁷In fact, Galileo does not relate distance and time directly. What he reports is an equality between ratios of distance and ratios of their respective time squared. "The spaces described by a body falling from rest with a uniformly accelerated motion are to each other as the squares of the time-intervals employed in traversing these distances." (Galilei, 1914, p. 174)

based on their admiration for the abstract and mathematical portions of Galileo's work. Only when the experiments are attempted are the actual procedural and material challenges manifest, and in ways which contemplation of the experiment cannot be guaranteed to identify. It is all too easy to fail to distinguish mechanical time-keeping and the skillful discernment of intervals of time. It is through practical attention (and attempting) that these small but crucial truths appear. If we are to look at science as being relevant to metaphysics, it will be in its ability to reveal the extent and limitations of these sorts of interactions.

What are we to make of the turn to practice in the present context? If we suppose that science licenses metaphysics we must ask after the nature of such a license, and in turn what sort of metaphysics it licenses. Once we begin to emphasize the active and embodied aspects of scientific practice it becomes less clear that there is anything like a cohesive ontological picture where one may step back and generalize the deep and unifying principles as one might do when contemplating theory. *What is it* then, that science licenses?

For the naturalist-pragmatist science itself *constitutes* metaphysical inquiry—but in a sense importantly qualified so as to divorce itself from the classical view which (for example) structuralists buy into. There is, this view asserts, nothing over and above the engaged practice of inquiry—we *do* metaphysics¹⁸ when we confront nature, when we skillfully construct delicate, causally isolated, novel micro-worlds, when we develop the skills and tools to not only ask, but *formulate* and *conceive of* questions which we 'put to nature', when we recognize the intimate interplay between our interventions on our world and our aims and ends—and when we engage in normative conversation about *which* aims and ends we ought to seek, and *how* we should like to actively transform our world.

¹⁸The obvious rejoinder here is that I am simply describing something wholly different from 'metaphysics' and so I should just abandon the term. I address this objection (to which I am sympathetic) in the conclusion of this chapter and offer some practical reasons why insisting on keeping the term is worthwhile for the naturalist.

4.1.2 METALINGUISTIC NEGOTIATION

Happily, and quite independently of the considerations above, philosophers working in what might be thought of as traditional analytic metaphysics have developed a deflationary view of metaphysical inquiry which casts it as a *distinctly* normative-conceptual enterprise of the sort which appears to be integral to a natural-scientific approach to metaphysics. Metaphysics, in these deflationary terms, is concerned with the adequacy and appropriateness of concepts, with an eye to “engineering [or] (re-)designing concepts to better serve certain functions.” (Thomasson, 2017a, p. 364)

Amie Thomasson has suggested a deflationary account of ontology which does not chalk up metaphysical talk as meaningless, but instead understands it as pragmato-normative. That is, ontological claims, though they occur in the object language, and appeal to justification which is ostensibly epistemic, are actually best understood as disguised meta-linguistic normative claims about concept formation, application, and adequacy. The appeals to what appears to be epistemic justification are themselves to be understood as normative assertions about use. Interestingly, though, one of Thomasson’s motivations for this sort of deflation is a perceived threat of metaphysics “falling prey to a rivalry with science.” (Thomasson, 2017b, p. 101) She seems to think that if ‘heavyweight metaphysics’¹⁹ goes toe-to-toe with science it is the heavyweight metaphysics which will come up short—“for isn’t it the purview of physics to discover the deep and fundamental facts about reality, and doesn’t it do so better than metaphysics?” (ibid.) So while Thomasson seems content to cede authority to discover ‘deep’ facts to the sciences, her approach is most interesting here for its application *to* scientific metaphysics—for the epistemic situation in science is not as sunny as she seems to think (a point which I’ve argued for in Chapter 1). Thomasson’s deflationary approach to metaphysics was not born from the motivating factors which concern us here. Nevertheless, I will make free use of this framework for ontology, since it matches nicely some features which a naturalist-pragmatist metaphysics sees the need for.

¹⁹This is Thomasson’s term (elsewhere she uses the term ‘serious’) for metaphysicians who “think of themselves as making discoveries about what really exists, and about [e.g.] the persistence conditions or modal properties of things of various sorts.” (Thomasson, 2017b, p. 102)

The fundamental concept at work in this account is that object language disputes which purport to be directly about ontology, can and should be analyzed in terms of *metalinguistic negotiation*. Introducing this idea Thomasson canvasses three curious features of ontological disputes: First, these debates “are typically not resolvable by empirical means—one cannot just wait for the decisive observation or experiment to determine whether [for example] a person can survive loss of memories or ... whether numbers exist.” (Thomasson, 2017b, p. 109) Second, even if it is suggested that these debates are ‘merely verbal’ and can thereby be dissolved by clarifying uses of language, “in the eyes of the disputants, they cannot be settled by such methods as these ... disputants in serious metaphysical debates typically deny they are each using the disputed term (‘number’, ‘person’) in a different sense.” (ibid., p. 109) And third, in cases where the disputants *do* recognize they’re using the term differently, *still* the debate is not resolved. The standard explanation of these odd dialectical features is due to their being about *genuine metaphysical facts*. Thomasson rejects this explanation (as mentioned above) on account of her wish to avoid a ‘rivalry with science’.²⁰ Instead, these features can be explained in terms of metalinguistic negotiation. (Plunkett and Sundell, 2013, p. 3)²¹ Here, the same features arise regarding disputes:

1. They don’t go away even if the disputants agree on all other ‘facts’—they can’t be empirically resolved just by correcting misinformation or adding empirical information;
2. They don’t go away even if the disputants agree about how the word is actually used, or are given full empirical information about how the word is used.
3. They don’t go away even if disputants recognize that they are using the term differently.” (Thomasson, 2017b, p. 110)

Take, for example, two friends watching the Winter Olympics, and arguing over whether figure skating is a sport:

The dispute is not about any of the ‘facts’ about what figures skaters do, what sorts

²⁰ We shall reject this explanation because it helps itself to the *a priori*.

²¹ “[B]y putting forward competing claims [about the term under consideration], speakers can, via metalinguistic uses of their terms, debate how it is those terms should be employed. We call a dispute like this—one that employs competing metalinguistic usages of an expression, and that reflects a disagreement about the proper deployment of linguistic representations—a *metalinguistic negotiation*.”

of training and skills are required, or how the competitions are judged. Disputants might agree about all that. One might interpret the disputants as using the term ‘sport’ in different ways (one of whom will apply it where ‘artistic impression’ scores play a role and the other won’t) and each speaking a truth in their own idiolect. But even if that is pointed out to them, they will not give up the dispute as merely verbal and go home. Similarly, even if we take them both (owing to semantic deference) to use the term with the same, public meaning, and could appeal to data from linguistics or experimental philosophy to show that (given actual usage) figure skating *does* (or doesn’t) meet the relevant criteria, that would not end the dispute. (Thomasson, 2017b, pp. 110–111)

The key here is to recognize that the dispute is poorly modeled as the sort of dialectical back-and-forth which we philosophers are wont to assume is the paradigm of discourse. That is to say, the dispute *is not* best understood as an *argument* in the philosopher’s sense. The claims made are better analyzed as a sort of *performance*.

In uttering ‘figure skating is/is not a sport’ each is *pragmatically communicating* views about whether the term ‘sport’ *should* be applied to figure skating. And such disputes may be very much worth having. For how we use words *matters*, both given their relations to other aspects of our conceptual scheme, and to our non-verbal behavior. Treating figure skating as a sport, for example, is connected to a range of types of societal honors and rewards—to appearing in the Olympics, to receiving sponsorships and television coverage, to honoring its practitioners in all the ways athletes are honored. What is at stake in arguments about whether figure skating is a sport, then, is a range of *normative* issues about how the skaters, fans, competitions, etc. *are to be treated*. (ibid., pp. 110–111)

The disputants are not *searching for*, let alone *discovering*, the truth about figure skating. In partaking in such a debate they are staking out normative positions, and *communicating* commitments

which they take to be valuable. Understood this way, we can easily explain the heretofore odd features of such a debate. Normative claims about what a term *should* mean, and how a concept *ought* to be used are simply not the sorts of things which could be countered by empirical reports. Likewise, even in the face of descriptive linguistic factors, normative disputes will not be resolved, since the motivation is, typically, to *change* the descriptive facts about usage (if such exist).

A naturalist-pragmatist view of metaphysics makes sense of ostensibly object-language claims about ontology in the sciences in these terms. Questions and debates such as “are there electrons?” or “only structure exists.” are best understood as staking out positions in a normative space, specifically as communicating *how such terms and concepts are to be used*. This, in turn, is expressive of what one takes to be valuable, both as an object of study (that is, how the practice should ‘go on’²²), but also in terms of our shared worldview and the distribution of valuations over various political, cultural, and ideological ends.

