e

Applications of small-e empiricism

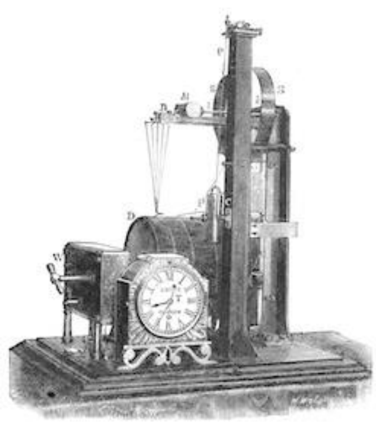

John D. Norton

2026

The cover image is of the Gray and Milne seismograph from John Milne, *Earthquakes and Other Earth Movements*. New York. D. Appleton and Company. 1886. p. 39,

Version: July 21, 2026

Preface

Preface

This volume is part of a larger project in which I seek to provide a revised version of empiricism that escapes the extremes of existing versions of empiricism and adapts empiricism to the practices of modern science. This revised version discards the inductive skepticism of a stricter “Big-E Empiricism” and allows that we can learn inductively contingent truths of science that are remote from experience. It also replaces the anthropocentric conception of experience as excitations of human sense organs with one that allows for the instrumental sensings now widespread in science. The elaboration and defense of this revised version of empiricism is found in my volume, *small-e empiricism*.

This revised version of empiricism is, in my view, benign. It preserves the best of existing empiricism in a way that I hope others with empiricist sympathies will find congenial and even automatic. While writing this account of small-e empiricism, however, it became clear to me that the consequences of this revised view for the larger literature in philosophy of science and epistemology are not benign. This volume presents those consequences. The volume can be read independently of the earlier *small-e empiricism*. I ask readers to recognize that the elaboration and defense of small-e empiricism is to be found in this earlier volume.

My practice is to acknowledge the help I have received from many during the writing of this volume in chapters to which they pertain. More general acknowledgments for assistance with the project overall is in the preface to *small-e empiricism*.

In all this, my greatest thanks remain to my partner in life, Eve. For the years of the writing of these texts, she tolerated a spouse whose head was buried in a computer screen and whose mind was in distant lands of thought, wrestling with sinister beasts. This work is dedicated to her. It is scant compensation for my inattention.

Contents

Contents

4. Inductive Fallibilism in small-e empiricism

Inductive inference is fallible. Since it is the only means we have of learning the truth of contingent propositions, the best we can affirm is that such a proposition is very well supported inductively by experience. We cannot absolutely affirm its truth.

5. Empiricism, Epistemology and the Gettier Problem

A definition of knowledge that requires truth becomes an idealization that we can never affirm applies. Discarding truth from its definition resolves the Gettier problem.

6. Gettier-Like Cases are Successes of Scientific Progress

Gettier-like cases arise commonly in the history of science. They are celebrated as great discoveries in the history of science. That their discovery did not meet the definition of knowledge as justified true belief is immaterial.

7. Small-r realism

Small-e empiricism admits a version of scientific realism that mediates between the unsustainable inductive skepticism of antirealism and the excessive inductive optimism of scientific realism. It asserts that inductive inference has the means to provide strong support for the contingent propositions of a science, but that whether any particular science has achieved that support must be left to a case-by-case analysis.

8. The Proliferation of Scientific Realisms

The present proliferation of different formulations of scientific realism extends well beyond the simplicity of the original motivation. These formulations are viable as long as they conform with small-r realism, but fail where they go beyond it.

1. Prospectus

Chapter 1

Prospectus

1. Overview

After a brief review of the doctrine of small-e empiricism, the first application of small-e empiricism to be explored concerns inductive inference itself. The distinctive feature of inductive inferences are that the go beyond what is compelled deductively by their premises. Whether any inductive inference succeeds depends essentially on further facts in the domain in which the inference is implemented. Most simply, we cannot generalize where some property applies, unless background facts in the domain sanction it as the sort of property that can be generalized. Since we can only learn of these background facts through experience, it follows that a determination of which inductive inferences apply in some domain is itself an empirical question. Accordingly, the material theory of induction treats inductive inferences empirically. Since the theory asserts that background facts determine fully which inductive inferences apply in any given domain, it treats the logic of inductive inference as something to be learned empirically from experience, domain by domain.

The next application derives from the fact that inductive inferences are fallible. At best they can tell us that some result is very likely. If an inference can assure us absolutely of its result, it has ceased to be inductive. It has become demonstrative or deductive. Since small-e empiricism asserts that inductive inferences are the only way we can learn of contingent propositions, it follows that the best we can affirm of a contingent proposition is the very great likelihood of its truth. An assurance of truth absolutely is beyond our reach.

There are many consequences of this fallibility. Definitions of knowledge that require truth become unattainable idealizations. We may believe a contingent proposition that is true, but we cannot assuredly affirm it. The safest course is to replace a concern for knowledge in our epistemology with a concern over which contingent propositions are well-supported inductively by experience. A notable consequence is that Gettier problems cease to be troublesome. Indeed, we shall see that Gettier-like cases in the history of science feature prominently in the progress of science. We also need to guard against misconstruals of inductive fallibilism. A pessimistic overreaction is a complete skepticism that refuses to accept the import of inductive inferences.

An optimistic underestimation treats the fallibilism of inductive inference as a challenge to be overcome. It leads to futile efforts to secure an unattainable level of certainty.

Misconstruals of the fallibility of inductive inference prove to be a driving force in the enduring debate between antirealism and scientific realism. Antirealists seek to misconstrue the mere possibility of error allowed by inductive fallibility into a certainty of error. They thereby conclude that all present science and even all future science must fail, even without considering the evidential support for our present sciences. Realists tend towards the opposite extreme. They seek ways to minimize or eliminate the import of inductive fallibility. Such efforts lead to excessively optimistic assessments of the truth of present theories.

A version of scientific realism—here called "small-r realism"—accommodates the fallibility of inductive inference by merely asserting that inductive inference provides the means through which contingent propositions in science can accrue very strong support. Whether any particular, present science enjoys such support must be determined on a case-by-case basis. It is not a matter to be proclaimed by a general principle. When small-r realism is specified carefully, it turns out to coincide almost exactly with small-e empiricism. The result is that the one position, small-e empiricism and small-r realism, provide a single position that mediates optimally between the traditional extremes of antirealism and realism.

Small-e empiricism maintains the traditional empiricists' opposition to non-empirical metaphysics. It is a milder opposition that avoids the extremes of earlier empiricist criticism. It does not dismiss non-empirical metaphysics as meaningless, as did the logical positivists. Rather the principal objection is that, in foregoing the guidance of experience, non-empirical metaphysics has little prospect of discerning which are the contingent truths of our world. If the contingent propositions proposed in the metaphysics lie beyond empirical scrutiny, then they are only constrained by logical consistency, which is too permissive to yield a definite result. It follows that debates among competing systems are inevitably indecisive. A second, less prominent criticism is the problem of "fatal abstraction." A founding assumption of metaphysics is that we can abstract away the particulars of a concept as it manifests in different sciences and still be left with a substantial body of universally applicable results. This assumption cannot be legislated in advance. Whether it holds must be determined on a case-by-case basis. It fails, I will argue, for the metaphysics of possibility and causation.

2. The chapters

Chapter 2 provides a brief reminder of the main claims of small-e empiricism, as developed and defended at greater length in my earlier *small-e empiricism*.

Chapter 3 responds to the fact that inductive inferences and relations of inductive support lie at the heart of small-e empiricism. Inductive inferences allow us to use experience to inform us of contingent facts beyond experience. Small-e empiricism requires that this expansion must be founded ultimately on experience. It is argued in Chapter 3 that the material theory of induction is the only account of inductive inference that allows this expansion to be founded fully on experience. For according to it, each inductive inference or relation of inductive support must be warranted by background facts in the domain of the induction. Those background facts must in turn be supported inductively. A small-e empiricist can trace their inductive support back to experience.

The troublesome alternatives are rule-based accounts of inductive inference that do not trace the warrant for their rules back to experience. In so far as that warrant fails, such accounts introduce a priori elements into the analysis. For example, it is common to assume that inductive relations of support are *always* probabilistic relations without ensuring that their applicability to a particular case is warranted by background facts. When that warrant is lacking, there is a serious risk of introducing spurious results that are artefacts of a misapplication of probability theory. What results is a violation of the core idea of empiricism, that all we learn of the world comes from experience.

Chapter 3 offers a brief summary of further details of the material theory of induction. The two accounts of the theory investigate how the material theory accommodates known but limited rules of inductive inference (in Norton, 2021); and how the individual relations of inductive support authorized by the material theory fit together in the large-scale structure of inductive support (in Norton, 2024).

Chapter 4 explores the consequences of the fallibility of inductive inference. Contingent propositions can be well-supported inductively by experience and even very strongly so supported. Since small-e empiricism asserts that inductive inferences are the only way that we can learn of these propositions, it follows that small-e empiricism can never go beyond very strong but fallible support for contingent propositions and be certain of their truth. This is not pessimism over the reach of empirical evidence since inductively very strongly supported results are quite likely true.

Chapter 5 recalls the view, widespread in epistemology, that a belief only rises to the level of knowledge if it is a belief in a truth: knowledge is justified, true belief. Insisting on truth in addition to justification in this requirement means that we define knowledge such that we can never affirm it positively of any of our beliefs in contingent propositions. Their truth is inductively opaque. In place of knowledge, it is argued, we should just consider that for which we have strong inductive support from experience. There would be no loss to our informal understanding of scientific knowledge, since we have never been able to affirm the truth of contingent propositions beyond securing strong inductive support for them.

Once we discard the demand for truth, we have resolved the celebrated challenge posed by Gettier problems to knowledge as justified, true belief. Gettier cases become benign. If there is some flaw in an apparently strong justification of a result, the worst outcome is that the result proves to be false, but we mistakenly think otherwise. The better outcome is that, nonetheless, by some unanticipated luck, the result turns out to be true after all. We are not mistaken in our belief, but only in how its truth comes about. That better case is the Gettier case.

Chapter 6 recounts how, unexpectedly to me, Gettier-like cases turn out to be pervasive in the history of science. The exact conditions of a Gettier case cannot be replicated in the history of science. In science, there is no omniscient narrator, familiar from the fictional narratives of Gettier cases, to declare what is the truth. This fictional narrator is replaced by the judgments of later stages of a science that assess which, in their view, of the results of the earlier stages are true or false. Similarly, unanticipated luck is replaced by a success of an unreliable method by unanticipated means. It turns out that these Gettier-like cases are common in the history of science and routinely consist of results that are celebrated successes. That the original discovery did not conform with an epistemologist's definition of knowledge is immaterial. What matters is the factual content of the discovery and that flaws in its justification can be repaired. It is part of the successful progress of science.

Chapter 7 develops the small-r realism sketched above. It asserts that, experience can provide strong inductive support for the contingent propositions of a science. Whether any science has achieved such support is a matter for further, case-by-case analysis. The core idea of small-r realism was already expressed in the early phases of the present tradition in the scientific realism literature from the mid twentieth century.

Chapter 8 reviews the proliferation of later formulations of scientific realism that seek to extend this original conception. It has been confounded by human concerns of the aims of science and the beliefs of scientists. A multiplicity of stratagems sought to minimize the import of inductive fallibility. They include the opaque notion of approximate truth. Selective realisms seek to identify, at a general level, which parts of a science deserve realistic construal and which do not. Those selected parts are identified as structure, entities and indispensable explanatory elements. In so far as these elements go beyond small-r realism and what inductive inference supports, they must fail. If so far as they do not go beyond it, they are superfluous additions to small-r realism.

Chapter 9 recalls how antirealism must commit to a skeptical anti-inductivism, which, it is argued, depends on misconstruals of the fallibility of inductive inference. The most important version, the pessimistic meta-induction, tries to infer from the failure of some past science to the necessary failure of all present and even future science. It neglects what is most important, the strength of present evidence for present science, in favor of the indirect and weaker evidence of past failures of earlier theories. Further, the meta-induction itself is found to fail as an inductive inference in every aspect that can be assessed.

In the final Chapter 10, small-e empiricism maintains the tradition of empiricist objections to non-empirical metaphysics. Its criticism retracts from the extremes of earlier empiricist objections, such as the logical positivists' declaration that metaphysical propositions are meaningless. It objects only that metaphysics without empirical guidance fails to be a reliable way of securing contingent truths. If the propositions sought are such as to evade empirical scrutiny, then they are constrained only by logical consistency. This constraint admits too many competing systems of propositions without providing the means to decide among them. What results are the intractable debates so characteristic of metaphysics. Examples of such debates concern the reality of possible worlds and the reality of past and future events. They are intractable because whether the predicate "real" applies in these literatures transcends empirical scrutiny.

Non-empirical constraints, such as simplicity and explanatory power, are ineffective. The material theory of induction finds that they are only epistemically potent if they are in turn founded on empirically supported background facts. The critique also rejects the proposal that non-empirical metaphysics serves to provide a framework within subsequent empirical investigations must proceed. It fails because such conceptual analysis is only useful if the concepts proposed are instantiated in the world. The range of possible concepts is so great that, as a practical matter, no purely a priori conceptual analysis can find those that are instantiated in the world without the guidance of experience. New theories, especially those in physics, are notable for requiring concepts quite unanticipated in earlier thinking.

A second, lesser-known objection to metaphysics is the problem of "fatal abstraction." It arises from the assumption in traditional metaphysics that we can abstract away all the specifics of a concept that distinguish its use in different sciences and still be left with a substantial, universal body of results. This assumption cannot be legislated in advance. It fails in the important cases of causation and possibility.

A new idea in the past few decades is "scientific metaphysics." It seeks to mimic the methods of empirical science and secure its results empirically. Much of its content is co-extensive with the empirically derived foundations of individual sciences. This use of empirical methods and results derived using them are quite welcome to small-e empiricism. What is unwelcome is the label "metaphysics." For it draws together the ill-fated project of non-empirical metaphysics with the fertile project of empirical science. The first accrues unearned legitimacy from its association with the second; and the second is compromised by its association with the first.