4.1.3 A PRAGMATIST VIEW OF CONCEPTS

The third and final thread which intersects our interests here is the philosophical outlook of the American Pragmatists (of whom I would count from Peirce up through—and past—Rorty). Pragmatism is of particular interest for the naturalist owing, first, to the clear line of ancestry through which the turn to practice (see above) can trace its origins, and second, to the overt invocation of pragmatism in Thomasson’s account of metalinguistic negotiation (especially its end-oriented nature), but also third, from the deeply interwoven threads of naturalism and pragmatism which one sees throughout its (pragmatism’s) history.

The commonality amongst pragmatists is their emphasis on and attention to the connection between *meaning* and *doing*. I have suggested, indirectly and directly, over the past chapters, that science, if it to be understood as adjudicating the articulation of concepts and their meanings at all, if it is to be accorded any sort of authority in this regard, then proper attention to the nature of that authority and adjudication must come by way of the particular activities that are special

²²cf. Wittgenstein, 2009, §151 ff.

to the sciences, both in the sense of activities which give rise to speaking in, and about science, and how speaking in and about science makes available new activities and practices.

This is perhaps an overly dense formulation of the picture I have in mind, and it will be the aim of the balance of this chapter to spell out this pragmatic view of a scientific metaphysics. How we arrived at this point, recall, was through examination of well-intentioned attempts to let science dictate our understanding of the metaphysical structure of the world. The failure of these attempts turns on a conflict between the naturalistic repudiation of the traditional metaphysical project, and the realist desire to reclaim the territorial holdings of that very same project.

The naturalist suspects traditional metaphysical profligacy for its shameless pandering to the conceptual home range of human beings—building up the purported deep structure of the world in terms of concepts and distinctions tied to our linguistic and material heritage. Such stories, the naturalist worries, give humans a false sense of self, an overlarge estimation of how ready-made the world is for creatures like us. However, the naturalist missteps if their response to the angel-on-pinhead-counting, joint-carving, from-nowhere-viewing of traditional metaphysics is to look into the buzzing confusion that is the natural sciences, searching for the *correct* story, an alternative story, which (as seems right to the naturalist) dislocates humanity from center stage, and instead tells the story in only its pure, natural details—*sub specie aeternitatis*.

My suggestion is that attention to the sciences in all their diverse, distributed, disjointed aspects, *proper* attention to their fruits, costs, struggles, missteps, and serendipities, reveals *no* story to extract, no correct list of facts that will set the more dreamy traditionalists straight. Science is not a continuation of and improvement on the idle dream of putting our mind into the correct relation with the world. Science is a renunciation of the article of faith that such a project is desirable, that such a relation is coherent and intelligible—the goal of inquiry is not a transcendent experience of Nature in place of God, the goal of inquiry is to facilitate the ongoing and continuous contact we already have with the natural world, of which we are an active part.

The view I present here sees the relationship between naturalism and pragmatism as one of intimate contact. Attention to science does two things: First it undermines the idea that there

is anything like a ready-made world and the idea that there is something like a transcendent authority legislating from above. Second, it refocuses our attention to the physical, embodied, transactional character of our relation to the world, and the influence on our conceptual scheme of such connected, concrete, contact with the world. These two aspects are themselves central features of the pragmatist tradition. This relationship between naturalism and pragmatism is expressed (albeit clumsily) in the following slogan: “Naturalism tells us to take science seriously, and science taken seriously tells us to be pragmatists.”

4.2 Metaphysics in Practice: Astrobiology and Life

It’s now time to put the picture of a pragmatic-naturalist metaphysics into practice. The topic of consideration for the remainder of this chapter is the field of astrobiology, specifically focused on the concept around which it is organized: **Life**. As we’ll see, there is considerable confusion within the field, with no real prospect for resolution of the difficulties these confusions give rise to. Consider the following quote:

The origin of life is widely regarded as one of the most important open problems in science. It is also notorious for being one of the most difficult. (Walker, Packard, and Cody, 2017, p. 1)

Why is it so difficult?

Our uncertainties in the most fundamental properties of life have left the origins-of-life researchers with a conundrum—*without an understanding of what life is, how can we approach understanding its origins?* Current hypotheses for the origins of life are motivated by the varied definitions for life (and tend to lie along disciplinary boundaries). One popular working definition for life is ‘a self-sustaining chemical system capable of Darwinian evolution’ (first emerging from the NASA Exobiology Discipline Working Group[†]), motivated by an assumption that it is genetic inheritance

[†]David W Deamer and Gail R Fleischaker (1994). *Origins of Life: The Central Concepts*. Jones & Bartlett Pub

with vertical descent that is life's most basic property. (Walker, Packard, and Cody, 2017, pp. 5–6)

There is an explicit need in the astrobiology community for a clear sense of *what* they're looking for, *what* they're studying, and in general *what* criteria to apply, not just in applications of theory, but in the practical design of observational techniques and exoplanetary science missions. For example, the inclusion of soil analysis experiments in the Viking landers and the Curiosity rover not only indicate our general interest in the possibility of past or present life on Mars, but reveal important conceptual constraints on what could count as evidence at all. In the two examples given (i.e. the experiments of the Viking missions and Curiosity mission) the choice was to search for the presence of organic compounds.²³ These design choices variously expose and foreclose specific possibilities for discovery, and in the present example these possibilities were narrowly prescribed—the Viking and Curiosity experiments would only find signs of life if that life was essentially identical to terrestrial life. This is not to criticize the choice, but only to make it explicit that in cases of seeking evidence of life, our understanding of that concept has real consequences for the possibility of discovery.

4.2.1 DEFINITIONS OF LIFE, THE PROBLEM

The discourse surrounding the search for life in astrobiology centers on the idea that what is needed is a *definition* of life. What we might call the received definition is the one stated in the quote above—often called the 'NASA' definition:

Definition 1. *Life is a self-sustaining chemical system capable of Darwinian evolution.*

There are numerous obvious counterexamples to this definition, and it suffers due to its vagueness in many respects. First, 'self-sustaining' does not appear to have any useful, context-free

²³In the case of the Viking landers several of the experiments in the 'biology instrument' were designed with the explicit expectation to measure metabolic products—identical to or in close analogy with metabolic processes of terrestrial organisms. Curiosity's soil-analysis package ('SAM') measured for organic chemicals similarly to the Viking experiment (gas chromatograph-mass spectrometry) and also measured isotope ratios in CO₂ and CH₄ to distinguish between inorganic and biological origins—again modeled on isotope ratios in terrestrial samples (see Brown et al., 1978; Wikipedia, 2017—respectively—for more details).

sense. Certainly organisms are not self-sustaining in any absolute sense, since it appears that all life catalogued to date requires some thermodynamic gradient to exist.²⁴ Even within such a gradient most²⁵ organisms ‘self-sustain’ for only finite spans of time. On the other hand there are non-living systems which appear to be likewise self-sustaining on time-spans which outlast the life-spans of some prokaryotes (and indeed some eukaryotes). Fluid vortices (e.g. tornadoes) can be relatively long-‘lived’ under ideal circumstances, the longest recorded tornado lasting ~ 3.5 hours, well past the minutes-long lifetime of some bacteria. Indeed, extraterrestrial examples of fluid vortices have much longer lifetimes, Jupiter’s ‘Great Red Spot’ being a notable example, with an age of at least 400 years.²⁶

Other definitions suffer from similar difficulties. Here is one from the chemist Leslie Orgel:

Definition 2. *Life is a complex, information transforming, reproducing object that evolves by natural selection. (in Raulin, 2010, p. 192)*

Another, from Luisa Damiano and Pier Luigi Luisi:

Definition 3. *A living system is a system capable of self-production and self-maintenance through a regenerative network of processes which takes place within a boundary of its own making and regenerates itself through cognitive or adaptive interactions with the medium. (Damiano and Luisi, 2010, p. 149)*

And one more from Radu Popa:

Definition 4. *Living entities are self-maintained systems, capable of adaptive evolution, individually, collectively or as a line of descen[t]. (Popa, 2010, p. 188)*

Edward Trifonov suggests an approach where we can extract a definition of life by analysis of commonalities across various definitions. Trifonov performs a frequency analysis on 123 defi-

²⁴For example, most plants exploit the solar-thermal gradient, converting high-quality sunlight to low(er)-quality waste products and heat in order to drive their metabolic processes. (see e.g. Wolfe, 2015)

²⁵With the possible exception of some members of the phylum Cnidaria which appear to be effectively immortal. (Martínez, 1998)

²⁶We will return to the ‘life-like’ character of vortices below.

nitions (embedded in their surrounding papers) and extracts nine properties (*definiencia*) for life, offering the following definition-by-consensus:

Definition 5. *Life is [a] metabolizing, material, informational, system with [the] ability of self-reproduction with changes (evolution), which requires energy and [a] suitable environment. (Trifonov, 2011, p. 260, emphasis original)*

Such examples can be multiplied. What I wish to emphasize in considering all five definitions is not so much their shared commonality of predicates ('self-sustaining', 'self-maintaining', 'self-reproducing', etc.) but the underlying structural character of these definitions. They all seek to offer perfectly general criteria for identifying an entity as alive—ostensibly what these definitions practically recommend is that when we wonder whether x is an instance of life we simply observe its properties, and compare them with the chosen definition. This is not, however, how these definitions are typically received. Instead, they are subjected to antecedently determined cases of living and non-living systems, and evaluated on how well they partition the cases.

But this treatment is somewhat mysterious if we are to take seriously that what our goal in developing a definition is to do is stipulate clear and operational criteria for the practical application of the concept in scientific inquiry. Notice that the counterexample-proffering behavior is distinctly *not* seen as legitimate in other exercises of definitional stipulation. When the International Astronomical Union defined a planet in such a way as to rule out Pluto, this was not seen as evidence²⁷ for the inadequacy of the definition. Compare this with the responses to (for example) the NASA definition. For example, Benner (2010) notes that, according to the NASA definition, a single rabbit doesn't count, but a breeding pair does—the implication being that *obviously* a single rabbit *should* count. If such definitions were *genuinely* stipulative then it should at least seem like a plausible option that we could count out single rabbits. What's gone wrong?

²⁷At least among IAU members and astronomers at large. The public reaction was notably different.

OUR VITALISTIC HERITAGE

Life, as a concept, is tightly linked with the folk-theory of vitalism. To see this, it's first worth noting that the question of "what is life?" was not a question of scientific urgency until relatively late in the scientific revolution. Indeed, it was only once the atomic theory, and the beginnings of biochemistry were underway that the question of life arose. Why is that? The central reason for the recognition of this need was that the atomic theory, and its ramifications for material objects, especially biological ones, threatened to undermine the central sustaining concept of life which had persisted since Aristotle—vitalism. Vitalism is, roughly, the thesis that life is an essence (or substance) which animates a material body, explaining the special *active* nature of living things, and explaining both their manifest differences from inanimate material as well as their transformation back into inanimate material. Such a view faced considerable pressure from the newfound abilities of biochemists, like the synthesis of urea, and ultimately, the discovery of DNA.²⁸

When vitalism was the dominant account of life, there was no need to ask after a stipulative definition. Certainly there were investigations into vitalism's nature, but the conceptual boundaries of life were easily limned—all and only those material objects imbued with the *élan vital*. Empirically tricky, but conceptually rock-solid. The contemporary difficulty which the sciences, particularly astrobiology, face is the disconnect between our language and concepts, which are still very much tied up in vitalism, and our scientific picture of the world, which leaves no room for any such thing.²⁹

The latent vitalism in our concept of life is revealed in our linguistic habits. 'Life' has the grammatical form of a sortal, designating a kind. Kind terms serve to group objects in virtue of their type or genus, which, in turn, an object possesses due to their having appropriate *essential* characteristics. This grammatical form encourages questions of the form "What is a living sys-

²⁸Prior to the synthesis of urea in a laboratory proponents of vitalism thought that there might be a distinction in the *substance* of living and non-living matter. Likewise in the case of *information* prior to the description of DNA.

²⁹In many respects this problem shares its structure of conflict between our habitual linguistic and conceptual practices and our scientific picture with the better-known mind-body problem.

tem?” modeled on questions like “What is a mammal?” or “What is a triangle?” Questions of this form are seeking the essence of a kind. The normal (natural language) use of the term ‘life’ carries with it the suggestion that it designates an essence—an essence which we have a pretheoretical capacity to grasp.

Thus, when an ostensibly stipulative definition is put forward, it is simply not treated as such—to be a satisfactory definition it should meet all and only my intuitive criteria for identifying life, criteria which are not forthcoming for the simple fact that they depend on something for which there is simply no room in our best scientific picture of the world. In our attempts to define life there is an essential conflict between our concepts (which are vitalistic) and our scientific picture of the world (which is, roughly, materialist).

One response to the recognition of this conflict would be to attempt to legislate our pre-scientific concepts of life away. We could be more explicit about the stipulative force of any given definition, and simply refuse to entertain possible counterexamples. I imagine the reader will agree that such a response is not attractive. Legislating concepts is rarely effective, and in those cases where it is (as in the IAU’s Pluto example) the antecedent concept has held considerably weaker intuitive appeal³⁰ and even then popular resistance is significant. The vitalistic picture is a strong one, bolstered by our habitual use of language and the manifest experience of the wide gulf between living and non-living matter. One valuable thing which our vitalistic concept permits is that it allows us to imagine alternatives—creatures wildly different from the things living on Earth, but which by virtue of we-know-not-what are equally alive. Our latent vitalism, seen in this light, possesses the virtue of keeping our idea of life conceptually open and receptive.

³⁰It should be of no surprise that the layperson does not have a robust notion of ‘planet’ already up and running in their heads. The concept of planets has been historically quite plastic considering the Sun and Moon have been classed as planets.

TRADITIONAL METAPHYSICS

Nevertheless the conflict between our vitalistic legacy and our scientific picture remains, and does appear to stand in need of some resolution. We should notice here that this is a genuinely metaphysical question—what is the ontology of life? However traditional metaphysics, caught up in offering an eternal, essential verdict, cannot adequately resolve the issue to the satisfaction of either the naturalist or the astrobiologist. Indeed, the traditional metaphysician can only oscillate between reductionism and vitalism. In analyzing our concepts of life traditional metaphysics seeks to offer definitive reasons for *what exists*. For example, in the vein of Canberra-Plan style analytic metaphysics, one might try to fit our locutions about life into our ‘best total theory’ and determine what such a theory quantifies over. But such a process must result in one of two options: Either our analysis of total theory militates against adding any entity or universal ‘life’ to our ontology—thus constituting a reductionist program, reducing away the concept;³¹ Or such an analysis would suggest *adding* to our ontology a new thing—something like a vital force.³²

Neither approach *resolves* the conflict between our concepts and our scientific picture. Instead they simply pick one side or the other. But insofar as ‘life’ is a meaningful predicate, and an acceptable scientific predicate at that, neither of these approaches should be seen as satisfactory for the naturalist. Why does the traditional metaphysical approach fail? As we’ll see below, the traditional approach to metaphysics is not equipped to recognize and handle the fundamentally normative character of the concept ‘life’. Traditional metaphysics (and naïve and sophisticated naturalisms, as we’ll see below), buying in, as it does, to a spectator theory of knowledge, sees itself as engaged in describing ‘life’ more accurately. But what resolution to the conflict requires is the recognition that the pragmatist metaphysician sees herself helping to *create* the concept ‘life’.³³

³¹Lewis’ Humean Supervenience is one example of such a reductionist program, leaving no room for ‘life’.

³²Mind-body dualism and panpsychism are two broad theories which take approaches in this vein.

³³The turn of phrase in the previous two sentences is borrowed from Rorty’s discussion of uniting pragmatism and feminism. The relevant sentence: “I want now to enlarge on my claim that a pragmatist feminist will see herself as helping to create women rather than attempting to describe them more accurately.” (Richard Rorty, 1998, p. 212) This sentiment applies, in my view, to philosopher’s engaged in the pragmatist project generally.

Insofar as the naïve and sophisticated naturalistic metaphysical projects canvassed in previous chapters share the spectator view of knowledge with the traditional approach, their specifically science-oriented metaphysics toolkit is equally ill-equipped to make progress in making sense of the astrobiologist's puzzle. Because they've received considerable attention already, I won't belabor the point, but it is worth outlining the general structure of their attempts.

Recall that naïve naturalized metaphysics seeks to 'read off' the correct metaphysics (in this case ontology) from our best scientific theory. But in this case, *what* scientific theory is that supposed to be? The problem of underdetermination we saw in Chapter 1 regarding special relativity and the metaphysics of time is even more starkly displayed in the present case. For not only are there similar difficulties of theory interpretation that we saw in the Putnam-Stein debate, but the space of candidate theories is more diverse and murky than before. The range of disciplines which purport to have some relevance to the concept of life ranges from geology, to population ecology, morphology and taxonomy, ethology, molecular biology and even physical chemistry. *Which* theory to 'read off' is far from clear, and any choice will reveal *antecedent* determinations of the concept **Life**.