References

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Norton, John D. (2024) *The Large-Scale Structure of Inductive Inference.* BSPS*Open*/University of Calgary Press.

2. small-e empiricism, Briefly

Chapter 2

small-e empiricism, Briefly

1. Introduction

This chapter presents a brief synopsis of the main ideas of small-e empiricism, as elaborated in my volume, *small-e empiricism*. That volume provides many illustrations of these ideas in the history of science. In the interests of brevity, none are given here. A reader interested in them should consult this earlier volume.

2. small-e empiricism in Relation to Earlier Empiricisms

Small-e empiricism preserves the defining conception of empiricism. It is that experience provides our only means of learning the truth of contingent propositions, such as form the claims of our sciences.

2.1 Inductive Skepticism Renounced

Big-E Empiricism, as I call it, adds to the defining conception of empiricism that the content of experience itself is all that we learn from experience. They are the only facts to which we have secure access. We cannot attribute a factual status to the results of a science that go beyond experience. A slightly softer version holds that all we learn from experience is the content of that experience itself and the possibility of factual content in further experiences of a similar character. The role in science of results beyond direct experience is purely instrumental. They serve, for example, to provide convenient summaries of our experiences and to mediate inferences from our experiences to further possible experiences.

Small-e empiricism discards this inductive skepticism in favor of the optimistic view that we can learn from experience contingent facts that are themselves remote from experience.

2.2 Privilege of Human Sense Organs Renounced

Traditional empiricism, as it emerged in the seventeenth and eighteenth centuries, characterized experience as the excitation of human sense organs: vision, hearing, smell, taste and touch. The subsequent processing of that experience was most commonly understood as undertaken by automatic mental processing.

Small-e empiricism allows for experiences that manifest in the excitations of human sense organs. It extends the concept to include purely instrumental sensings, such as are now commonly employed throughout the sciences. The habits of mind of a cognizer are an unreliable means of discerning the import of experience. There is no guarantee that they preserve or extract truths from experience. These habits are replaced in small-e empiricism by logical inferences, as defined by an applicable logic. Deductive inferences are those that preserve the truth of their premises. Inductive inferences are those that can preserve its likelihood.

3. Experience Reconceived

These revisions to empiricism require three modifications to our conception of experience.

3.1 The Propositional Representation of Experience

First, traditional empiricism undertook investigations that are now better treated by branches of newer sciences. The precise operation of our human sense organs is better treated by empirical physiology. The details of the habits of the mind and it processing of ideas are both better left to empirical psychology. An empiricism adapted to the practices of modern science should engage with experience as it manifests within scientific work. That manifestation is in the form of propositions in scientific writing that are declared to be reports of experience. These propositional reports of experience have been formulated so as to shed idiosyncrasies of individual human sensings and aberrations of human mental processing. They are presented as universal results that should be replicable by any competent observer.

3.2 Experience Includes Instrumental Sensings

Second, instrumental sensings expand the repertoire of what counts as experience beyond the excitations of human sense organs. Here small-e empiricism follows the near universal trend over the past century and more in the sciences of replace human sensings by instrumental sensings. Astronomers, for example, have long ceased to make primary observations by gazing into the eyepiece of an optical telescope. For over a century, their gaze and sketches of what they saw have been replaced by photographic captures of the telescopes' findings. Such recordings provide an enduring record. Human sense organs play no role in the processes that create them.

Similarly, it is now increasingly common for the processing of these instrumentally sensed results to be conducted automatically by algorithmic computations, so that human mental process are further removed. Such algorithmic processing has become a necessity in many modern sciences. The massive volume of instrumentally collected sensings, especially in astronomy and physics, preclude human processing.

3.3 Experience as a Process

Third, experience is not characterized as a localized sensing event. Experience is conceived as a physical process that connects us with the system of interest. The processes considered are familiar. The propagation of electromagnetic waves includes propagations in the entire electromagnetic spectrum, from radio waves and microwaves, through the visible spectrum, to X-rays and gamma rays. Other propagations include sound propagation, such as the propagating noise of the 1883 Krakatoa volcanic explosion, and the propagation of seismic waves in the Earth after seismic events. These processes also include the propagation of substances. The simplest case is just our sense of smell identifying a distant, odor-producing source. Other examples include diffusion of samples in chromatographs and the material propagation of particles in cosmic rays from distant astronomical creation processes.

The various stages of these experiential processes admit propositional descriptions. The direction of physical propagation is from the system of interest towards us. Inductive inferences proceed in the opposite direction. Propositions describing stages of the process closer to us provide inductive support for propositions describing stages more remote from us, closer to the system of interest.

4. The Terminal Obsession Escaped

The reconception of experience as a process resolves one of the most vexing problems facing Big-E Empiricism. That view is founded on the idea that we can distinguish propositions representing experience from so-called "theoretical" propositions that pertain to conditions more remote from experience. Since the distinctively experiential propositions would be associated with the terminus of small-e empiricism's experiential processes, the enduring but long-failing efforts to provide a precise characterization of them is what I call the "terminal obsession."

The terminal obsession depends on the assumption that being designated as experiential is a monadic property of a proposition. Either it is or is not experiential. Once we reconceive experience as a process, this monadic property is replaced by a binary relation. Stages of the experiential process and propositions describing them are closer to or farther from the system of interest. There is no absolute division between the experiential and the non-experiential. Rather, propositions representing experience in scientific writing describe a conveniently chosen stage of the experiential process that is suitably close to us.

The stage reported is provisional. Doubt may be cast onto the propositional report through suspicions over the inferences to it from stages of the experiential process closer to us. It may be, for example, that these prior inferences introduced theoretical assumptions whose security is suspect. The straightforward remedy is to revert or "wind back" to a stage of the experiential process closer to us. This procedure of winding back is used routinely in science. It enables a practical escape from the exaggerated perils of what is called, in the philosophy of science literature, the "theory ladenness of observation."

5. The Doctrine

The formulation of small-e empiricism presumes this revised conception of experience and is:

> *Science is empirical.* The propositions of experience provide inductive support for the truth of the contingent propositions of science; and the inductive support of experience is the only means of providing support for these contingent propositions.

This formulation identifies inductive inference as the mode of inference from experience. In so far as the contingent propositions of a science have a scope greater than the confines if their supporting experience, these propositions are logically stronger than the experiences that support them. It follows that the support cannot be deductive, but must be inductive.

In principle, any logic of inductive inference can be employed. However, as will be reported in Chapter 3, the material theory of induction is uniquely well adapted to small-e empiricism.

The doctrine has two parts. The first asserts a positive: the capacity of inductive inference from experience to provide support for the truths of contingent propositions in science that are themselves remote from experience. What this part does not assert is that any particular science has in fact secured strong support from experience for its central claims. All the part asserts is that inductive inferences from experience have the capacity to provide such support. Whether it has been realized in any particular science must be established on a case-by-case basis.

The second part of the doctrine asserts a negative: these contingent proposition can only secure support from experience. This negative applies on the largest scale. We may have results remote from experience that are deduced from other results also remote from experience. For example, Kepler's laws of planetary motion are deduced from Newton's system of the world. Small-e empiricism asserts that a fuller examination of the support for these propositions collectively will find that support to reside ultimately in experience.

6. Arguments for small-e empiricism

6.1 For the Positive Claim

The immediate and most obvious argument for the positive claim is the extensive record of success of empiricism in the sciences of past millennia. Time and again experience has provided us the foundation for novel science. Empirical evidence has great power. It can force us to accept results that may initially discomfit us greatly. This was the case in the early twentieth century, when empirical results, such as concerning atomic spectra, forced an acceptance of quantum theory.

That experience has this power follows from the reconception of experience as a physical process that connects us continuously with the system of interest. For we learn of the system in the most direct way, by connecting with it. The authority of experience resides in that connection.

Finally, since small-e empiricism is inductively optimistic, it allows for strong support for the real existence of the processes comprising experience. That is, empirical investigations can affirm that there really are light and sound waves that inform us of distant systems of interest. In short, small-e empiricism supports and is supported by an ontology of experiential processes.

6.2 The Weakness of Big-E Empiricism

Big-E Empiricism is, from the perspective of small-e empiricism, an extreme, unsustainable doctrine. In additions to the difficulties indicated above, we can see that the principal arguments for small-e empiricism do not extend well to Big-E Empiricism. Support for the historical success of empirical investigations arises for both Big-E and small-e empiricism. However, because of its inductive skepticism, Big–E Empiricism cannot locate the authority of experience in the directness of the physical connection to the system of interest. For Big-E Empiricism treats these empirical processes as purely instrumental inventions. It similarly reserves judgment on whether the systems of interest are real, as opposed to another part of an instrumental scheme for relating experiences to possible experiences.

Big-E Empiricism asks us to accept the epistemic privilege of some states of excitation of human sense organs. It is true that these excitations have been fertile historically in allowing us to infer to other possible experiences. What remains unanswered why these purely mental events have this power. It cannot be explained by asserting, as does small-e empiricism, that processes like light and sound really inform us of facts in the world well beyond experience. For Big-E Empiricism is distinctive in refusing to accept their factual character. Big-E Empiricism cannot provide a foundation for the authority of experience comparable to that offered by small-e empiricism.

6.3 For the Negative Claim

Arguing for the negative claim, that experience alone can provide support for the contingent results of experience, presents a greater challenge. It is not just that we must address other known modes of investigation. We must also address unknown modes. How can we preclude the efficacy of some other mode that is undiscovered and may even elude the imaginations of our greatest writers of fiction?

We can, it turns out, have some confidence that no such novel mode is nearby. Over two millennia of the most energetic investigations into nature have failed to provide a serviceable alternative; and the reason for this failure has general scope.

To begin, other modes have been considered and used. They can be categorized as Platonic insight; innate or intuitive ideas; and oracular revelation. Their universal failure has been recounted in Chapter 12 of *small-e empiricism* with many illustrations. That they fail is not immediately apparent from a record of their history. All these other modes claim successes. A closer examination of them, from Plato to Einstein, shows that these successes are illusory. These other modes can recover results *retrospectively*, once the results have already been found by empirical methods. There is however a universal failure *prospectively* of these modes to introduce novel results independently that survive empirical scrutiny. They are only ever wise after the fact.

A stark illustration of the comparative weakness of these other modes is given when, as is often the case, empirical investigations conflict with one of their certainties. Such conflict is inevitably resolved in favor of experience. Empirical results remain the ultimate authority.

Both empirical investigations and these other modes are fallible and can be corrected. However, the methods of empirical investigations are distinctive in containing within themselves the means to effect the correction. Empirical investigations of the experiential processes can correct mistaken appraisals of details of these processes. The more dramatic reversals of empirical science come about when new evidence arises from experience in novel domains. The reversals proceed, as empiricism requires, by accepting to the import of the evidence of experience. These other modes have no comparable means to correct their errors. They must just say that Platonic vision was clouded or innate of intuitive thoughts were somehow misapprehended.

That there are such differences between empirical and other methods can be accounted for by small-e empiricism itself. Empirical investigations provide us a rich ontology of the processes that can inform us of the world. Those that can so inform us have been exploited in small-e empiricism. The ontology is notable for lack of any processes that would allow for the successful operation of Platonic vision, or of oracles, or of innate or intuitive ideas, unless these last have in some way already been informed indirectly by experience.

What of an as yet known mode? How can it inform us of a system of interest if the mode does draw on a physical connection with that system; or draw on such a connection with other systems from which we can infer to the condition of the system of interest? Any mode that proceeds this way is, by definition, within the scope of small-e empiricism.

7. Conclusion

Small-e empiricism provides us with a portrait of how experience can inform us of the world. It tells us the extent of its reach and the limits of its reach. The remainder of this volume will consider how accommodating these capacities affects existing analyses in philosophy of science and epistemology.

3. The Material Theory of Induction

Chapter 3

The Material Theory of Induction

1. Introduction

Inductive inference and relations of inductive support play a central role in small-e empiricism. In working with small-e empiricism, one may choose any serviceable account of inductive inference. There is, however, one account that is especially well adapted to small-e empiricism. It is the material theory of induction, as elaborated in my two works, Norton (2021, 2024). It is especially well adapted to small-e empiricism because the theory requires that inductive inferences are themselves to be warranted by background facts; and an empiricist will then seek support for those background facts in experience.

The basis of this unique adaptation is outlined in more detail in Section 2. The main claims of the material theory of induction are summarized in Section 3. Section 4 provides a brief inventory of applications. The concluding Section 5 recalls briefly the difficulty of using universal rules of induction.

2. Empiricism and the Material Theory

The essential function of any account of inductive inference is to amplify the scope of propositions we have already secured evidentially. Such a theory may allow us to infer from premises of lesser generality to those of greater generality. Or the theory may just tell us that the evidence of such and such propositions provides support, weak or strong, for some proposition of greater scope. However it comes about and whatever form it takes, an account of inductive inference may only bear the name if it authorizes some expansion of what we learn of contingent facts of the world.

In serving in this capacity, an account of inductive inference falls within the scope of empiricism. For an inductive logic allows us to expand what we learn from experience. Empiricism requires that this expansion must itself derive from experience. The expansion cannot rely on an a priori principle for then it contradicts the basic doctrine of empiricism.

For example, it is common in practical science to use a condition of simplicity to expand empirical results. We may have a finite set of data points that represent what we have found empirically in some finite set of cases. We expand to infinitely many cases by plotting a curve through the points, when they are drawn on graph paper. Of all possible curves, it is routine to justify the particular one selected as the simplest. We may imagine that this justification relies on some a priori metaphysics of simplicity. We might then take it as some sort of deep metaphysical truth that the world is simple; and this truth allows the expansion from a finite set of cases to an infinite set.

If this justification were to succeed, it would violate the doctrine of empiricism. It would require us to accept that our empirical explorations need to be supplemented by a special fact about the cosmos not itself derived from experience. Closer examination in Norton (2021, Ch. 6) has sought to show that there is no such additional fact of simplicity, in so far as it is to be conceived as a cosmic fact not itself inferred from experience. The crucial decision of which curve to fit is actually determined by background facts and will differ from case to case. Attempts to formulate a single universally applicable notion of simplicity, applicable in all cases, either fail empirically or are so qualified to protect them from falsification that they become vacuous.