The sophisticated approach (here in the form of structuralism) will admit that such a plurality and diversity of options leads to an unacceptable role to antecedent *choice* in our determination of the correct ontology. Instead, they seek some underlying, objective, suitably abstracted feature of the theoretical descriptions which they take to stand apart from the distorting influences of human subjectivity in interpretation. For the structuralist this will be the move to reducing the concept life to a concept of invariant nomological-mathematical structures—what Ladyman and Ross term 'real patterns'.³⁴

This reduction to stable structures results in a dilemma. Either the structuralist offers criteria for which stable structures are to be considered life, or the structuralist simply declines to pick out which stable structures count as life. If they pick the first horn, then they must offer *some*

³⁴The term itself borrowed from Daniel Dennett's paper of the same title. (Dennett, 1991)

reason to identify some particular class of structures as properly 'life'. But this just recapitulates either the traditional metaphysician's problem of reducing the concept away or adding to our ontology, or the naïve task of selecting which theory is to guide us in this choice. On the other hand, the second horn is simply another attempt at reduction (this time in terms of structure rather than matter), and will likewise fail to acknowledge the important and valuable role the concept life plays.

Thus, if we're committed to doing justice to the astrobiological project of studying the possibility of life of different kinds and in different contexts than our terrestrial one, the sophisticated and naïve metaphysician will give no better an answer than the traditional metaphysician.

LIFE IS NOT A NATURAL KIND

Instead of approaching the question as seeking necessary and sufficient conditions, and struggling when we find that there is no satisfactory list which captures the *intuitive* understanding of life, while respecting a naturalistic closure of ontology, we should recognize the normative role the term 'life' plays in discourse. Life is an achievement, and it marks significance of some sort. This much is recognized by all. But we err then we proceed to *first* identify the kind **Life**, and only then say 'Ah! That's why these things (falling under this kind) have a certain importance for us.'

This sort of response inverts the conceptual order—things are significant for various reasons, and by virtue of those reasons we come to class those things as a kind, **Life**. This makes sense of our use of kind terms, yet emphasizes the principle function of kind terms which is to organize antecedently significant objects. The task of ontology (and metaphysics generally) is to expose to inquiry and discourse what constitutes significance, how we ought to and how we do articulate such significance, and if we might want to change what is of significance for us.

4.2.2 A PRAGMATIC APPROACH

What should count as alive? What possibilities should we countenance, which should we close off? Should we open up our concept as wide as it will go, or are there values and ends which are

better served drawing the lines more narrowly? We will consider some edge cases of life, and look at how various views of importance, values, and ends might be appealed to in counting them as life or not. These considerations are not intended to be definitive or comprehensive—the goal here is to simply exhibit the texture of such a naturalist-pragmatist metaphysical inquiry.

VIRUSES

Viruses are a classic example of an edge case for living systems. They are distinctively biological in terms of their general chemical structure, as well as in terms of their general biochemistry. On the other hand viruses seem to not count as living because they fail to *do* certain things. The typical explanation you'll find is cashed out in terms of activity and autonomy:

Living organisms [are] thought to require a degree of biochemical autonomy, carrying on the metabolic activities that produce the molecules and energy needed to sustain the organism...

Viruses, however, parasitize essentially all bio-molecular aspects of life. That is, they depend on the host cell for the raw materials and energy necessary for nucleic acid synthesis, protein synthesis, processing and transport, and all other biochemical activities that allow the virus to multiply and spread. One might then conclude that even though these processes come under viral direction, viruses are simply nonliving parasites of living metabolic systems. (Villarreal, 2008)

Nevermind the awkwardness of describing viruses as “nonliving parasites”.³⁵ The focus of this passage is on autonomy, and how viruses lack that autonomy. But in what way do they lack autonomy? They cannot replicate on their own, but what organism can? Even single-celled prokaryotes need a nutrient rich environment to divide, and the interdependency of living organisms on their mutual products and bi-products is one of the central sticking points in understanding how such a system began in the first place. Rather, I imagine the sort of autonomy which is

³⁵Parasitism, to my ears at least, seems to be a *quintessential* property diagnostic of life. Perhaps this just goes to show the conceptual haziness surrounding the classification of viruses.

being appealed to here is a sort of internal activity. Single cells of organisms are able to exploit external thermodynamic gradients in order to maintain internal gradients which in turn drive protein synthesis and other functional chemistry. Viruses appear to lack this ability, ‘hijacking’ those capacities within host cells.

But *why autonomy?* Why this focus on internal activity and production as diagnostic of life? Notice that there is nothing intrinsic to the activity described obviously identifies it as essential for life. While I have neither the means nor the space to do so here, I suspect a strong case could be made tracing these notions of autonomy back to the political and social liberalism of 15th and 16th century European enlightenment philosophy. The individual, and their autonomous choice, action, and production are seen as *important* and *valuable*. To be autonomous is an achievement and valuable in its own right. Seen in this light our reluctance to attribute life to viruses is due to their passivity and parasitism. They lack the characteristics that indicate the specific value we place in life—they don’t *deserve* the categorization.

To bolster this view we might look at a related but importantly different case—the status of mitochondria. The reason that mitochondria are interesting in this respect is due to both their shared passive and ‘parasitic’ nature with viruses, yet also being biochemically similar to prokaryotes. Mitochondria are wholly dependent on their host cells, but within their host they operate in a relatively autonomous manner. They have their own DNA and reproduce independently of the host cell (though in many cases the reproduction is regulated by the host). Are mitochondria alive? I think the general sentiment is ‘no’—and this answer due to their inability to survive and reproduce outside their host.³⁶ Autonomy, or the lack of it, seems to be the common underlying factor which unites viruses and mitochondria as non-living.

But what are the alternatives? How might our reordering of values and importance change our willingness to call viruses alive? We may first note that attitudes towards the status of viruses are already plastic—the discovery of *Mimivirus*, other large viruses, and their associated

³⁶But yet again, such a criterion appears rather flimsy: tapeworms and gut flora cannot survive outside their hosts, a wide variety of organisms cannot live outside their narrowly defined habitats, and of course fish cannot live out of water!

virophages³⁷ have caused a reassessment of their classification. (see, for example Forterre, 2010) Viruses have been observed to build their own molecular factories within the infected cell's cytoplasm to synthesize the components for further virions (the viral particle which exists outside host cells). (Miller and Krijnse-Locker, 2008; Kuznetsov et al., 2013) This is in sharp contrast to the widespread picture of the virion disappearing upon infection of the cell, with only its genetic material being transmitted. While this is the case in the widely studied bacteriophages, it appears that this is not generally true. (Forterre, 2010, p. 156) What's remarkable about these 'factories' is that they are *sui generis* mechanisms, built and controlled by the virus itself. The recognition of the central role played by these factories has motivated a shift in identifying the locus of 'virus-hood' from the virion to the intracellular mechanisms which exhibit much of the behavior and activity which living cells exhibit. Forterre envisions the infected cells as the proper virus stage.

[O]ne can conclude that infected eukaryotic cells in which viral factories have taken control of the cellular machinery became viruses themselves, the viral factory being in that case the equivalent of the nucleus. By adopting this viewpoint, one should finally consider viruses as cellular organisms. They are of course a particular form of cellular organism, since they do not encode their own ribosomes and cell membranes, but borrow those from the cells in which they live. (ibid., pp. 156–157)

It's not my purpose to argue in favor or against such a view here. Indeed, we should note that this extension of the concept life to the domain of viruses is a rather conservative one. It is through the recognition that viruses *are* in fact more autonomous and self-productive than previously thought that they get included in the concept. What is interesting for us in the present context is whether and under what conditions we might want to attribute life to viruses, not by focusing on what they are, but by considering their importance and relevance to us.

For example, viruses occupy various positions of significance in our view of the world. For instance we are now coming to learn that viruses are the most abundant biological entities on earth, and their general ecological effects are only beginning to be explored. It is estimated that

³⁷Small viruses which parasitize the biomolecular mechanisms of the larger viruses.

viruses kill 20% of the ocean's microbial biomass every day. (Suttle, 2007)³⁸ Given their potentially immense influence on the ecosphere, their position and role as alive or not may well be subject to change.

More concretely, our choice to consider viruses as living could be sensitive to their role as disease agents and in gene therapy applications. As pathogens they are of relevance to us in terms of their danger, and potential to harm us. Insofar as viruses are seen as threat, it may be useful and indeed important to think of them as living. A living thing has a telos which can be subverted and an agency which can be impeded. Alternatively, insofar as viruses are seen as a tool and resource it might be more important to think of them as non-life—inert matter we can exploit. If and when we begin to use viruses in medical contexts as potential gene therapy vectors, it will perhaps become important to distance ourselves from a conception of them as alive. Notice that similar distancing can occur in cases where the status of entities under consideration is far less ambiguous. In factory-farming or cosmetic testing contexts it may be (psychologically) important for the workers to frame their work and interaction with the animals in such a way as to diminish their proximity to being 'life'. In these cases we might find such distancing morally blameworthy, precisely because the animals under consideration are *important* to us. We are not, when criticizing the practice of factory-farming, or articulating the inert, passive aspect of a gene-therapy virion-vector, describing some essence more accurately, but taking part in creating the very category of life itself, again, based on what matters *to us*.