Curve fitting problems should be treated on a case by case basis, according to the specific background facts of the problem at hand. In many cases, fitting a linear or near linear curve is indicated by what we already know loosely of the system under investigation. For example, in exploring the relationship of the pressure of a gas with its density under ordinary conditions, we antecedently expect some proportionality in the two quantities and that a linear relationship is likely to be a pretty good representation of it over a small range of parameter values. We would be prepared to admit small higher order corrections if the data demanded it. That allows us to explore step-wise up a hierarchy of higher order polynomials—quadratics, cubics, etc.—if they secure a better fit without obvious overfitting.

Would we ever *start* with an infinite order polynomial? In the abstract, one may think not. However, we would so start if the process is known to be cyclic in time, such as planetary orbits or ocean tides. For then we would use a combination of trigonometric functions of time, such as sines and cosines; and they are infinite order polynomials. We might start with a single sine or cosine and then add further sines or cosines to secure a better fit.

The invocation of simplicity is a convenient way to circumventing the need to explain in detail why a particular choice of curve is deemed appropriate. It is tedious in familiar cases to explain why we think it plausible that a linear relation is a good starting point. It may be especially so if our reasons are only vaguely formulated. Even so, our choice of curves must be governed by what we already know factually of the system in question. Those background facts are essential if any curve or hierarchy of curves is to be justified. Without them, all we have is a collection of dots on a page of graph paper. We can have no idea of which curve should connect them or even if a curve should connect them at all. In sum, to require simplicity in curve fitting is a convenient surrogate for the import of background facts that govern the case at hand.

The basic assertion of the material theory of induction is that all successful inductive inferences conform with this pattern. It may appear that that the inductive inference or judgment of inductive support is justified by some general rule of inductive inference. In so far as the induction is successful, the general rule can only be an intermediate whose real inductive power derives from background facts. In this way, according to the material theory of induction, any inductive expansion of experience is in turn warranted by further facts; and those warranting facts are themselves to be secured from experience.

Thus the material theory of induction provides an account in which the totality of what we learn from experience is grounded in the facts of experience themselves. In so doing, it conforms with the requirements of small-e empiricism in the most thorough-going manner.

3. Main Claims of the Material Theory of Induction

Since the two volumes, *The Material Theory of Induction* (2021) and *The Large-Scale Structure of Inductive Inference* (2024), provide a detailed account of the material theory, all that is needed here is a brief inventory of the main claims. The principal claims of the material theory as listed in (2021, Ch. 1) are:

Inductive inferences are warranted by facts not by formal schema.

All induction is local. It is contextual.

Inductive inference is generically variegated and imprecise.

Inductive risk is assessed and controlled by factual investigation.

Inductive inference is material at all levels.

When the material theory is applied to the overall structure of relations of inductive support in a science or science in general, (2024, Ch. 2) lists four principal claims:

1. Relations of inductive support have a nonhierarchical structure.

2. Hypotheses, initially without known support, are used to erect nonhierarchical structures.
3. Locally deductive relations of support can be combined to produce an inductive totality.
4. There are self-supporting inductive structures.

Subsequent chapters use these four claims to support further results:

The inevitable circularities in the relations of inductive support in a mature science are benign. (Ch. 3)

The dynamics of inductive inference allow a rich body evidence to single out a mature scientific theory as the uniquely best supported. (Ch. 4)

The traditional problem of induction is dissolved by the material theory, since the theory does not employ universal rules of inductive inference. (Ch. 6)

4. Some Applications

The material theory of induction is offered as the account of all successful inductive inference. As a result, a primary burden of the theory is to demonstrate that it is adequate to existing inductive practices, in so far as they are cogent. In particular, this means that any universal rule of induction in the present literature must have a basis in background facts, in so far as that rule can be used successfully. A major task of *The Material Theory of Induction* was to meet this challenge. It addressed a series of familiar rules of inductive inference and showed how the successes of each were derived from background facts.

Enumerative induction is addressed in Chapter 1. It argues that there is no general license that allows us to infer from particulars to their generalizations. All such inferences, if cogent, are authorized by background facts in the domain of application; and these background facts specify just which aspects of the particulars may be generalized in each domain. Chapter 3 addresses the "gold standard" of empirical science, the replicability of results. It argues that there is no general principle of replicability. There are exceptions in all combinations to the import of successful and failed replication. The individual successes and failures in specific cases, however, can be accounted materially by specific background facts.

Chapter 4 argues that efforts to provide a general schema of analogical inference leads to a regress to ever more complicated rules, as the schema seek to accommodate each new, anomalous example. Analogical inferences are not warranted by such formal rules, but by facts of analogy, where these facts are specific to each case. The general idea of the epistemic potency of what are called “epistemic virtues” or “epistemic values” is sustainable, it is argued in Chapter 5, only in so far as these values are not taken as freely chosen ends in themselves. Rather they can only succeed when they are criteria whose epistemic potency derives from background facts. Section 2 above recalled the account given in (2021, Ch. 6-7) of the example of simplicity. It has epistemic import only in so far as it is a surrogate for specific background facts in each domain of application.

Abduction or inference to the best explanation presents special problems for analysis in Chapters 8-9, since there is no sufficiently well-articulated account of explanation that can sustain the epistemic function abduction requires of it. A careful examination of many standard examples of inference to best explanation finds that the inductive successes in these cases can be recovered from a simple model which assigns no special epistemic powers to explanation. Rather the successes derive from the adequacy to the evidence of the favored hypothesis; and the failure of a foil. The latter fails either by being contradicted by the evidence or by accruing an unsustainable inductive debt of unsupported suppositions.

The concluding Chapters 10-15 address probabilistic accounts of inductive inference, such as a Bayesian account. The material theory does authorize probabilistic inductive inferences and probabilistic inductive support. It should do so, since such probabilistic accounts have been immensely successful over a wide range of domains. The material theory does not support the common idea that *any* indefiniteness or uncertainty can be well represented by a probability measure. Instead, the material theory requires that a probability measure can be employed in any particular inductive problems *only if background facts peculiar to the problem warrant the use of the probability measure*. In physics, a physical theory may authorize a probability measure. In social science applications, probabilities may be introduced by the assumption that some individual has been randomly sampled from a population; or that the individual may be treated as such. If this authorization is missing, the analysis is at serious risk of inductive fallacies, whose spurious results are an artefact of a misapplied probability measure.

This appraisal of probabilistic induction is supported by noting that all arguments for probabilism are necessarily circular. They are deductive inferences whose premises must be logically at least as strong as the conclusion of probabilism, so that the assumption of probabilism must be included already in the premises. To illustrate how other calculi of inductive inference are possible, several chapters describe cases in which the background facts authorize a non-probabilistic inductive logic. They include systems whose chance properties are akin to an infinite lottery machine (Ch. 13), non-probabilistic cases of indeterminism in physics (Ch. 15) and some quantum systems (Ch. 16).

5. Conclusion

Inductive inference and relations of inductive support lie at the heart of small-e empiricism. The advantage of treating them materially is that the material theory of induction is uniquely adapted to the core commitment of empiricism, that we only learn of the world through experience of it. The two combine to form a well-integrated system.

The risk to small-e empiricism of employing a rule-based account of inductive inference is that a rule, employed in a context in which it is not warranted, introduces factual content without proper empirical support. For example, the rule of enumerative induction authorizes us to infer from the proposition that some *A*'s are *B*, to the proposition that all *A*'s are *B*. The well-known difficulty is that the rule can only be used successfully for very few of the possible *B*'s in each circumstance. If we apply it with *B*'s outside these very few, we arrive at spurious generalizations about *A*'s.

Enumerative induction is a simple and transparent instance of this failure. It can happen with any rule-based account, if the rule is applied inappropriately, that is, if it is applied without proper material warrant in background facts. For example, if we use probabilities to represent complete neutrality of evidential support, the "disjunctive-inductive" fallacy follows, as described in Norton (2010). Assume that our evidence is completely neutral over an exhaustive, mutually exclusive set of outcomes, $O_1, \ldots, O_{100}$. The evidence provides nothing of use in discriminating among them. If we represent that neutrality by a roughly uniform probability distribution over all one hundred outcomes, we would then find that the probability is very high that the outcome that obtains is one of the first ninety-nine. That spurious conclusion is an artefact of the misapplication of probabilities. The appropriate assessment should just be the assumption made at the outset, that the evidence gives us no basis for any conclusion over the comparative likelihood of the outcomes, beyond the original complete indefiniteness.

References

Norton, John D. (2010) "Cosmic Confusions: Not Supporting versus Supporting Not-." *Philosophy of Science*, **77**, pp. 501-23.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Norton, John D. (2024) *The Large-Scale Structure of Inductive Inference.* BSPS*Open*/University of Calgary Press.

4. Inductive Fallibilism in small-e empiricism

Chapter 4

Inductive Fallibilism in small-e empiricism

1. Introduction

Inductive fallibilism asserts that inductive inferences are always fallible. No matter how strong the inductive support for some contingent proposition, there is always the possibility that the proposition is false, even if its falsity is extremely unlikely. According to small-e empiricism, the only way we can learn the truth of contingent propositions is by their inductive support from experience. If we combine inductive fallibilism with small-e empiricism, it follows that we can never secure inerrantly the truth of a contingent proposition. In the best case, we can be assured only that its falsity is very unlikely.

The goal of this chapter is to articulate inductive fallibilism in the context of small-e empiricism, its consequences and how they have been neglected or misread. Section 2 will review how inductive fallibilism arises from the nature of inductive inference and that its import in small-e empiricism is that we can never be assured absolutely of the truth of a contingent proposition. Section 3 will recall how, in the nineteenth and twentieth centuries, the fallibility of inductive inferences came to be accepted as an inescapable constraint on science. Section 4 will argue that inductive fallibility is no basis for inductive skepticism. Fallibility is not inductive failure. Section 5 argues that inductive fallibility is a routine part of the familiar dynamic of science, as it progresses through revolutionary changes. Among the totality of our best supported sciences, it is inevitable that somewhere the unlikely failure will eventually happen and a significant revision will be required. Finally, Section 6 argues that a consequence of the opacity of truth is that a serviceable definition of scientific knowledge requires only that its contingent results have very strong inductive support from experience. Their truth cannot be demanded.

2. Inductive Fallibilism in small-e empiricism

This chapter will develop the consequences of combining two results. The first arises from the nature of inductive inference. The second in a basic supposition of small-e empiricism.

First, that inductive inference is fallible follows directly from how inferences are divided into the deductive and the inductive, that is, the demonstrative and the non-demonstrative.

Deductive or demonstrative inferences are, by definition truth, preserving. If the premises are true, then the conclusion must also be true. Inductive inferences are non-demonstrative, by definition. They are ampliative in the sense that they support propositions that extend beyond their premises. The price for this gain is that the truth of the premises does not assure the truth the conclusion. The likelihood that the conclusion is true may merely be enhanced, even if very slightly. In the better cases, the truth of the premises may render the truth of the conclusion very likely. If the induction is expressed in relations of inductive support, such as conditional probabilities, the strongest relation of support must still leave open the possibility that the result supported is false. It follows that we can never be assured absolutely of the truth of a result on the basis of inductive support.

Second, the core doctrine of small-e empiricism, as developed in Chapter 2 above, asserts that the inductive support of experience is the *only* means of providing support for the truth of the contingent propositions of science. If we allow that this restriction applies to contingent truths in general, it now follows that we can never be assured absolutely of the truth of a contingent proposition.

We should not overestimate the peril that comes with this lack of absolute assurance. In familiar, established sciences, we have levels of inductive support that are so strong that, for all practical purposes, they can and should be taken as assuring us of truth. However, a careful and proper statement of the level of security of even the best inductively supported results is that there is always a possibility of error, even if extremely unlikely in any particular case.

3. The Rise of Fallibilism in Philosophy of Science

In his historical survey of demarcation criteria for science, Laudan (1983) identified a long-standing tradition of infallibilism prior to the nineteenth century. Good science had, to use his Kantian-inspired expression, "apodictic certainty." That is, its results were derived from foundations so secure as to be beyond dispute. This was the ideal found in the writing of Aristotle, Bacon, Newton and others. That the ideal could be achieved in practice was not so universally accepted. Hume's celebrated critique of inductive inference was an extreme expression of this reserve.

It is in the nineteenth century, according to Laudan's synopsis, that the unattainability of this ideal was advanced more widely as a fundamental limit on science. The term "fallibilism" was introduced for it by Charles Sanders Peirce in then unpublished writings. Here is one synoptic formulation (1960, p. 60):[1]

> … On the whole, then, we cannot in any way reach perfect certitude nor exactitude. We never can be absolutely sure of anything, nor can we with any probability ascertain the exact value of any measure or general ratio.
>
> This is my conclusion, after many years study of the logic of science; and it is the conclusion which others, of very different cast of mind, have come to, likewise.

Shortly after in this document, Peirce gave one of many characterizations of what he had called the "doctrine of fallibilism" (1960, p. 61). He wrote:

> Nor does fallibilism say that men cannot attain a sure knowledge of the creations of their own minds. It neither affirms nor denies that. It only says that people cannot attain absolute certainty concerning questions of fact.

Although Peirce is now generally credited with introducing fallibilism into philosophical writing, he was not the first nineteenth century writer to do so. Earlier, W. Stanley Jevons had advanced similar themes in his analysis of the "principles of science." His text (1874) promoted the use of hypotheses and the inverse method of probability, which we would now call "Bayesian." He wrote: (1874, p. 416)

> … throughout inductive logic, we deal only with probabilities. There is no infallible mode of arriving at the absolute truth, which lies beyond the reach of human intellect, and can only be the distant object of our long continued and painful approximations.

Chapter 6 of *small-e empiricism* reported that Jevons had some discomfort with what he characterized as empirical science when it was compared to his preferred deductive science. However, this preference did not overrule his fallibilism. He wrote of the relation between the two forms of science: (1874, p. 166)

> The second great portion of scientific knowledge consists of facts which have been first learnt in a purely empirical manner, but have afterwards been shown to follow from some law of nature, that is, from some highly probable hypothesis.

Even the more secure science rested on foundations with only probabilistic security.

[1] Identified by the editor as from an untitled manuscript or manuscripts of c. 1897.

In the twentieth century, this probabilistic formulation of fallibilism came to be widely accepted. Reichenbach noted the new view: (1938, p. vi)

> The idea that knowledge is an approximative system which will never become "true" has been acknowledged by almost all writers of the empiricist group; but never have the logical consequences of this idea been sufficiently realized. The approximative character of science has been considered as a necessary evil, unavoidable for all practical knowledge, but not to be counted among the essential features of knowledge; the probability element in science was taken as a provisional feature, appearing in scientific investigation as long as it is on the path of discovery but disappearing in knowledge as a definitive system.

The earlier view was then dismissed as fictional:

> Thus a fictive definitive system of knowledge was made the basis of epistemological inquiry…

My impression is that fallibilism is widely accepted in both philosophy of science and philosophy in the abstract, as a general matter of principle. Yet, as we shall see in the chapter to follow, in the concrete, its specific implications are not always recognized.

4. Fallibilism is not Failure-ism

4.1 Fear of Failure

Inductive fallibilism allows for the possibility that even results strongly supported inductively can turn out to be false. The mere fact of that possibility is no basis for skepticism about inductive support. The easy fallacy is to infer from the possibility of falsity to the presumption of a complete absence of support. We shall see in Chapter 9 below that something like this fallacy lurks behind some antirealist sentiments.

Inductive fallibilism allows for a range of strengths of inductive support. There will be many results whose inductive support is weak. It is prudent to suspend judgment on them. There will be many others that enjoy very strong inductive support. That just means that they are unlikely to be false. If we follow Hume's maxim from Chapter 12 of *small-e empiricism* and proportion belief to the evidence, it would be *im*prudent to suspend judgment.

There is nothing especially perilous in following Hume's maxim. It is quite familiar from ordinary life. A 20-year-old can, on average, expect to live a further 55.63 years.[2] That is 20,319 days. As time passes, each day, that person will make an inductive inference to the conclusion that this day is not their last. The inference will, on average, yield a correct conclusion 20,318 times and an incorrect result once. We can make many inferences at this level of certainty without any great peril.

Eventually, if we make enough fallible inferences, we can expect a failure. As a surrogate for these many inferences by one person, consider a population of 810 20-year-olds. Each has a probability of death in one year of $0.001235 \approx 1/810$. If everyone of this population of 810 20-year olds infers to their survival each day for a year, we have a total of 810x365=295,650 inferences. On a rough average, only one of them will infer to a falsehood.

We are similar position with the inductively strongly supported results of our sciences. It is possible and even expected, once we consider enough such results, that one will fail. However, overwhelming, most will not and it would be folly to proceed as if all were about to fail.

4.2 Feynman on Flying Saucers

That inductive fallibility is not fatal to science is widely recognized in the sciences, but most commonly only tacitly. Here is an instance of its explicit recognition by Richard Feynman, a Nobel laureate and founding figure in twentieth century quantum theory. (Feynman, 1965, pp. 165-66, his emphasis)[3]

> Some years ago I had a conversation with a layman about flying saucers—because I am scientific I know all about flying saucers! I said 'I don't think there are flying saucers'. So my antagonist said, 'Is it impossible that there are flying saucers? Can you prove that it's impossible?' 'No', I said, 'I can't prove it's impossible. It's just very unlikely'. At that he said, 'You are very unscientific. If you can't prove it impossible then how can you say that it's unlikely?' But that is the way that *is* scientific. It is scientific only to say what is more likely and what less likely, and not to be proving all the time the possible and impossible.

[2] Statistic for males for 2022 from US Social Security Administration. https://www.ssa.gov/oact/STATS/table4c6.html

[3] For video of Feynman reading this text see https://www.openculture.com/2011/08/richard_feynman_the_likelihood_of_flying_saucers.html

> To define what I mean, I might have said to him, 'Listen, I mean that from my knowledge of the world that I see around me, I think that it is much more likely that the reports of flying saucers are the results of the known irrational characteristics of terrestrial intelligence than of the unknown rational efforts of extra-terrestrial intelligence.' It is just more likely. That is all. It is a good guess. And we always try to guess the most likely explanation, keeping in the back of the mind the fact that if it does not work we must discuss the other possibilities.

4.3 Bridgman's Fallacy

We might well wonder if the fallacy noted above could arise among serious scientists. Might they mistakenly infer from the possibility of falsity to the presumption of a complete absence of support? There is, it turns out, a striking example of it from the pen of a leading scientist and it was published in the leading journal in philosophy of science. Here is the founder of the modern doctrine of operationism and another Nobel laureate, Percy Bridgman, writing on "Impertinent Reflections on History of Science" in the journal *Philosophy of Science*. He wrote: (1950, p. 63)

> The thesis that the past has an objectively real existence is devastatingly shattered by the consideration that there is no way whatever by which it can be proved that the entire universe was not created five minutes ago. For if at the moment of creation five minutes ago all the fossil records and all the imprints of memory in the minds of men had been created as they were at that moment, there would be no way of distinguishing one situation from the other. In the absence of the possibility of implementing any distinction between the two situations the distinction itself loses meaning.

Bridgman's thesis obtains only if his "proved" requires infallible demonstration. If instead we consider the relative strengths of evidential support, then the two situations are readily separated. Take the fantasy proposal of creation five minutes ago. There is *no evidence at all*, specifically, for a creation event five minutes ago with all evidence of an older world somehow spuriously implanted within it. Now take the prosaic proposal of an ancient world without mysteriously implanted deceptions. We have abundant and powerful inductive evidence for it across the sciences, notably in geology and biology. If we accept that we should be guided by the evidence, the fantasy proposal is self-defeating. The very statement of the fantasy proposal contains within it the assertion of the existence of compelling evidence for the prosaic proposal.

5. Fallibilism and the Dynamics of Scientific Progress

A direct application of this last consideration casts an optimistic light on scientific progress. That widely accepted scientific theories can fail is a commonplace in the history of science. We shall see in Chapter 9 below that this failure is too easily misread as a basis for antirealist skepticism about all science. Of course, there are many cases in which a well-accepted theory fails and does so because the evidence for it never was strong.[4] However even inductively well-supported theories can prove to be false. It is unlikely that this would never happen. That would require the epistemic dice to fall in our favor every time we make an inductive gamble. Such good fortune can only be expected in an extremely unlikely, fantasy world. In the real world, in our real history, the epistemic dice is rolled very often. Failures will eventually arise no matter how scrupulously we followed the evidence.

Thus, at each stage of a science as it develops, some—we hope most—of its results are correct, but some are not. We cannot then be sure of which are correct. At the time, the best we have is strong inductive support for the results and a lingering sense that this support is fallible. Then the science moves on, commonly through the finding of new empirical evidence. Through it, we learn that some of the results of the past stage of science were incorrect, in spite of their enjoying what we then thought[5] was strong inductive support. They are replaced with new results better supported by the expanded evidence. The process repeats. Our science grows. We correct past errors and replace them with newer, well-supported propositions. At no stage can we be sure that we have secured the final truth. The best we can say is that the inductive evidence for our most secure results is so strong that their falsity, individually, is extremely unlikely.

---

[4] We each likely have favorite examples. I have long argued that the claim of a non-trivial connection between information and thermodynamic entropy is false and that it fails for lack of empirical evidence.

[5] Recall that, according to the material theory of induction, judgments of inductive support derive from facts in the relevant domain. If we adjust our appraisal of which those facts are, we must make corresponding adjustments to our appraisal of relations of inductive support.

The outcome is an optimistic appraisal of scientific revolutions. That a widely accepted theory of the past is overturned is, by itself, no basis for inferring that its earlier proponents committed some sort of irrationality or error in evaluating of the evidence. Rather, even if scientists follow the import of evidence scrupulously, eventually errors will creep somewhere in their results.

The next chapter will return to this dynamic in the discussion of the Gettier problem, where it will be illustrated by examples drawn from the history of science. I will argue that it is routine for successes of the earlier stages to be Gettier-like cases, when judged by the results of the later stages, and that these Gettier-like case results are carried forward to the new stage. Here, Gettier cases are not a dire epistemic threat, but a familiar and celebrated part of the progress of science.

6. Conclusion: Defining Scientific Knowledge

How are we to define "scientific knowledge" in a way that is compatible with small-e empiricism? It is tempting to insist that a result can only count as scientific knowledge if it is true. The difficulty with this necessary condition is that truth is opaque. That is, we have no way of assuring that some specified contingent proposition is true. The best we can do is to affirm that the result has very strong inductive support from experience. That makes its truth very likely, but not infallibly so. We can do nothing more.

If we insist on truth as a condition for scientific knowledge, we are left in an awkward position. We cannot positively affirm that *any* of the results of even our best science constitute scientific knowledge. At best, they are merely candidates for scientific knowledge.

If the term "scientific knowledge" is to apply to any of the results of our best sciences, we must retract from this unrealizable ideal. Small-e empiricism provides the natural retraction. Scientific knowledge consists of a body of contingent propositions concerning natural states and processes that have strong support inductively from experience. Their truth is not assured absolutely, but is very likely. Under this reduced definition, the core results of our best sciences do count as scientific knowledge. That is the way it should be.

References

Bridgman, Percy W. (1950) “Impertinent Reflections on History of Science,” *Philosophy of Science,* **17**(No. 1), pp. 63-73.

Feynman, Richard (1965) *The Character of Physical Law*. Cambridge, MA: The M.I.T. Press.

Jevons, W. Stanley (1874) *The Principles of Science: A Treatise on Logic and Scientific Method.* New York: MacMillan.

Laudan, Larry (1983) “Demise of the Demarcation Criterion,” *Physics, Philosophy and Psychoanalysis: Essays in Honor of Adolf Grünbaum*. Eds. R. S. Cohen and L. Laudan. Dordrecht: Reidel.

Peirce, Charles Sanders (1960) *Collected Papers of Charles Sanders Peirce.* Volume I. *Principles of Philosophy*. Eds. C. Hartshorne and P. Weiss. Cambridge, MA: Belknap Press.

Reichenbach, Hans (1938) *Experience and Prediction*. Chicago: University of Chicago Press. First Phoenix edition, 1961.

5. Empiricism, Epistemology and the Gettier Problem

Chapter 5

Empiricism, Epistemology and the Gettier Problem

1. Introduction

The preceding chapter investigated the version of fallibilism that arises in small-e empiricism. It applies specifically to relations of inductive support. They cannot provide an absolute assurance of the truth of contingent propositions in science. This fallibility, it was argued, does not compromise the authority of inductively well supported science and even has a visible role in the dynamics of scientific progress.

Fallibilism also arises in another literature, in epistemology in general philosophy. Its formulation and goals differ from those of small-e empiricism. Where fallibilism in small-e empiricism primarily concerns abstract logical relations of inductive support, fallibilism in epistemology primarily concerns knowledge and beliefs, that is, mental states, and allows for many sources of error beyond the fallibility of inductive inference specifically. Fallibilism in some form is widely accepted as an abstract principle in the epistemology literature, but, we shall see, its consequences are resisted.

The difference is clearest in their approaches to defining knowledge. Inductive fallibilism in small-e empiricism led, as a secondary result, to a deflation of the concept of scientific knowledge. A result belongs within scientific knowledge, if it is inductively well supported by empirical evidence. It does not need to be true.

In contrast, a controlling problem in epistemology is to determine how to preserve the requirement that beliefs must be of truths, if they are to count as knowledge. Trouble arises when, in addition to stipulating that knowledge requires truth, it is expected that satisfaction of the other parts of the definition will assure that the belief is of a truth. For then the problem becomes insoluble. Even the strongest inductive justification of a contingent proposition does not assure its truth; and we have no other means to assure the truth of a contingent proposition. The clearest manifestation that this literature seeks to solve the insoluble is the persistence of inconsequential debates over Gettier cases. The assurance that a belief is of a truth is secured by the unrealizable fiction of an omniscient narrator who just declares the belief in of a truth.

In this chapter, Section 2 reviews some of the versions of fallibilism in the epistemology literature. Truth of the proposition believed is still almost universally demanded of a definition of knowledge, even though the epistemological literature on knowledge has provided no means through which we can be assured that truth has been attained. Section 3 reports some minority attempts in that literature to deflate the concept of knowledge in a way similar to that of small-e empiricism. Section 4 reviews the "KK Principle," that to know something requires that we know that we know. Its failure follows from the opacity of truth for contingent propositions. Section 5 considers two ways that inductive fallibilism can be misconstrued. The pessimistic response is inductive skepticism and the idea that we cannot eliminate rationally fantasy scenarios of deceiving demons. It is mostly resisted in the epistemology literature. The optimistic response is that fallibility can be circumvented if only we try harder. It has led to sustained but futile attempts to evade fallibility.

Sections 6 and 7 report the most prominent failure in the epistemology literature to accept the consequences of fallibilism and opacity of truth. It is the immense literature on Gettier cases. In them, the justification for a true belief succeeds not because the justification operated as intended, but success came by luck or happenstance. They are taken to threaten the adequacy of the definition of knowledge as justified true belief. For they reveal what fallibilists should know, even the strongest of inductive justifications do not assure us of truth. Repeated efforts to protect the definition seek a way to distinguish cases in which strong justification does properly indicate truth and those in which it does not. The inductive opacity of truth entails that all these efforts fail and must fail. That is already evident in the Gettier case narratives, for the truth of a belief is recovered not from further inductive investigations, but from the declaration of a disembodied, omniscient narrator.

Section 9 argues that the enduring problems of this Gettier case literature are resolved if we do what inductive fallibility compels. We reduce the definition of knowledge to well-justified belief. Then Gettier cases cease to be troublesome and are the less worrisome cases. The fallibility of inductive inference is most troublesome when there is strong inductive support for a falsehood. The better case arises when there is some failing in the relations of inductive support, but the result nonetheless turns out to be true. They are the Gettier cases. Section 10 offers a brief summary conclusion.