4.2.3 JOVIAN TORNADOES

We can extend our inquiry of the bounds of life and look even further afield. Tornadoes, I suspect, would not even come to mind when most people are considering 'edge cases' for life. But are they alive? Again, if we are committed to the naturalist credo, and are committed to taking science seriously, then we should accept that this determination is not, and cannot, be made by finding

³⁸Suttle, 2007 is a fascinating and sobering reminder of the immense lacuna in our knowledge of the microbiome—especially of the oceans. Suttle estimates there are approximately 10^{30} virions present in the oceans, and 10^{23} viral infections occurring every second.

essences, by clarifying concepts, by seeing more clearly. Instead, what significance—or lack of it—do they hold for us? What role do they occupy in our own activities, in our own action and planning? My contention is that if our activities changed, if our ordering of values differed, if our ends, actions, plans involved tornadoes³⁹ in some different way, then it could become more plausible, indeed, obvious that tornadoes were alive.

To see this, let us first ask what reasons we give when explaining why tornadoes are not alive? They are not composed out of the typical bio-macromolecules, they do not appear to reproduce, they lack complex internal structure, and they are temporally and physically *ephemeral*. They are so foreign in their composition and structure from those things we easily classify as life that it becomes difficult to even think of reasons that don't seem laughably trivial. For one thing, due to their ephemeral nature, we often think of tornadoes as phenomena, not objects—*they're just clouds, or really, something clouds do*—we might say. They are as different from animals, or even virions, as something could be—they're water droplets, dust particles, and thermodynamic gradients. *How*, we might ask, *could all that add up to life?*

They could 'add up' to life—the response goes—in just the same way that hydrogen, carbon, a few other light atoms, and a thermodynamic gradient do. Contrary to the thought that tornadoes are mere phenomena, fluid vortices are distinctive self-organizing, self-individuating physical systems. Two clear examples of self-organized individuation come in the form of Bénard cells, and Taylor vortices, both thermodynamically driven fluid systems.⁴⁰ The former are the convective cells which arise in a thin layer of viscous fluid placed in a thermal gradient.⁴¹ The cells spontaneously form, often in the form of regular polygons (typically hexagons), and persist indefinitely in the continued presence of a gradient. While the viscous fluid⁴² is relatively homogeneous in molecular composition, one should not underestimate the complexity inherent in the coherent structure of the cells—from the random heat-motion of the molecules in the fluid's conductive mode, the transition to convective cells involves a remarkable structural organization. Schnei-

³⁹As will become clear below, I really have in mind here the general class of fluid vortices.

⁴⁰Bénard cells are admittedly not technically vortices, but their behavior is of relevance.

⁴¹There are a variety of excellent videos illustrating this phenomenon online, see e.g. (Anwar, 2016; Moore, 2016).

⁴²Demonstrations typically use a viscous oil, spermaceti (sperm whale) oil in Bénard's original case.

der and D. Sagan (2005) describe Bénard's original experiment and the conduction-to-convection transition:

Although the intermolecular distances of whale oil are on the order of 10^{-8} centimeter, the liquid resolved into organized structures 0.1 centimeter in size: a simple heat gradient brings some hundred billion billion (10^{20}) molecules into lockstep. They line up, show coherent motion. Such correlation among molecular trajectories and speed [is] striking. (Schneider and D. Sagan, 2005, p. 115)

One could, if one wished, describe such coherent motion as autonomous, individuated, goal-oriented activity. What goal do these cells orient towards? Simply, gradient reduction. The cells are specially suited to dissipate the thermal gradient placed across them,⁴³ and in the absence of a continued thermal input the cells disappear.

Even more 'life-like' in their properties are Taylor vortices—these are toroidal vortices set-up by a pressure gradient within the space between two concentric counter-rotating cylinders.⁴⁴ These vortices display a number of remarkable characteristics:

As the speed of the inner cylinder [changes] the flow patterns [tend] to change discontinuously and irreversibly. These transitions [occur] at certain well-defined speeds. At a single rotation rate there [can] be as many as twenty different stable states with differing numbers of vortex pairs and waves. Each stable state [is] characterized by a distinct number of vortices. And the number of vortices [depends] on the experiment's history ... [the] fluid systems of Taylor vortices have an implicit memory of their initial conditions [and] their previous states. (ibid., pp. 128–129)

These organizational features apply to vortices in a wide array of contexts, including tornadoes, which are driven by thermal gradients in the atmosphere. (see Renno, 2008) The characteristics

⁴³One should note that this behavior is a surface-layer property of the fluid itself, and not dependent on gravitational orientation as one might first expect. That is, the convection is driven by surface tension, not buoyancy. (see Schneider and D. Sagan, 2005, p. 119)

⁴⁴In standard demonstrations the outer cylinder is held fixed and the inner cylinder is rotated freely. Again, the internet is an excellent resource to see demonstrations of Taylor flow, e.g. (Ibanez, Swinney, and Rodenborn, 2014; Sahni, 2016; Schaefer, 2014).

of spontaneous self-individuation, persistence, and coherent internal structure appear on these larger terrestrial scales (hurricanes are a more convincing example of this than shorter-lived tornadoes), and as mentioned above, vortices at even larger spatial and temporal scales are well-known, for example: Jupiter's Great Red Spot.

I'm not interested in trying to argue that life is best defined in terms of local, low-entropy, gradient-reducing systems (whereas Schneider and D. Sagan (2005) explicitly *are*). Indeed, claiming as much simply recapitulates the search for an essence to life, a search I think is fruitless and misguided. Instead, I've focused on these examples in order to motivate more imaginative, alien, alternatives. For even Schneider and Sagan, in their description of Bénard cells state, as though it were obvious, that these "[c]onvection cells [are] far simpler than life". (ibid., p. 116) But again I ask, 'why?' and, 'what would have to change for our opinion to change?'

The possibility of an alien alternative I have in mind is one in which we might imagine an exo-ecology suitably foreign, but suitably familiar. By familiar here I don't mean the sort of familiarity we breed when we imagine 'silicon-based life' in analogy to 'carbon-based life', or when we speculate about a biochemistry which exploits liquid methane the way our biochemistry exploits liquid water.⁴⁵ In both of these cases we're seeking analogies to the chemo-physical makeup of cellular life. It's not obvious to me that we can learn much from focusing solely on these analogies. Rather the familiarity I suggest is phenomenological, behavioral, functional familiarity.

It has long been a topic of speculation as to whether life could exist in the atmospheres of the gas giants. C. Sagan and Salpeter (1976) suggest the possibility of ecological niches for 'sinkers', 'floaters', and 'hunters' in the Jovian atmosphere. (ibid., p. 737) But even here the central assumption is that such systems are to be modeled on the sort of living systems we are most familiar with:

In the following discussion we will consider three comparable ecological niches on Jupiter: The primary photosynthetic autotrophs, which must replicate before they are pyrolyzed, will be described as sinkers. A second category of larger organism, which

⁴⁵This latter speculation has been specifically motivated by the discovery of methane oceans on Saturn's moon Titan.

may be either autotrophs or heterotrophs but which actively maintain their pressure level, will be described as floaters. A third category of organisms actively seek out other organisms; we call these hunters ... Finally, there is a category of organisms which live almost at pyrolytic depth. They are scavengers, metabolizing the products of thermal degradation of other organisms. (C. Sagan and Salpeter, 1976, p. 747)

Instead, imagine a radically different system. Suppose we find, on Jupiter, the other gas giants, or further afield, a collected system of fluid vortices. We observe them, and discover that they are long-lasting (say, on the order of 10 years), and bear sufficient identity conditions for us to track individual vortices over time. In tracking them we notice that vortices interacting with the fluid medium of the surrounding atmosphere grow in size, and above a certain size grow unstable and split into two smaller vortices. We even notice that these vortices have certain individuating characteristics (perhaps periodic undulations in the vortex itself⁴⁶) which are preserved with variation when a vortex divides. Finally, imagine that we detect the presence of a signaling system amongst the vortices, permitting communication (as rudimentary or sophisticated as you'd like) between vortices sufficiently close.⁴⁷ We could add any number of additional features to our hypothetical Jovian vortices (perhaps the group tends to migrate to areas with steep gradients; perhaps the heritable patterns differentially affect motility, durability, signaling sensitivity; etc.).

Would this cluster of vortices—supposing such a configuration as imagined is physically possible—be alive? I don't know. But that's precisely the point—it would be, I think, a genuinely open question whether such things should be counted as life, were we to discover them somewhere. One might object on bio-essentialist grounds, that insofar as they don't share our chemistry—even in analogy, they are not to be counted as life. That's certainly a line we could take, but then it seems to make the prospects and potential for finding extraterrestrial life at once dull and remote. The reason that the question of extraterrestrial life has such a grip on the imagination is because we hold open the prospect of beings *like us*, but profoundly, and paradoxically, *absolutely alien*.