2. Fallibilisms in the Epistemology Literature

The doctrine of fallibilism seems to be widely accepted by epistemologists in an enormous literature. Or so we are frequently told. The opening line of the abstract to Cohen (1988) is frequently cited and quoted with approval. It asserts: “The acceptance of fallibilism in epistemology is virtually universal.” Brown (2018, p. ix) has found, however, that fallibilism is fading with a rise of infallibilist thinking. That there is or, perhaps, was such widespread agreement in as fractured a community as that of the epistemologists is possible because there are so many versions of fallibilism. Here are two forms of it.

2.1 Human Frailty

The most anodyne version of fallibilism is a simple concession that we humans are fallible and that this fallibility must affect our claims to knowledge. Leite (2010, p. 370) expresses it as:

> Fallibilism. In the broadest sense of the term, an anti-dogmatic intellectual stance or attitude: an openness to the possibility that one has made an error and an accompanying willingness to give a fair hearing to arguments that one’s belief is incorrect (no matter what that belief happens to be about).

and

> In some recent discussions, the term “fallibilism” stands for the thesis that human beings are fallible about everything (or just about everything) they believe …

Hetherington (n.d., Section 5) leaves little doubt as to the expansive scope of this sense of fallibilism. In a list whose incompleteness he recognizes, he provides seven possible sources of error in belief formation. They are “misusing evidence,” “unreliable senses,” “unreliable memory,” “reasoning fallaciously,” “intelligence limitations,” “representational limitations” and “situational limitations.”

2.2 Incomplete Evidential Support

This attribution of the origin of fallibility to human frailty is remote from the narrower sense of inductive fallibility developed in this small-e empiricism. This latter sense is closer to a more precise version of fallibility commonly found in epistemological writing. Here are a few versions of it drawn randomly from the literature over many decades. Although the fallibility arises through a frailty in the means of justification, the core focus is still knowledge. Does this frailty still allow us say that we know something?

> A better way to understand fallibilism is as the claim that:
> (F1) It is possible for *S* to know that *p* even if *S* does not have logically conclusive evidence to justify believing that *P*. (Feldman, 1981, p. 266)

> … a fallibilist theory allows that S can know q on the basis of r where r only makes q probable.” (Cohen, 1988, p. 91)
> … the weaker principle which permits S to know q on the basis of r provided r makes q sufficiently probable. (Cohen, 1988, p. 92)

> More commonly, however, “fallibilism” is used as a name for a thesis about knowledge and justification: that we can have fallible justifications for our beliefs, and that it is possible to know that something is the case even if one has only a fallible justification for believing it. (Leite, 2010, p. 370)

> [Most contemporary epistemologists] claim that one can know even if one lacks evidence which guarantees the truth of what’s known. They embrace ‘fallibilism’: one can know that p even though one’s evidence does not guarantee the truth of p. (Brown, 2018, p.2)

These formulations of fallibilism typically omit mention of whether fallibilism allows us to attribute knowledge to belief in evidentially well-supported falsehoods. In the broader literature, truth is taken to be an uncontroversially necessary condition for knowledge.[6] My sense is that most fallibilists agree. Reed (2012, p. 585, his emphasis) makes this condition explicit in his introductory formulation of fallibilism:

---

[6] Ichikawa and Steup (2017, Section 1.1) “Most epistemologists have found it overwhelmingly plausible that what is false cannot be known.”

> Roughly stated, the basic idea is that the subject can know something even though it *could have been false*. This is not the same as saying that the subject can know something that is false–it is very widely accepted by philosophers that, if a belief counts as knowledge, it is true.

2.3 Comparison

We can now summarize how inductive fallibilism within small-e empiricism differs from the formulations of fallibilism in the epistemology literature.

Inductive fallibilism does not arise from any human frailty. It inheres in the non-demonstrative nature of inductive logic itself.

The general literature in epistemology formulates fallibilism in human terms of beliefs and, especially, as a problem for the special human belief state of knowledge. The narrower formulation of inductive fallibilism in small-e empiricism asserts fallibility as a property of the logical relation of inductive support. A connection with knowledge only arises derivatively if we proceed to consider the import of inductive fallibility for beliefs.

The general literature on fallibilism is concerned to preserve the necessity of truth for knowledge, in light of the challenge of fallibilism. The import of inductive fallibility for small-e empiricism is that this requirement of the necessity of truth should be dropped. It is an idealized requirement whose satisfaction will always remain opaque to us.

Considerable attention is given in the fallibilism literature to the question of whether necessary truths, such as those of mathematics, can be known, when fallibilism is adopted. Inductive fallibilism is limited to considerations of the inductive support of contingent propositions.

3. Knowledge Demoted

The mainstream view in epistemology seems to be that knowledge should retain its centrality. However, there have been repeated suggestions, to me welcome, that this centrality is a mistake. I do not have a thorough survey of these proposals. Here are a few of them, collected during my reading.

Jeffrey (1968), in his Bayesian analysis of belief, "Probable Knowledge," writes: (p. 166)

> The obvious move is to deny that the notion of knowledge has the importance generally attributed to it, and to try to make the concept of belief do the work that

> philosophers have generally assigned the grander concept. I shall argue that this is the right move.

Kaplan (1985) captures the essential idea that the transition from strong evidential support to knowledge adds nothing accessible to our investigations. He wrote: (p. 355, his emphasis)

> Suppose again that you have carried out inquiry, come to believe that *P* on good evidence, and that you now ask, with this new analysis of knowledge in mind, "But do I *know* that *P*?" Once again, there is nothing to find out. Insofar as you are satisfied that your belief in *P* is well founded, you will ipso facto be satisfied that you have not inferred *P* from a false premise–otherwise you would not think you had good reason for concluding that *P*.

Perhaps the most colorful expression is Papineau's (2021) "The Disvalue of Knowledge."[7] Mimicking Russell's celebrated riposte on causation, he writes in his abstract: (p. 5311)

> I argue that the concept of knowledge is a relic of a bygone age, erroneously supposed to do no harm. I illustrate this claim by showing how a concern with knowledge distorts the use of statistical evidence in criminal courts, and then generalize the point to show that this concern hampers our enterprises across the board and not only in legal contexts.

4. The KK Principle

That inductive fallibility precludes us from affirming absolutely that we know something has some affinity with what is called the "KK principle" in epistemology. The principle asserts that if we know something, it must also be the case that we know that we know it. The principle was introduced in a formal analysis of a knowledge operator "*K*" in Hintikka (1962, p. 104). It asks if *Kp*, knowing proposition *p*, is equivalent to *KKp*, knowing that one knows *p*. The principle has been the subject of extensive debate. For a short survey, see Hemp (n.d.).

The assessment of the principle for present purposes is quite simple. If a belief must be of a truth to count as knowledge, then the KK principle fails for the reasons already given. That our belief in a contingent proposition is of a truth is not something that we can ascertain, beyond an assurance of very-strong inductive support. To require that we can know that we know in this case just repeats the same problem. To know that we know a contingent proposition requires

[7] I thank Joachim Horvath for directing me to this paper.

again that we can ascertain the truth of the proposition, which is precisely what inductive fallibilism in small-e empiricism precludes.

If we adopt the weaker alternative that knowledge just requires strong inductive support from experience for the contingent proposition believed, then the KK principle raises no special questions of foundational importance. That a believer knows that they know some proposition requires the believer to determine that the proposition is well supported inductively. That determination requires some facility with inductive logic that many do not have. According to it, many people may believe the earth orbits the sun, but they do not know it since they are unable to reconstruct the evidential relations that support the belief. An astronomer however can know the result since an astronomer has the expertise to explore these evidential relations. The distinction between the two cases is clear without any mention of "knowledge."

To designate only the astronomer's belief as knowledge, as the KK principle would require, is an arbitrary stipulation. All that is at issue is the assigning of an honorific term "knowledge." The important foundational concept is that of strong inductive support from experience. It determines how well we are informed of contingent propositions. A believer's awareness of the details of the evidential case is immaterial. What matters is that the evidential case has been made.

5. Misconstruals of Inductive Fallibilism

The most significant and enduring failure of the epistemology literature has been to insist on a definition of knowledge that requires the truth of its subject and to make that definition central to epistemology. There are two further misconstruals of inductive fallibilism. One is to adopt a far-reaching pessimism that arises most often in popular work: if we cannot ascertain truth inerrantly, then we have no basis to affirm anything. The other is a far-reaching optimism: that inductive fallibilism compromises our capacity to affirm truth inerrantly is merely a temporary obstacle that more careful analysis will surmount. Both of these misconstruals, I am pleased to report, have been identified and disputed by fallibilists in the epistemology literature.

Both misconstruals appear elsewhere. In later chapters below, I argue that they underlie both sides of the realism debate in philosophy of science.

5.1 Fallibilist Pessimism

Fallibilists in epistemology routinely seek to refute the idea that fallibilism leads to skepticism about knowledge. Cohen (1988) bases the refutation on the notion of "relevant alternatives." We judge that we know something by considering only relevant alternatives. Presumably the possibility of an artful deceiving demon with extraordinary powers is not among these alternatives. The difficulty Cohen then faces is to articulate a sufficiently precise sense of "relevant." This proposal is one among many. See Carrier (1993) for another account and Hetherington (n.d., Sections 8-10) for a survey of various forms of skepticism that may arise within fallibilism.[8]

My view is that a simple dictate of rationality precludes us taking deceiving demon scenarios seriously. We should follow what the evidence supports. Deceiving demon scenarios are, by construction, ones in which the evidence strongly supports the prosaic conclusion that world is just as it appears. They provide us no evidence at all for a fantastic, deceiving demon.

5.2 Infallibilist Optimism

To some and perhaps many, that we can never be assured of the truth of a contingent proposition is intolerable. They harbor the view that the fallibilism that supports it must be defective and may even result from some derangement in the thinking of fallibilists. That at least is the import of the metaphor used by David Lewis (1996, p. 550) in his defense of infallibilism if we are to take his light-hearted remarks seriously:

> We are caught between the rock of fallibilism and the whirlpool of scepticism. Both are mad! Yet fallibilism is the less intrusive madness.

After allowing that "people can get used to the most crazy philosophical sayings imaginable," he proceeds to the plea (his emphasis):

> If you are a contented fallibilist, I implore you to be honest, be naïve, hear it afresh. 'He knows, yet he has not eliminated all possible sources of error.' Even if you have numbed your ears, doesn't this overt, explicit fallibilism *still* sound wrong?

[8] Hetherington's (n.d., Section 1) summarizes: "Almost all contemporary epistemologists will say that they are fallibilists. Yet the vast majority of them also wish not to be skeptics."

Over a century after the revolutions of relativity and quantum theory shook our faith in absolute security of our best-established science, Lewis' visceral aversion to the ineliminability of error sounds quaint and nineteenth century.

Lewis is not alone in these sensibilities. Views such as his have promoted a spirited debate over fallibilism and infallibilism in epistemology. Brown (2018) is a monograph length repudiation of the rise of infallibilism and a defense of fallibilism.

Proposals for infallible knowledge of contingent propositions cannot work for all the reasons that support inductive fallibilism in small-e empiricism. I write with some weariness of the inevitable failure of attempts to formulate a cogent infallibilism. They cannot escape the fallibility of the only way we can learn of contingent propositions, inductive support. Of course, it is always possible to contrive convoluted narratives that mask the inevitable failure. Since the proposal must depend on unsupported contingent assumptions or covert, fallible inductive inferences, they will be found and the failure revealed. It is inevitable and requires only that we take the trouble to look for them. I think of the belief that this inevitability can be escaped if only we put more effort into it, the "try-harder fallacy."

Lewis' (1996) proposal, "Elusive Knowledge," is as good an illustration of this mode of failure as any. Lewis makes it easy in so far as he explicitly repudiates the essential role of justification in knowledge. It is, he argues (p. 551), neither necessary nor sufficient. He concludes:

> The link between knowledge and justification must be broken.

If ever we suspected that Lewis was no empiricist, this pronouncement leaves no doubt. To an empiricist, it has an air of magical thinking. We can somehow, it seems to say, know deep truths of the world without experiences of them. To support his infallibilism, Lewis (p. 551, his emphasis) offers the following definition:

> Subject *S knows* proposition *P* iff *P* holds in every possibility left uneliminated by *S*'s evidence; equivalently, iff *S*'s evidence eliminates every possibility in which not-*P*.

Prima facie, the definition sets an unachievable standard for knowledge, since the assured elimination of every possibility other than *P* is breathtaking in scope. Inductive inference cannot achieve it. What saves the definition, temporarily, is a further qualification that some possibilities are to be ignored. Lewis writes (p. 554, his emphasis):

> Our definition of knowledge requires a *sotto voce* proviso. *S knows* that *P* iff *S*'s evidence eliminates every possibility in which not-*P*–Psst!–except for those possibilities that we are properly ignoring.

This addition allows Lewis to escape the artifices of Cartesian skepticism. That a deceiving demon has made his cat invisible, he asserts (p. 562), can be properly ignored.

Which possibilities can be ignored? Lewis lays out an elaborate account in eight rules. The first three specify possibilities that cannot be ignored: "rule of actuality," "rule of belief" and "rule of resemblance."[9] They are followed by five rules about what may be ignored: "rule of reliability," two "rules of method," "rule of conservatism" and "rule of attention."

This totality of eight rules presents an inviting challenge to respondents. There is considerable scope for criticism of each rule individually and, no doubt, still further scope for concerns about how all eight are to be combined consistently. It is tempting to engage in this criticism. The weaknesses of the proposals are low hanging fruits, ripe for the picking. The temptation should be resisted. There is sufficient ambiguity throughout the proposal for it to sustain endless, inconclusive debate. This is how dubious work can generate an apparently energetic literature that can never yield anything definite.

Rather, the failure of the analysis can be seen without engaging in any of the details of these rules, which is why I have not given their formulations. Lewis' project is to provide an account of infallible knowledge. Lewis notes (p. 554) that the truth is not included in his definition of knowing. However, it is secured by the first rule of actuality: "It follows [from the rule] that only what is true is known." That is, following Lewis' rules leads to truths and it does so *infallibly*.

How can the rules achieve this?[10] One possibility is that they somehow tap into a resource that enables us to divine truths of the world without a basis in experience. That Lewis

[9] This last rule is Lewis' response to Gettier problems. We do not know the Gettier proposition since, Lewis argues by considering several Gettier cases, we cannot ignore the possibilities in which the Gettier proposition is false. (pp. 557-58)

[10] I ignore here the escape that Lewis's system operates vacuous in delivering no truths or delivering only very few in highly contrived truths. It is just another way in which the system would fail.

has found such a resource is not evident in his analysis. In Chapter 12 of *small-e empiricism*, in arguing for small-e empiricism, I made the case that there is no such resource.

From the examples in Lewis' text, it is evident that he expects us to secure knowledge through experience and thus, inevitably, his rules take us to truths beyond experience. To recall the examples in Lewis' first paragraph, how else are to know what food penguins eat and that the Australian rules football club, Essendon, won the 1993 Grand Final. That is, Lewis' rules are a contribution to inductive logic and a remarkably convoluted one. Since any inductive logic is non-demonstrative, it is necessarily fallible. The artifice of precluding some possibilities as ignorable does not save it. For inferring that they are ignorable is itself an inductive inference and thus fallible.

I will forgo the exercise of identifying precisely where and how the rules fail to be infallible guides to truth. For identifying any specific failure invites responses that trigger the interminable and inconsequential debates that convoluted proposals like Lewis' engender. The details do not matter. The failure is inevitable.

It seems to me that Lewis' analysis has already shown clear signs of decay. In my analysis of the material theory of induction (2021), I found a common mode of failure afflicting rule-based accounts of inductive inference.[11] Each starts with a simple rule. Counterexamples threaten it. They are parried by adding further conditions to the original, simple rule. Then counterexamples arise for these further conditions, so that still further conditions are needed. The cycles continues without a definitive termination. The rule-based account is in a death spiral. That Lewis' account requires *eight* rules suggests that it is already well-advanced in the death spiral.

[11] See for example the analysis of analogical inference in (Norton, 2021, Chapter 4).

6. The Gettier Problem

6.1 Its Persistence

It would be comforting to imagine that the prevalence of fallibilism in epistemology has led to a relaxed attitude to the requirement of truth for knowledge.[12] We all surely agree that truth is our goal. Might we be prepared to adopt a conception of knowledge that allows us to say that we know something without stipulating that its truth must be included in the definition of knowledge? Might epistemologists shift their emphasis to the strength of a justification in a belief and accept that truth will follow, mostly but not invariably?

The longevity of writing on the Gettier problem is a strong indication that this comfort is, alas, just my wishful thinking. In that literature, as we shall see, there is a near universal agreement that the truth of its subject is part of the *definition* of when a belief is knowledge. Before outlining the details of the Gettier problem, it is worth assessing just how massive is the attention it has attracted. Gettier's paper was published in 1963. In an early 21st century review,[13] Hetherington (n.d., a) noted:

> There is no consensus, however, that any one of the attempts to solve the Gettier challenge has succeeded in fully defining what it is to have knowledge of a truth or fact. So, the force of that challenge continues to be felt in various ways, and to various extents, within epistemology.

The situation has persisted over 50 years after Gettier's paper. Hetherington (2016, p. ix) noted that the Gettier problem literature was characterized by "widespread frustration" that…

> … epistemologists as a group, it seems, remain as far as they have ever been from agreeing on why the Gettier Moral[14] is true.

[12] My thanks to the many who have indulged my earlier attempts to understand the Gettier problem, including Joachim Horvath, Kareem Khalifa, James Norton and Timothy Williamson. My apologies to those among them who tried to save me from my Gettier heterodoxy.

[13] The latest reference in this undated article is from 2006.

[14] [of the inadequacy of defining knowledge as justified, true belief]

General search results affirm the extraordinary attention the problem has attracted. Gettier's (1963) paper had been cited over 6500 times[15] by November 2025, according to Google Scholar. An ordinary Google search on November 7, 2025, on the term "Gettier problem" returned 566,000 hits; and on "Gettier case" returned 511,000 hits. These are extraordinarily large numbers, given that the issue is highly specialized and likely only to attract attention from epistemologists.[16]

One might imagine that a provocative paper like Gettier's would initially attract attention after publication, that in the subsequent few decades the literature would find a generally accepted accommodation and then it would fade from the literature. There are indications that something like the reverse has happened. A Google n-gram plot of the presence of the terms "Gettier problem" and "Gettier case" in Google Books' digitized library for the period 1960 to 2022 (Figure 1.) indicates a marked rise in work discussing the terms in the decade 2010-2020.

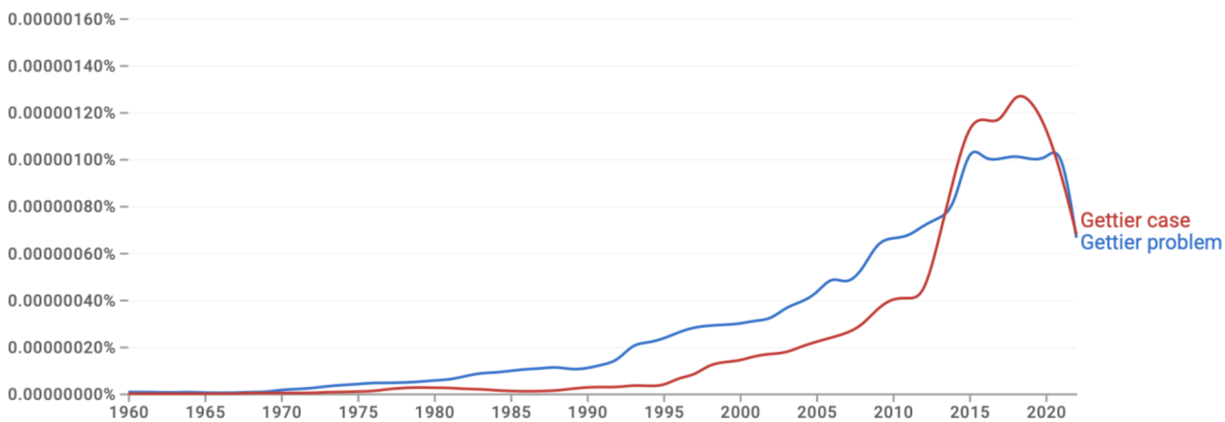

Figure 1. Distribution of the terms "Gettier problem" and "Gettier case."

Whether it is deserved or not, this level of attention and fragmentation of responses shows that the Gettier problem has posed an enduring and vexing problem for epistemology.

[15] Specifically, 6594 citations on November 9, 2025.

[16] For this reason, I despair of a literature search that does justice to the breadth of writing on the topic and apologize to readers whose favorite Gettier problem paper has been overlooked here.

6.2. The Problem

Gettier's (1963) paper argues for the insufficiency of a definition of knowledge that, it is generally allowed, has been recognized since antiquity. It is that a believer *S* knows a proposition *P* if and only if three conditions are met: *P* is true; *S* is justified in believing *P*; and *S* believes *P*. In a short, familiar slogan, knowledge is justified, true belief. The case for its insufficiency is made by displaying counterexamples, in which these three conditions are met, but which we are supposed to see intuitively do not count as knowledge. In them, *S* is justified in believing *P*, but the mechanism through which the justification operates fails. Nonetheless some other circumstance intervenes and *P* turns out to be true after all. Or so an omniscient narrator assures us.

The repertoire of examples has grown large over the decades, as discussion of the problem persisted. An example due to Bertrand Russell, presented prior to Gettier's paper, is typical. In my version of it, in the town square, I notice that the clock on the tower reads noon. Its reading justifies my belief that it is noon. What I do not know is that the clock had stopped a while ago and was left fixed at the noon reading. As it happens, it is noon, so my belief is of a truth after all. Russell (1923, p. 170-71) concludes of his version of the example, in line with the general view of Gettier cases in the literature, that:

> … this man acquires a true belief as to the time of day, but cannot be said to have knowledge. … Such instances can be multiplied indefinitely, and show that you cannot claim to have known merely because you turned out to be right.

The repertoire of Gettier cases in the literature consists, as far as I can see, of situations in ordinary life. The precise notion of justification employed is commonly left open. In the event, the notion is inductive, which keeps Gettier cases within the compass of the present discussion. Most commonly, the justification for a belief is an induction from the evidence of human sense experience. In the first of Gettier's (1963) cases, Smith comes to believe contingent facts about who secured a position. The evidence is an assurance by a company president that Jones will be selected for a position and that Smith has counted ten coins in Jones' pocket. In the other of Gettier's examples, Smith believes Jones owns a Ford since he has always seen Jones drive a Ford. In other examples, recalled in Hetherington (n.d. a, Section 4), we see what looks like a sheep but is in fact a shaggy dog; we strike matches we believe on past experience always ignite, but do not; we see what looks like a barn in a field, but is in fact merely a barn façade. Ichikawa and Steup (2017, Section 3) offer the case of someone misidentifying a mirage as water, where fortuitously there is water. It is essential to the examples that there is an omniscient narrator who declares the truth.

6.3. Escapes

Over the decades, many solutions to the Gettier problem have been proposed. What they generally share is a retention of the requirement that knowledge requires truth. Ichikawa and Steup (2017, Section 3, their emphasis) note:

> Few contemporary epistemologists accept the adequacy of the JTB analysis. Although most agree that each element of the tripartite theory is *necessary* for knowledge, they do not seem collectively to be *sufficient*.

For over half a century since Gettier's paper, epistemologists have sought to repair the characterization of knowledge as justified true belief, almost invariably by adding further conditions. The range of proposals is so great that no synopsis is possible here. We can however get some sense of its enormity by recalling the categories of responses in surveys of the Gettier problem. Hetherington (n.d. a) lists as "attempted solutions": "Eliminating Luck," "Eliminating False Evidence," "Eliminating Defeat," and "Eliminating Inappropriate Causality"; and as "Attempted Dissolutions": "Competing Intuitions" and "Knowing Luckily." Ichikawa and Steup (2017) list as possible fourth conditions "No False Lemmas," "Sensitivity," "Safety," "Relevant Alternatives," "Reliabilist Theories of Knowledge," "Causal Theories of Knowledge," "Epistemic Luck," "Methodological Options," "Virtue-Theoretic Approaches," "Knowledge First," "Pragmatic Encroachment" and "Contextualism."

Among many recent proposals, I found noteworthy attempts to strengthen the justification for knowledge. Horvath and Nado (2021) define a threshold of "normality" above which justification must rise before it suffices for knowledge. de Grefte (2023) argues that the original definition of knowledge as justified true belief survives if we require that the justification arises from a reliable method.

Prominent among all these proposals is the idea that Gettier cases do not count as knowledge because their agents arrived at the truth by luck. This proposal seems to capture well a sense of what goes wrong in Gettier cases and, for that reason, seems to be quite widely accepted. A popular version was defined by Prichard (2005, p. 146):[17]

> *Veritic epistemic luck*
>
> It is a matter of luck that the agent's belief is true.

That knowledge of truths cannot be acquired by luck is an anti-luck principle, which can be formulated in positive terms as a principle of safety (p. 71)

> *The safety principle*
>
> For all agents, φ, if an agent knows a contingent proposition φ, then, in most nearby possible worlds, that agent only believes that φ when φ is true.

It is a version of Sosa's (1999) safety principle, but now expressed in terms of possible world semantics. Further reflection led Prichard (2009, p. 41) to add a second condition:

> … an ability condition of some sort—i.e., a condition to the effect that the true belief was gained via the employment of the agent's reliable cognitive abilities.

The notion of an agent's ability leaves the content of the condition open. A more precise version is found in de Grefte (2023, p. 533):

> VERITIC LUCK: A belief is veritically lucky if and only if it is a matter of luck that the method one used to form one's belief produced a true belief.

This is the understanding of veritic luck that I will use below.[18]

6.4 What the Escapes Cannot Do

What can these escapes hope to achieve? Here an obvious point bears repetition and emphasis. No matter what condition is added to the definition of knowledge to tame the Gettier cases, the methods of justification available remain inductive. Therefore, they are fallible and can always lead us into belief in a falsity. No escape can eliminate this enduring difficulty of inductive support. The best that all these efforts can achieve is some sort of comfort that the true belief arising in Gettier cases fails to be secured in a way that merits calling it "knowledge." The

---

[17] The term "veritic" is, as far as I can tell, an adjectival derivative of "verity" meaning truth. It is idiosyncratic and does not appear in the dictionaries I consulted. Its use was rare and sporadic until Prichard's introduction of it into the literature.

[18] A minor point is that the "only if" conjunct is too strong in that it rules out veritic luck that cannot be attributed some aspect of a method.

presumption is that there is a way to distinguish true beliefs secured meritoriously from those not so secured and that this way can be captured in simple formulae or rules. That the debate remains open over 60 years after Gettier's paper suggests that this presumption is false.

6.5 Why the Fuss?

The Gettier problem literature is driven by a sense that Gettier cases pose a profound challenge to the concept of knowledge. As the discussion elsewhere in this chapter indicates, I do not sense that they are troublesome for a fallibilist empiricist. My puzzle has been to understand why my indifference is exceptional, even while epistemologists now generally accept fallibilism and have some sympathies with empiricist ideas. It is not really a question in philosophy, but one in the psychology and sociology of epistemologists. Hetherington (2012, p. 217) has given an answer that I find plausible:[19]

> A methodological mistake is allowing them not to notice how they are simply (and inappropriately) being infallibilists when regarding Gettiered beliefs as failing to be knowledge. There is no Gettier problem that we have not merely created for ourselves by unwittingly being infallibilists about knowledge.

This matches my sense that the acceptance of fallibilism by epistemologists is limited and does not extend to an important consequence of fallibilism, the opacity of truth.

In earlier centuries, it seemed quite possible to attain truth and recognize that it was attained. In the eighteenth century, it was widely felt that Euclid had done just this with the truths of geometry and Newton had done the same with truths of dynamics. The challenge to philosophy was to find an infallibilist epistemology that capture the methods they used. Might they employ a Cartesian deductivism? Euclid deduced his geometry from self-evident truths. Or might we look to some development of Kantian notions of synthetic a priori truths? They assure us of the truth of Euclid's geometry and Newton's mechanics.