⁴⁶See Ibanez, Swinney, and Rodenborn, 2014 for a remarkably clear example of wave patterns appearing in Taylor flow.

⁴⁷Skyrms (2010), for example, outlines how signaling systems can arise from remarkably simple physical conditions.

On the other hand, should we, inspired by the above imagined scenario, start to see tornadoes, hurricanes, and the Great Red Spot as life? Surely, it would not be contrary to reason to think so. No logical contradictions arise. But perhaps such an expansion of the concept of life dilutes it in the same way that a bio-essentialism stifled it. The reason I have offered up the above scenario of the Jovian vortices is because, to my intuitions at least, such a collection of *things* is suitably alien, and suitably familiar, that there seems to me a real question.

How do resolve such a question? Perhaps we don't, or perhaps we don't until we actually are confronted by such a discovery. In the face of such a discovery, does such a curious population of vortices hold any importance, significance for us? It's not my purpose, nor do I think it's possible, to attempt to convince the reader one way or another in this question. My point is that regardless of which way we decide, it will depend, fundamentally on a *choice* made by us, a choice about how to speak, about whether to *expand* some certain concept or not.

ONTOLOGY AS ALTERNATIVES

And finally, how does this affect the field of astrobiology? How has such speculation improved the conceptual lot of those attempting to study the conditions of the possibility of life elsewhere in the universe? It is, principally, a call for conceptual liberalism and pluralism. I side with Wilson and Friedman in expressing the hope that “philosophy might once again serve as a trenchant ‘critic of concepts’ in the mode of our great philosophical forebears.” and that it might serve “as the fierce *enemy of ideological complacency*.” Wilson, 2013, p. 152, emphasis original

In the case of astrobiology, and the definition of life, does the pragmatist metaphysician have any conclusions to offer? In the spirit of my philosophical forebears, I would counsel an attitude of tolerance and an experimental spirit. Insofar as the astrobiological conversation about defining life encourages imaginative and exploratory alternatives, then nothing need be done. But insofar as persistent and pervasive assumptions about the nature of a ‘definition’ shape the discourse surrounding the study of life, it is the role of the philosopher to call attention to these assumptions *qua* assumptions. This is not to say that it is the role of the philosopher to disprove, discredit,

defeat, or undermine these assumptions—but to bring them to light, and to help recognize the potential for *different* assumptions, and the resulting differences in our aims and ends.

The role ontology plays in discourse is to underwrite a picture of our world in which we can place our values, aims, and ends. Cordoning off the concept **Life** to a particular corner of phenomena is not *descriptive* of the world, but *reflective* of those antecedent values, aims, and ends. My suggestion for a distinctly naturalist-pragmatist metaphysics, and what I've attempted to model above, is to resist the inference from *a* picture to *the* picture of the world.

The pragmatist-metaphysician's role is to explore and offer alternatives, by examining what values drive our adoption of any particular world-picture, and imagining alternative values and alternative world-pictures which might better suit us on any particular occasion. And this involves reflecting and conversing about our shared self-conception and common goals, and making a conscious *decision* about which concepts should we use, which ends we should value. With Dewey we can ask, "What would [philosophy/metaphysics] be if it ceased to deal with the problem of reality and knowledge at large?" And with Dewey we can answer,

In effect, its function would be to facilitate the fruitful interaction of our cognitive beliefs, our beliefs resting upon the most dependable methods of inquiry, with our practical beliefs about the values, the ends and purposes, that should control human action in the things of large and liberal human import. (Dewey, 1930, pp. 36–37)

In the case above this exercise is relevant to our search for extraterrestrial life. It is important to recognize the plasticity of our concept of life—and that such plasticity is a function of human activity, values, ends, and choices, not a static and eternal universe. We delude ourselves when we seek to grasp *a priori* the necessary and sufficient conditions for life anywhere in the universe. Instead of the traditional question: "*does* this count as life?" a scientific metaphysic prepares us to ask "*should* this count as life?"

4.3 Conclusion: Why ‘Metaphysics’?

The reader may, upon reflection, be left wondering whether the argument put forward here, and in its entirety throughout the previous chapters, is really simply much ado about nothing. That is, it seems like I am at pains to describe an engaged attitude toward science, one in which the philosopher has a role in criticizing concepts, and, crucially, exploring the possibility space for what new concepts we might create as our needs change. I’ve argued that such a criticism of concepts can only be licensed and be useful if it comes from an understanding, not merely of theory, but of the embodied, practical activities of building, testing, running, measuring, and ultimately—creating new ways for the practice to go on.

In addition I’ve argued that, in contradistinction to the engaged attitude defended above, what now passes as ‘naturalized’ metaphysics, fundamentally fails to take the most important and profound lessons of science to heart—lessons ultimately about the nature of inquiry, the fluidity and problematic-relative content of our concepts, our fundamental naturalness as human animals, and the illusory character of authority. I’ve argued that the core assumptions of traditional metaphysics have no place in the naturalist project, and I’ve argued that those purported naturalized metaphysics carry over precisely these core assumptions.

Why, then, do I insist on claiming the mantle of metaphysics for myself? The reason is simply that I *do* believe that metaphysics is meaningful and valuable. It is, in my view of things, the vital activity of merging our thin and meager empirical orbit, that narrowly circumscribed horizon of contact with our surroundings, into the rich and vast realm of which experience, testimony, and discourse tell stories. Metaphysics, in this sense, is story-telling, and we should care about the stories we tell ourselves and each other. What I am fighting is a sort of epistemic hubris in that story-telling, the idea that we have a good grasp on very much at all, that we can rest satisfied once we tell the one story, the final story, the story which is ‘true to the facts’. One unfortunate thing the sciences have invited, coordinate with their remarkable and profound successes, is our false sense of thinking our grasp of the world is simple and almost finished. But, in truth, we have no idea:

Physics and philosophy are at most a few thousand years old, but probably have lives of thousands of millions of years stretching away in front of them. They are only just beginning to get under way, and we are still, in Newton's words, like children playing with pebbles on the sea-shore, while the great ocean of truth rolls, unexplored, beyond our reach. (Jeans, 1942, p. 217)

Metaphysics, once rid of the idea that we aim to get our representation of the world to mirror the world, is more deeply about understanding, of knowing one's way around, of seeing what hangs together. The value of a scientific metaphysics, is not that science delivers a true and justified and static picture of a fixed and unchanging reality of which our mind mirrors, but that science, the embodied, engaged, practical activity, is epistemologically destructive. What I mean by this is that engaging with the world does less to affirm belief, or even dissolve one belief only to solidify another, and more to instill the sense of multiplicity, of complication and subtlety, of our conceptual involvement, and the contingency which rests on the questions and problems which are relevant to us. This is metaphysics as a social and cultural project, to make us more sensitive to the world, and to bring us more in touch, not with some ethereal, fixed, immutable, law-bound Nature, but with the material world of which we—and our language and concepts—are an inseparable part. Metaphysical inquiry on this view consists of holding our concepts and expectations accountable to the world and to others, and to seek alternatives, to negotiate new and different questions to ask, and new and different concepts to employ.

Appendices

Appendix A: The Geometry of Simultaneity Spaces

The core concept to appreciate is how the simultaneity space of an observer ‘tilts’ as their relative velocity becomes non-zero.⁴⁸ This is easily seen by examining some representative space-time diagrams. In figure 1 we a space-time diagram of an object α represented by the vertical solid line (α ’s world-line). For reference we have marked out an event o on the world-line and have drawn the light-cone for α at o .⁴⁹




Figure 1: A protocol for measuring events by light signal.

⁴⁸I owe an enormous debt to Robert Geroch’s book *General Relativity from A to B* 1981 for the manner of reasoning and visual depiction I employ here.

⁴⁹For expository convenience we shall speak of this world-line as *our* world-line.

We should like to know when events not on our world line (i.e. elsewhere in space) happen with respect to our local timekeeping. In order to do this we need to establish a concrete protocol for interaction with such events. Here is one such protocol—send out a light signal at some point s on our world-line. We'll start our clock at this point, and call it τ_s . We then keep track of the elapsed time, and record how much has passed when we receive the reflected signal back at our world-line at point r , with the elapsed time τ_r .

The only information about the event e we have access to is the total elapsed time from sending our signal to our reception of the reflection off e . From this we need to determine when e ‘happened’—which event on our world-line between s and r occurred at the same time. We choose to stipulate that the event e happened at the same time as the point on our world-line indexed by the time $\frac{1}{2}\tau_{r-s}$, i.e. at the half-way mark between sending the signal (s), and receiving its reflection (r).

We generalize this protocol: for every point x on our world-line there is a class of events not on the world-line such that the time index $\frac{1}{2}\tau_{r-s}$ assigned by the protocol to each event is equal to the time index of x . Such a class of events satisfies an equivalence relation which we may call the *simultaneity* relation—we will call the class of some point p its *simultaneity space*. As shown in figure 1, we can see that the simultaneity space of the event o includes e .

Now let's apply this measurement protocol to an object β moving (with constant velocity) relative to our frame. In a space-time diagram this is depicted as a straight line with some non-zero slope. In figure 2 we see that β is moving from left to right. At point s β sends out a light signal, and at some later point r receives its reflection. Just as before, β calculates the elapsed time between r and s , determining that the event e occurred simultaneous with the point on their world-line indexed by time $\frac{1}{2}\tau_{r-s}$.

Figure 2 shows the class of all such simultaneous events by the slanted dotted line passing through e . Such a simultaneity space appears rather odd compared with the example of figure 1. Indeed, if α were coincident with β at the point o , it is clear that e would, according to α , share a very different simultaneity space. Why are we justified in thinking that this is a correct




Figure 2: Our measurement protocol applied to an object β traveling with constant velocity (left to right).

representation of β 's simultaneity space? Because β is traveling with a constant velocity, the principle of equivalence applies—there is no reason not to think of β as at rest while the observer is in motion. In β 's frame the application of the protocol looks identical to that of figure 1—indeed, we can apply a ‘boost’ which transforms the situation in figure 2 such that it preserves every possible physical measurement, but is represented in a diagram identical to figure 1.

Figure 3 shows the relative ordering of the events e and e' according to two observers, α and β . We can see that α and β begin their measurement when they coincident at point s . β then receives a signal from e at point r , and calculates the point on their world-line simultaneous with e . α receives signals from e and e' both at point r' , and also calculates that they are simultaneous (as we would expect). Finally β receives a signal from e' at r'' and again calculates its simultaneity space.




Figure 3: Two observers in relative motion will not agree on the ordering of events. According to α , event e' and e occur simultaneously (in the simultaneity space A_α), while β will calculate that e' occurs (in simultaneity space B_β) significantly later than e (in C_β). (Braces highlight the $\frac{1}{2}\tau_{r^*-s}$ calculation.)

Appendix B: Extending Putnam's Argument

To see how Putnam's argument entails that the relation R is the universal relation, it is sufficient to see that two arbitrary points in space-time can be linked by intersecting simultaneity spaces. In figure 4 the point o on β is simultaneous with the point p on γ (according to our usual mea-




Figure 4: Diagram showing that the relation R holds between a point o and an arbitrary point q in o 's causal future.

surement protocol). Likewise, the point p is simultaneous with the point q on α . Allowing that R is a transitive relation on simultaneity spaces it follows that o and q satisfy such a relation, i.e., the event q is determinate at o , even though q is in o 's timelike future.

Works Cited

- Anwar, Tarique (2016). *Benard-Marangoni instability*. Youtube. URL: <https://youtu.be/v2vMXmuC818>.
- Armstrong, D. M. (1983). *What is a Law of Nature?* Cambridge University Press.
- Baker, David John (2010). "Symmetry and the Metaphysics of Physics". In: *Philosophy Compass* 5.12, pp. 1157–1166.
- Benner, Steven A. (2010). "Defining Life". In: *Astrobiology* 10.10, pp. 1021–1030. DOI: 10.1089/ast.2010.0524.
- BonJour, Laurence (1998). *In Defense of Pure Reason*. Cambridge University Press.
- Brading, Katherine, Elena Castellani, and Nicholas Teh (2017). "Symmetry and Symmetry Breaking". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2017. Metaphysics Research Lab, Stanford University.
- Brau, James E. (2017). *The Arecibo Message as a 73 by 23 character message*. URL: <http://pages.uoregon.edu/jimbrau/astr123/Notes/ch28/73by23.html>.
- Brown, Frederick S et al. (1978). "The Biology Instrument for the Viking Mars Mission". In: *Review of Scientific Instruments* 49.2, pp. 139–182.
- Compos, Daniel G. (2011). "On the Distinction Between Peirce's Abduction and Lipton's Inference to the Best Explanation". In: *Synthese* 180.3, pp. 419–442.
- Carroll, Lewis (1895). "What the Tortoise Said to Achilles". In: *Mind* 4.14, pp. 278–280.
- Chakravartty, Anjan (2013). "Naturalized Metaphysics". In: *Scientific Metaphysics*. Ed. by Don Ross, James Ladyman, and Harold Kincaid. Oxford University Press, pp. 27–50.
- Damiano, Luisa and Pier Luigi Luisi (2010). "Towards an autopoietic redefinition of life". In: *Origins of Life and Evolution of Biospheres* 40.2, pp. 145–149.
- Deamer, David W and Gail R Fleischaker (1994). *Origins of Life: The Central Concepts*. Jones & Bartlett Pub.
- Dennett, Daniel C. (1991). "Real Patterns". In: *Journal of Philosophy* 88.1, pp. 27–51.
- Dewey, John (1930). *The Quest for Certainty*. London: G. Allen & Unwin.
- Drake, Frank and Carl Sagan (1997). *The Search for Extraterrestrial Intelligence*. URL: <https://www.scientificamerican.com/article/the-search-for-extraterre/>.
- Drake, Stillman (1975). "The Role of Music in Galileo's Experiments". In: *Scientific American* 232.6, pp. 98–105.
- Earman, John and John Norton (1987). "What Price Spacetime Substantivalism? The Hole Story". In: *British Journal for the Philosophy of Science* 38.4, pp. 515–525.
- Fine, Kit (2012). "What is Metaphysics?" In: *Contemporary Aristotelian Metaphysics*. Ed. by Tuomas E. Tahko. Cambridge University Press, pp. 8–25.

- Flanagan, Owen (2012). *Buddhism and Naturalism with Guest Owen Flanagan*. Ed. by Mark Linsenmayer. Podcast Audio. URL: <https://www.partiallyexaminedlife.com/2012/03/26/episode-53-buddhism-and-naturalism-with-guest-owen-flanagan/>.
- Forrest, Barbara (2000). "Methodological Naturalism and Philosophical Naturalism". In: *Philo* 3.2, pp. 7–29.
- Forterre, Patrick (2010). "Defining life: the virus viewpoint". In: *Origins of Life and Evolution of Biospheres* 40.2, pp. 151–160.
- French, Steven (2014). *The Structure of the World: Metaphysics and Representation*. Oxford University Press.
- Galilei, G. (1914). *Dialogues Concerning Two New Sciences*. Trans. by H. Crew and A. de Salvio. Macmillan.
- (1967). *Dialogue Concerning the Two Chief World Systems, Ptolemaic and Copernican*. Trans. by Stillman Drake. 2nd ed. History of science: Univ. of California Press. University of California Press. ISBN: 9780520004504.
- Geroch, R. (1981). *General Relativity from A to B*. Phoenix Books. University of Chicago Press. ISBN: 9780226288642.
- Goodman, Nelson (1983). *Fact, Fiction, and Forecast*. Harvard University Press.
- Hall, Ned (2010). "David Lewis's Metaphysics". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Summer 2016.
- Harman, Gilbert (1965). "The Inference to the Best Explanation". In: *Philosophical Review* 74.1, pp. 88–95.
- Healey, Richard (2007). *Gauging What's Real*. Oxford University Press.
- Hinchliff, Mark (2000). "A Defense of Presentism in a Relativistic Setting". In: *Philosophy of Science* 67.3, p. 586.
- Humphreys, Paul (2013). "Speculative Ontology". In: *Scientific Metaphysics*. Ed. by Don Ross, James Ladyman, and Harold Kincaid. Oxford University Press, p. 51.
- Ibanez, Ruy, Harry L. Swinney, and Bruce Rodenborn (2014). *Taylor Couette (unstratified) Looking at different instabilities*. Youtube. URL: <https://youtu.be/ygW630nzDIg>.
- James, William (1892/1920). *Psychology: The Briefer Course*. Henry Holt and Company.
- Jean, James Hopwood (1942). *Physics and Philosophy*. Dover Publications, p. 217.
- Kornblith, Hilary (1994). "Naturalism: Both Metaphysical and Epistemological". In: *Midwest Studies in Philosophy* 19.1, pp. 39–52.
- Koyré, Alexandre (1953). "An Experiment in Measurement". In: *Proceedings of the American Philosophical Society* 97.2, pp. 222–237.
- Kuznetsov, Yuri G et al. (2013). "Morphogenesis of mimivirus and its viral factories: an atomic force microscopy study of infected cells". In: *Journal of virology* 87.20, pp. 11200–11213.
- Ladyman, James and Don Ross (2007). *Every Thing Must Go: Metaphysics Naturalized*. Oxford University Press.
- (2013). "The World in the Data". In: *Scientific Metaphysics*. Ed. by Don Ross, James Ladyman, and Harold Kincaid. New York, USA: Oxford University Press, pp. 108–150.
- Lakoff, George and Rafael E. Núñez (2000). *Where Mathematics Comes From How the Embodied Mind Brings Mathematics Into Being*. Basic Books.
- Lange, Marc (2002). *An Introduction to the Philosophy of Physics*. Blackwell, pp. 285–286.
- Lewis, David (1986a). *On the Plurality of Worlds*. Blackwell Publishers.
- (1986b). *Philosophical Papers: Volume 2*. Oxford University Press.