The alarm over Gettier cases indicates to me a lingering attachment to this older, infallibilist project. For the Gettier cases suggest a mismatch between our best methods and truth

[19] Hetherington does not adopt my approach of dropping the requirement of truth from the concept of knowledge. Rather he finds that Gettier cases are thought mistakenly to pose a problem for justified true belief through a modal fallacy. It conflates partial and holistic characterizations of Gettier cases through a tacit assumption of infallibilism.

that would trouble infallibilists, but not fallibilists. If we look for it, we can find attachments to infallibilism in at least some responses to the Gettier problem. An early version of the anti-luck escape from Gettier cases is provided by Unger (1968). His formulation is that knowing precludes accidental truths: (p. 161)

> For any sentential value of *p*, a man knows that *p* if and only if it is not an accident that the man is right about its being the case that *p*.

This wording leaves open whether accidents allow the attaining of truth by a means compatible with fallibilism, that is, by a means only assured of success with very high probability. Unger goes to some pains to rule out this fallibilist possibility. He imagines a well-shuffled deck of a billion white cards and just one of another color. He asks if a man can know that he will win a bet if the top card is white. Unger (p. 162) inclines to answer that he does not:

> Still, we may also find ourselves saying that he cannot really know that he has won until the color of the card is actually revealed. Rather, the increase to such high chances of success furthers our readiness to apply our analytic condition, to say that it is not at all accidental that the man is right (assuming of course that he is right). But again, and equally I think, our willingness here is not so complete as it might be.

The one in a billion chance of failure is enough to preclude knowledge. Whether Unger's attachment to infallibilism is shared more widely, even in some deep corners of epistemologists' thinking, goes beyond what I can ascertain.

7. The Tenacity of the Opacity of Truth

7.1 The Generality

Each of the various forms of the escapes to the Gettier problem amount to a variant definition of what it is to know. Considerable energy has gone into debates over which is the best definition. In so far as these definitions require truth as a condition for knowledge or that truth follows from the definition, these debates are futile. For they are debates over definitions that fail to meet a condition of adequacy for a definition of scientific knowledge developed in the previous chapter: there must be at least some results in our best science that can be designated as belonging within scientific knowledge. To recall, the difficulty is that even the most prized items of our best science are only supported by fallible inductions. We cannot positively affirm their truth. The best we can affirm is very strong support such that we deem their truth very likely. It follows that *none* of the definitions of knowledge that require truth enable us to affirm that anything in our best science is knowledge.

7.2 Illustrations

Consider again my walk in the town square, where I notice that the clock on the tower reads noon. I am now well justified in believing that it is noon. That is all I can say. It will be that way if the clock is well functioning and it is noon; if the clock is stopped and it not noon, so my belief is mistaken; or if the clock is stopped and it just happens to be noon, so it is a Gettier case.

There is a temptation to imagine that we can somehow step outside of the case at hand and determine which is the real situation. That is an illusion. We can only assemble more evidence that allows us to be well justified in our belief about the situation. Or we may find evidence that may undo that justification. We will never ascertain the truth with absolute certainty. We will forever be trapped in the narrative of the case.

To see this, continue the above case. After reading the clock on the tower and forming the belief that it is noon, I glance towards a shop window and notice a second clock that also reads noon. I am reassured in my belief that it is noon. Can I now affirm infallibly the truth of my belief? No. I am still trapped in the narrative. It may well be that this clock is functioning well, so I now have further justification that it is noon. Or I may just be in an extended Gettier case. It may turn out, let us say, that all the clocks in the square are synchronized to a single standard clock. All of them are set to noon temporarily as part of regular maintenance of the synchronization system. If it is not noon, my belief would be mistaken. If it is noon, I am now in an extended Gettier case.

This situation can be continued indefinitely. No matter how the story unfolds, we will only ever have fallible support for our belief; or we may end up with evidence that properly shakes our confidence in that belief. That would be the case if the clock in the shop window were to have read 11am. We will never have an infallible affirmation of the truth of the belief. In the real world, there is no omniscient narrator whose authoritative, disembodied voice murmurs through the square for us to hear: "… as it turns out, the clock has stopped at noon, but it just so happens that it is noon."

That we can never step out of the narrative and adopt the position of an omniscient narrator has long been noted in the literature. Using examples similar to those above, Zagzebski (1994, p. 73) concludes:

> As long as the truth is never assured by the conditions which make the state justified, there will be situations in which a false belief is justified. I have argued that with this common, in fact, almost universal assumption, Gettier cases will never go away.

8. Cartesian Doubt

This analysis makes clear that definitions of knowledge that require truth do not just fail the condition of adequacy for scientific knowledge. They also fail to be adequate to the most prosaic of examples. Famously, Descartes, in his First Meditation, asked provocatively: (1641, p. 14)

> Indeed, that these hands themselves, and this whole body are mine—what reason could there be for doubting this?

We would like to be able to affirm that we *know* we have hands. However, even our seemingly most secure beliefs such as these are only supported inductively and may fail to be true. In the "phantom limb syndrome," amputees continue to have the sensation of the presence of the amputated limb. The sensation is not veridical. There is something along the lines of a reverse syndrome in "body integrity identity disorder." Suffers become convinced that some body part just should not be part of their body. They doubt, in agreement with Descartes' provocation, that the hands they sense are theirs.

As far as I can see, none of the Gettier inspired amendments to the definition of knowledge as justified true belief escape this difficulty. Briefly, amendments that require justification to meet stringent, further conditions fail, since under them justification remains inductive and thus fallible. Can we dismiss the threat of inductive failures by casting them out as requiring possibilities that can properly be ignored, as does Lewis in Section 5.2 above? The suggestion faces the same problem. The decision over what can be properly ignored is itself inductive and thus fallible. It might seem proper to ignore the possibility that we are mistaken in sensing that we have hands. That may be seem so until we learn of phantom limb syndrome. More significant examples arise in science. A nineteenth century physicist might well have felt it proper to ignore the rather extreme possibilities countenanced by the quantum theory of the following century.

There is, I believe as indicated above in Section 5.1, a quite serviceable response to Cartesian skepticism. It is that our evidence provides overwhelming support for a world that is as it seems and no evidence for the fantasies of a brain-in-a-vat world. It does so without requiring an absolute affirmation of the truth of our beliefs. If we know that for which we have very strong evidence, then we know we are in a world that is just as it seems.

9. Gettier Cases are Benign

The remedy offered in this chapter to the travails of Gettier cases is straightforward. Attempts to augment the definition of knowledge for contingent propositions as justified true belief cannot overcome the opacity of truth. It is futile to try. The fallibility of inductive inference cannot be erased by even the most artful of added conditions.

Precisely because of its opacity, there is no loss to us if we drop truth from the definition of knowledge. In its place, I have suggested above, we should put beliefs that are well supported inductively by the evidence of experience without demanding their truth. Call such beliefs "knowledge" if you like. Or, to avoid confusion, just call them "beliefs well-supported by experience." We lose nothing, since all we ever have had are such well-supported beliefs. We have no means to go beyond what strong inductive support provides. We do not have an omniscient narrator to consult.

The outcome is that Gettier cases cease to be troublesome. They just become another instance of the familiar problem of dealing with the fallibility of inductive support. And they are a less troublesome case. The most common difficulty is that a result that is well supported by the evidence turns out to be false. Something went awry so that what was judged to be a reliable mode of justification failed. The usually safe assumption that the clock in the square is running well was not safe. If the frozen reading is wrong, I form a false belief about the time and will suffer the consequences.

There is a better case. The inductive support we relied for our belief may not function as we thought, but some other circumstance may intervene so that the belief turns out to be right after all. This is the better case. It is the Gettier case. If we have freed ourselves from the worry that the Gettier case troubles our definition of knowledge, all that remains is relief that things worked out. Or so we should think, in so far as we are able in the event to determine that ours is a Gettier case. I may be mistaken in assuming that the clock in the square is running well, but the belief that it was noon is of a truth. No harm was done.

10. Conclusion

Fallibilism entered the mainstream of thought in philosophy of science over a century ago, after relativity theory and quantum theory overturned what centuries before it had deemed unassailable certainties. An important outcome of this chapter is that epistemology has still to accommodate this most momentous of changes in science. In failing to do so, it continues to err in giving knowledge a central position in the epistemology of science, where the truth of a belief is deemed necessary for it to be knowledge. It so doing it has given pride of place to belief states whose attainment will forever be opaque for the contingent propositions of a science.

My sense is that the epistemology literature persists in this formulation since it has been unable to discard an earlier optimism in which the truth of contingent propositions could be ascertained absolutely. It avows an adherence to fallibilism but not to its import. The outcome is that epistemology has taken on the impossible task of giving a precise account of knowledge that still allows us to attribute knowledge to the content of our best science. Once fallibilism is adopted, there is little choice but to replace what we know assuredly in science with that for which we have the strong inductive support of experience. To persist with attempts to make sense of scientific knowledge as true belief is to engage in a project that cannot succeed. That unhappy fact, I believe, lies at the heart of the continuing struggles in epistemology over knowledge.

References

Brown, Jessica (2018) *Fallibilism: Evidence and Knowledge.* Oxford University Press.

Carrier, L.S. (1993) "How to define a nonskeptical fallibilism," *Philosophia*, **22**, pp. 361–372.

Cohen, Stephen (1988) "How to be a Fallibilist," *Philosophical Perspectives, Vol. 2, Epistemology*, pp. 91-123.

de Grefte, Job (2023) "Knowledge as Justified True Belief" *Erkenntnis*, **88**, pp. 531–549.

Descartes, René (1641) *Meditations on First Philosophy*. Trans. M. Moriarty. Oxford: Oxford University Press, 2008.

Feldman, Richard (1981) "Fallibilism and Knowing that One Knows," *The Philosophical Review*, **90**, pp. 266-282.

Gettier, Edmund L. (1963) “Is Justified True Belief Knowledge?” *Analysis*, **23**, pp. 121-123.

Hemp, David (n.d.) “The KK (Knowing that One Knows) Principle,” *The Internet Encyclopedia of Philosophy* https://iep.utm.edu/kk-princ/

Hetherington, Stephen (n.d.) “Fallibilism,” *The Internet Encyclopedia of Philosophy*, https://iep.utm.edu/fallibil/

Hetherington, Stephen (n.d., a) “Gettier Problems,” *The Internet Encyclopedia of Philosophy*, https://iep.utm.edu/gettier/

Hetherington, Stephen (2012) “The Gettier-illusion: Gettier-partialism and infallibilism” *Synthese*, **188**, pp. 217–230.

Hetherington, Stephen (2016), *Knowledge and the Gettier Problem*. Cambridge: Cambridge University Press.

Hintikka, Jaakko (1962) *Knowledge and Belief*. Ithika, NY: Cornell University Press.

Horvath, Joachim and Nado, Jennifer (2021) “Knowledge and Normality,” *Synthese* (2021) **198**, pp. 11673–11694.

Ichikawa, Jonathan and Matthias Steup (2017) “The Analysis of Knowledge,” *The Stanford Encyclopedia of Philosophy* (Fall 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.), https://plato.stanford.edu/archives/fall2024/entries/knowledge-analysis/

Jeffrey, Richard (1968) “Probable Knowledge,” *Studies in Logic and the Foundations of Mathematics.* **51**, pp. 166-190.

Kaplan, Mark (1985) “It's Not What You Know that Counts,” *The Journal of Philosophy*, **82**, pp. 350-363.

Leite, Adam (2010) “Fallibilism,” in J. Dancy, E. Sosa and M. Steup, *A Companion to Epistemology*. 2nd ed., Wiley-Blackwell, pp. 370-75.

Lewis, David (1996) “Elusive Knowledge,” *Australian Journal of Philosophy*, **74**, pp. 549-67.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Papineau, David (2021) “The Disvalue of Knowledge,” *Synthese.* **198**, pp. 5311–5332.

Prichard, Duncan (2005) *Epistemic Luck*. Oxford Clarendon Press.

Prichard, Duncan (2009) “Safety-based Epistemology,” *Journal of Philosophical Research*, **34**, pp. 33-45

Reed, Baron (2012) “Fallibilism,” *Philosophy Compass* 7/9, pp. 585–596.

Russell, Bertrand (1923) *Human Knowledge: Its Scope and Limits*. London: George Allen and Unwin. Repr. 1948.

Sosa, Ernest (1999) “How to Defeat Opposition to Moore,” **33**, *Supplement: Philosophical Perspectives,* 13, *Epistemology*, pp. 141-153.

Unger, Peter (1968) “An Analysis of Factual Knowledge,” *The Journal of Philosophy*, **65**, pp. 157-170.

Zagzebski, Linda (1994) “The Inescapability of Gettier Problems,” *The Philosophical Quarterly*, **44**, pp. 65-73.

6. Gettier-Like Cases are Successes of Scientific Progress

Chapter 6

Gettier-Like Cases are Successes of Scientific Progress

1. Introduction

The Gettier cases reviewed in the preceding chapter have attracted considerable attention in the epistemology literature. They have, as far as I know, essentially no presence in the literature in philosophy of science. That may be, we might suppose, because standard examples of Gettier cases involve fanciful scenarios that are remote from actual science and are artfully contrived to make a specific point. Even in the fictional worlds of their invented scenarios, they are by construction extremely rare. They require *two* unlikely occurrences to fit together in just the way needed. First, an otherwise reliable mode of justification for some result must fail to operate as intended; and, second, some further unlikely coincidence allows the result still to be true. Even that is not enough. An essential part of the scenario is that there is a disembodied, omniscient narrator who just declares the truth for our convenience.

These considerations may lead to the expectation that Gettier cases are so rare as to be irrelevant to actual scientific practice. This chapter will show otherwise. Gettier-like cases, it turns out, are pervasive in the history of science; and their results are routinely celebrated as scientific successes, even when their Gettier-like character is recognized.