- Lowe, E. J. (1998). *The Possibility of Metaphysics: Substance, Identity, and Time*. Oxford University Press.
- Martínez, Daniel E (1998). “Mortality Patterns Suggest Lack of Senescence in Hydra”. In: *Experimental Gerontology* 33.3, pp. 217–225. ISSN: 0531-5565. DOI: [https://doi.org/10.1016/S0531-5565\(97\)00113-7](https://doi.org/10.1016/S0531-5565(97)00113-7). URL: <http://www.sciencedirect.com/science/article/pii/S0531556597001137>.
- Maudlin, Tim (2007). *The Metaphysics Within Physics*. Oxford University Press.
- McDowell, John (1994). *Mind and World*. Cambridge: Harvard University Press.
- Miller, Sven and Jacomine Krijnse-Locker (2008). “Modification of intracellular membrane structures for virus replication”. In: *Nature Reviews Microbiology* 6.5, pp. 363–374.
- Moore, Nick (2016). *Rayleigh–Bénard convection cells*. Youtube. URL: <https://youtu.be/gSTNxS96fRg?t=24s>.
- NASA (2017). *About NAI*. NASA. URL: <https://nai.nasa.gov/about/>.
- Plunkett, David and Timothy Sundell (2013). “Disagreement and the Semantics of Normative and Evaluative Terms”. In: *Philosophers’ Imprint* 13.
- Popa, Radu (2010). “Necessity, futility and the possibility of defining life are all embedded in its origin as a punctuated-gradualism”. In: *Origins of Life and Evolution of Biospheres* 40.2, pp. 183–190.
- Prior, Arthur N. (1970). “The Notion of the Present”. In: *Stadium Generale* 23, pp. 245–248.
- (1996). “Two Essays on Temporal Realism: Some Free Thinking About Time”. In: *Logic and Reality: Essays on the Legacy of Arthur Prior*. Ed. by B. Jack Copeland. Oxford University Press, pp. 47–51.
- Putnam, Hilary (1967). “Time and Physical Geometry”. In: *Journal of Philosophy* 64.8, pp. 240–247.
- Quine, W. V. (1953). “Two Dogmas of Empiricism”. In: *From a Logical Point of View*. Harvard University Press. Chap. 2, pp. 20–46.
- (1955). “Posits and reality”. In: *Reprinted in The Ways of Paradox and Other Essays* 2, pp. 246–254.
- (1969). “Epistemology Naturalized”. In: *Ontological Relativity and Other Essays*. New York: Columbia University Press.
- Raulin, François (2010). “Searching for an Exo-life in the Solar System”. In: *Origins of Life and Evolution of Biospheres* 40.2, pp. 191–193.
- Renno, Nilton O (2008). “A thermodynamically general theory for convective vortices”. In: *Tellus A* 60.4, pp. 688–699.
- Rickles, Dean (2016). *The philosophy of physics*. John Wiley & Sons, pp. 180–185.
- Rorty, R. (1979). *Philosophy and the Mirror of Nature*. Princeton University Press.
- Rorty, Richard (1998). “Feminism and Pragmatism”. In: *Truth and Progress: Philosophical Papers, Volume 3*. Cambridge University Press, pp. 202–227.
- Ross, Don, James Ladyman, and Harold Kincaid (2013). *Scientific Metaphysics*. Oxford University Press.
- Russell, Bertrand (1912). *The Problems of Philosophy*. Home University Library.
- Sagan, Carl and Edwin E Salpeter (1976). “Particles, environments, and possible ecologies in the Jovian atmosphere”. In: *The Astrophysical Journal Supplement Series* 32, pp. 737–755.
- Sahni, Akshita (2016). *Instability of Taylor-Couette flow in a hand-driven cell*. Youtube. URL: <https://youtu.be/6LRPv0YRf18>.

- Schaefer, Bryce (2014). *Turbulent Taylor-Couette Flow Experiment*. Youtube. URL: <https://youtu.be/BNiocOsgxW8>.
- Schneider, E.D. and D. Sagan (2005). *Into the Cool: Energy Flow, Thermodynamics, and Life*. University of Chicago Press. ISBN: 9780226739366. URL: <https://books.google.com/books?id=u4wb1EQWQD0C>.
- Schopenhauer, Arthur (1819/1966). *The World as Will and Representation*. Trans. by E. F. J. Payne. Vol. II. Dover Publications.
- Settle, Thomas B. (1961). “An Experiment in the History of Science”. In: *Science* 133.3445, pp. 19–23.
- Shamos, M.H. (2012). *Great Experiments in Physics: Firsthand Accounts from Galileo to Einstein*. Dover Publications.
- Sherman, Paul D. (1974). “Galileo and the Inclined Plane Controversy”. In: *The Physics Teacher* 12.6, pp. 343–348.
- Sider, Theodore (2001). *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford University Press.
- Sklar, L. (1992). *Philosophy of Physics*. Dimensions of philosophy series. Oxford University Press. ISBN: 9780198751380. URL: <https://books.google.com/books?id=x11i9dvn7D0C>.
- Sklar, Lawrence (1974). *Space, Time and Spacetime*. University of California Press.
- Skyrms, Brian (2010). *Signals: Evolution, Learning, and Information*. Oxford University Press.
- Soler, Léna et al. (2014a). “Introduction”. In: *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. Ed. by Léna Soler et al. Routledge, pp. 1–43.
- eds. (2014b). *Science After the Practice Turn in the Philosophy, History, and Social Studies of Science*. Routledge.
- Stein, Howard (1968). “On Einstein–Minkowski Space–Time”. In: *Journal of Philosophy* 65.1, pp. 5–23.
- Suttle, Curtis A (2007). “Marine viruses—major players in the global ecosystem”. In: *Nature Reviews Microbiology* 5.10, pp. 801–812.
- Thomasson, Amie L. (2017a). “Metaphysics and Conceptual Negotiation”. In: *Philosophical Issues* 27.1, pp. 364–382.
- (2017b). “What Can We Do, When We Do Metaphysics?” In: *The Cambridge Companion to Philosophical Methodology*. Ed. by Giuseppina D’Oro and Søren Editors Overgaard. Cambridge Companions to Philosophy. Cambridge University Press, pp. 101–121. DOI: 10.1017/9781316344118.007.
- Trifonov, Edward N (2011). “Vocabulary of definitions of life suggests a definition”. In: *Journal of Biomolecular Structure and Dynamics* 29.2, pp. 259–266.
- van Fraassen, Bas C. (2002). *The Empirical Stance*. Yale UP.
- Villarreal, Luis P. (2008). *Are Viruses Alive?* Scientific American. URL: <https://www.scientificamerican.com/article/are-viruses-alive-2004/> (visited on 12/2017).
- Walker, SI, N Packard, and GD Cody (2017). “Re-conceptualizing the origins of life”. In: *Phil. Trans. R. Soc. A*. DOI: 10.1098/rsta.2016.0337.
- Wikipedia (2017). *Sample Analysis at Mars*. Wikipedia, The Free Encyclopedia. URL: https://en.wikipedia.org/w/index.php?title=Sample_Analysis_at_Mars&oldid=807054520 (visited on 01/2018).

- Wikipedia(a) (2017). *Communication with extraterrestrial intelligence* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 25-November-2017]. URL: https://en.wikipedia.org/w/index.php?title=Communication_with_extraterrestrial_intelligence&oldid=809413025.
- Wikipedia(b) (2017). *Gauss's Pythagorean right triangle proposal* — *Wikipedia, The Free Encyclopedia*. [Online; accessed 25-November-2017]. URL: https://en.wikipedia.org/w/index.php?title=Gauss%27s_Pythagorean_right_triangle_proposal&oldid=808757745.
- Wilson, Mark (2013). "What Can Contemporary Philosophy Learn From Our 'Scientific Philosophy' Heritage?" In: *Scientific Metaphysics*. Ed. by Don Ross, James Ladyman, and Harold Kincaid. Oxford University Press, pp. 151–181.
- Wittgenstein, Ludwig (2009). *Philosophical Investigations, 4th Edition (Trans. Hacker and Schulte)*. Wiley-Blackwell.
- Wolfe, Joe (2015). "Cellular Thermodynamics: The Molecular and Macroscopic Views". In: *eLS*, pp. 1–12.
- Worrall, John (1989). "Structural Realism: The Best of Both Worlds?" In: *Dialectica* 43.1-2, pp. 99–124.