Two aspects of Gettier cases do not readily translate into real scientific cases. To see how these Gettier-like cases arise in the history of science, Section 2 reviews surrogates for these aspects. First, in place of the omniscient narrator are the judgments of a later stage of the science as to which of the results of the earlier stage are true or false. Second, there is no simple analog in the sciences for success through unanticipated luck. In its place is that success came even though the scientist was using methods that are found in later investigations to be unreliable.

Section 3 reports that Gettier-like cases are not just possible in science, but pervasive. It notes further how the variability often arising in methods used in them follows from improvements in our understanding of prevailing facts. It also reports that the presence of these cases has not been an occasion for alarm in science, even when their Gettier-like character has been discovered.

Section 4 collects seven cases in which important and celebrated discoveries in science turned out to be Gettier-like. The examples recounted are: Bradley's discovery of stellar aberration, Dalton's discovery of atomic theory, Carnot's discovery of the founding principles of thermodynamics, Thomson's discovery of the electron, Bohr's discovery of the quantum mechanics of atoms, Einstein's discovery of the general theory of relativity and Hubble's determination of the Hubble law. In each case, the Gettier-like result is not just preserved as the science progresses. They are routinely celebrated as major achievements in the foundations of a science.

Section 5 concludes by comparing the reception of these cases in the epistemology and science literatures. That the original discovery might not conform with an epistemologist's definition of knowledge is immaterial to the scientists. What matters is that the result was discovered, that we can secure an evidential foundation for it and that the original discoverer does deserve credit.

2. The Character of Gettier-Like Cases in Science

2.1 A Surrogate for the Omniscient Narrator

The first difficulty in identifying Gettier cases in science is that we have no omniscient narrator to tell us when such a case has arisen. The dynamics of scientific progress supplies a suitable surrogate. We will never have an absolute affirmation of the truth of core propositions in science. As science progresses, each stage can identify what it considers to be the errors of the earlier stages and which new results are to be accepted as its truths. Gettier-like cases will be those in which an earlier stage is, informally, lucky in its determination of a truth, as judged by the next stage.

2.2 Method Replaces Luck

Luck, or more precisely, *unanticipated* luck is the key fact in a Gettier case.[20] The difficulty in scientific examples is to determine when some truth has been achieved by

[20] Is there such a thing as *anticipated* luck? A beachcomber, sifting the sand each day, does not know, on any particular day, whether luck will turn up some valuable. However, the beachcomber anticipates that if the sifting is carried out repeatedly, day after day, luck will eventually come their way.

unanticipated luck. In the contrived examples from ordinary life, we have a default expectation on the range of possibilities in the fictional Gettier scenario. If we are willing to distribute chances over them, we can then determine whether some particular instance has a small chance and thus can be lucky. It is plausible from default assumptions about how I might behave that my visit to the town square (in the previous chapter) could reasonably happen over a wide time interval. It does seem unlikely and thus lucky that my visit coincides precisely with noon. It is unanticipated luck since I have no reason to think that the precise timing of my visit has special import.

It is, however, quite unclear that a similar concept of luck can be implemented in the cases discussed below in the history of science. Niels Bohr, for example, used a collection of ideas from 19th century electrodynamics, applied inconsistently, to arrive at some key results of the coming quantum theory of the atom. He also used these same ideas to derive many more results inconsistent with the coming theory. We can loosely assign chances to the different times for my visit to the town square. Over what possibilities do we assign chances so that Bohr's successes are lucky?

A better way to identify Gettier cases in the history of science, is supplied by de Grefte's (2023, p. 533) reliabilist characterization above of Gettier cases. He wrote: "…it is a matter of luck that the method one used to form one's belief produced a true belief." The key idea for me is that luck is defined in terms of what the *method* can properly deliver. We can discern the methods used by a scientist and we can determine retrospectively whether the scientist had any basis for expecting those methods to provide results that are important to the next stage of science.

A Gettier-like case in science in one in which a method that turns out in retrospect to be ill-adapted to the science nonetheless proves to deliver a truth, as judged by the next stage. Among the many parts of the method is one that just happens to work by the retrospective judgment of a later stage of the science. That is, some component of the method used replicates something that works according to the later science. The scientists of the time cannot know of the future validation of this component, since they do not know the future science. Call it an "unanticipated saving factor." It is the distinctive feature of Gettier-like cases in science. They are, informally speaking, the lucky occurrence that makes the case Gettier-like.

2.3 Actual Performance Replaces Possibility

In the Gettier literature, it is routine to make judgments of luck by means of possible world semantics. A success is lucky if the same success would not be realized in some suitably defined nearby possible world. I am loath to use this approach to assess the adaptation of the methods. My concern over judgments concerning possible worlds is their malleability. We need to determine which possible worlds are closer to the actual and which are farther away; and we need to determine what it is for the conditions and methods of the actual world to be realized in them. My suspicion is that there is sufficient elasticity in the complications of these determinations that a creative philosopher could manipulate them to get whichever result happens to suit their purposes.

Here the fact that the examples are real and not fictional stories gives us a better way to proceed. In each case, we can identify the methods actually used by the scientists; and we can determine their prospects. In the Gettier-like cases, we will see that the methods used routinely returned many results deemed false in the retrospective view of later stages. Thus, we have no good basis for expecting these methods to return a result amenable to the later stage. Considerations of luck have been replaced in the analysis by considerations of the actual performance of the method. There is no need to entertain fragile counterfactuals as to how the method would perform in some imaginary possible world. We consult the actual history.

3. Gettier-like Cases are Common

These general reflections open the possibility of Gettier-like cases in scientific progress. It does not establish their frequency. When I began looking for them in the history of science, I was astonished at how often they appear. Gettier cases are *pervasive* in the history of science. The following section recounts briefly a selection of them from a range of sciences with whose history I have some comfort. The examples selected are notable for their importance in the history of science and for their subsequent celebration as major discoveries in science. The examples are: Bradley's discovery of stellar aberration, Dalton's discovery of atomic theory, Carnot's discovery of the founding principles of thermodynamics, Thomson's discovery of the electron, Bohr's discovery of the quantum mechanics of atoms, Einstein's discovery of the general theory of relativity and Hubble's determination of the Hubble law.

3.1 Variability of Methods

We will find that the methods used in the historical examples change over time. This is an expectation of the material theory of induction. Central to the methods are judgements of inductive support. These judgments are in turn dependent on the facts that prevail in each domain. Thus, the scientists' assessments of these relations of support depends in turn on their assessments of which are the facts of their domains. Since these last assessments are fallible, it will turn out commonly that methods that they deem reliable are not so. That variability is what is responsible for the prevalence of Gettier-like cases in the progress of science.

3.2 There is No Cause for Alarm

The sense of alarm surrounding Gettier cases in the epistemology literature is misplaced in the context of science. The worst we can say of the Gettier-like cases in the history science is that the evidential case made for their results was, at the time of their discovery, imperfect. They were still discoveries, clearly important to subsequent developments in their science. For this they are celebrated by later scientists. Whether the scientists originally making the discoveries knew them in conformity with some epistemologist's definition of knowledge is immaterial. What matters is the discovery.

Here, I do not intend a skeptical moral about discovery in science. Rather, I intend an optimistic moral. These cases show how science can proceed successfully, in spite of the fallibility of its methods. We see scientists proceeding responsibly with the best methods then available, as they should, and that later work is able to identify the errors and correct them. In all the cases considered below, the Gettier-like successes do not result from pure happenstance. Rather, unknown to scientists at the time, there is something right about the methods they are using. What they cannot then distinguish are which are those parts and which are not. That is a retrospective judgment. In the case of Bohr's model of the atom, for example, his key results of the stability and energetic discreteness of his electron states was properly a reflection of the stability and discreteness of the observed atomic spectra. Bohr would have recognized this at the time, but he would not have been able to identify just which of his remaining results were erroneous, for his methods did not allow it.

4. Gettier Cases in the History of Science

In order to make it possible to recall a range of Gettier cases in the history of science, I have given only abbreviated accounts of each.

4.1 Bradley's Discovery of Stellar Aberration

James Bradley (1729) reported his discovery of stellar aberration to the Royal Society. According to it, the apparent angular position of a star is altered slightly as a result of the motion of the Earth. In the simplest case, in which the star is located perpendicular to the Earth's motion, we subtract vectorially[21] the Earth's velocity v from the velocity of the light c and find the angular displacement in the direction of the Earth's motion to be $\alpha = v/c$. Discussion of this result in *small-e empiricism*, Chapter 11, Section 12, found that the general formula is $\tan \alpha = \sin \theta (v/c) / [1+(v/c) \cos \theta]$, where θ is the angular elevation of the star in the direction of motion of the Earth. For Bradley's case of $v << c$, this general formula reduces to the one Bradley reported, $\sin \alpha / \sin \theta = v/c$.

Bradley's discovery, as recalled in *small-e empiricism*, Chapter 11, was a celebrated result and one of the few experimental results that Einstein reported as decisive for his discovery of special relativity. It is, on our best present understanding, a truth. However, Bradley's method was flawed. For he derived the result by Newtonian addition of velocities. Famously, according to special relativity, this method fails for motions close to or at the speed of light. Under the Newtonian rule, the motion of the Earth has to be added vectorially to the speed of a light signal to produce an altered speed for the light. A postulate of special relativity is that the speed of light remains unchanged, even if measured from a moving Earth. Special relativity offers a different rule for the composition of velocities in place of the Newtonian rule of vector addition that preserves the constancy of the velocity of light.

Bradley's result survived, however, because of an unanticipated saving factor. The relativistic[22] and Newtonian rules for velocity addition give different results for the composed speeds, but the *same* result for the angular deflection under the small v/c conditions of aberration.

---

[21] To use the vector concept, not introduced until the following century.

[22] The relativistic rule yields $\tan \alpha = (\sin \theta (v/c) / [1+(v/c) \cos \theta])\sqrt{1 - v^2/c^2}$. The correction to the Newtonian result is a term second order in (v/c) and thus negligible for stellar aberration.

There is an intermediate stage in the history of astronomy in which the Gettier-like character of Bradley's discovery was explicitly recognized. Bradley's Newtonian addition of velocities applied as long we could conceive of light as the rapid propagation of Newtonian light corpuscles. In the nineteenth century, this Newtonian theory of light was replaced by a wave theory. The Newtonian addition of velocities could no longer be applied in a straightforward way to the direction of wave propagation and it became quite unclear as to how the effect of stellar aberration was to be recovered. Arthur Berry, writing a history of astronomy at the end of the nineteenth century, expressed puzzlement over Bradley's success. His account captured rather perfectly its Gettier-like character with the verbiage I have emphasized below in his remarks (1898, p. 265):

> The curious inference may be drawn that, if the more correct modern notions of the nature of light had prevailed in Bradley's time, it must have been very much more difficult, if not impracticable, for him to have thought of his explanation of the stellar motions which he was studying and *thus an erroneous theory led to a most important discovery*.

The difficulty for the wave theory had been resolved in electrodynamic theory, though apparently not to Berry's satisfaction, by H. A. Lorentz' 1895 theorem of corresponding states. The mathematics of that theorem was then incorporated by Einstein into his special theory of relativity. Stellar aberration is, we now recognize, a relativistic effect, arising from the relativity of simultaneity.

4.2 Dalton's Discovery of Atomic Theory

John Dalton's (1808, 1810) *New System of Chemical Philosophy* is celebrated as the beginning of modern atomic theory in chemistry. Dalton proposed facts that we still believe lie at the foundations of chemistry: that each element is composed of identical atoms of that element; and that the atoms combine to form the familiar compounds of ordinary chemistry, such as water, muriatic acid, ammonia and so on.

Our concern is with the methods that Dalton used to arrive at his results. What is clear is that whatever methods he was using, they were quite unreliable and produced false results throughout his work. Part I of Dalton's *System* (1808) was concerned with heat and the physical properties of bodies formed from atoms. He determined (Chapter II) that the atoms of gases consist of particles enveloped in atmospheres of heat, where the latter was conceived as an element, caloric, following Lavoisier's system of elements. The enveloped cores are spherical and thus he could characterize a gas as arranged as a neat pile of shot. He used this static model to account for the elastic behavior of gases under changes of pressure and temperature, in direct conflict with the kinetic theory of heat that emerged a few decades later.

Dalton's treatment of the chemical properties of material in Part II (1810) consisted largely of extensive reports of chemical experiments and what could be inferred from them. What little he could infer of the details of the atomic constitutions of these chemicals was, again, largely incorrect. He did not recognize that elemental hydrogen, oxygen, nitrogen and so on consist of diatomic molecules. He had no developed theory of chemical bonding to account for how compounds formed. In its absence, he determined that water is a compounding of *one* atom of hydrogen with *one* atom of oxygen. In modern language, that amounts to the formula HO and not the familiar H_2O.

Methods dependent on this factually erroneous portrayal of materials are quite unreliable. A specific example are the rules he developed in Part I, p. 214, to determine the ratios of atoms in each compound. Those rules were responsible for the determination of water as HO. They enjoin us to use the simplest ratios possible; and lead directly to the one-to-one ratio of hydrogen to oxygen.[23]

With methods so unreliable, there can be no expectation that Dalton could secure the credit for introducing modern atomism to chemistry. My sense is that the key factor lay in his timing. The idea that all matter is atomic has been proposed routinely since antiquity. All such proposals were ill-fated prior to Lavoisier's epochal determination of the elements in his *Elements of Chemistry* (1790). After it, success is assured for the first chemist to propose an atomic constitution for the chemical element, no matter how flawed their other results may be.

4.3 Carnot's Discovery of the Foundations of Thermodynamics

Sadi Carnot's (1824) *Reflections on the Motive Power of Fire*, is, in my view, one of the greatest works of science of all eras. Its pre-eminence lies in it not just laying the foundations of the new science of thermodynamics, but in introducing a novel form of theorizing that is

[23] See Norton (2024, Ch. 11) for an account of the half century of work needed to establish the correct atomic weights and molecular formulae.

distinctive of thermodynamics.[24] He recognized that it is possible to have a general theory of efficiency covering heat engines of all types, not just steam engines, and no matter their operating fluids or other details. The maximum capacity of these heat engines to produce motive power for each unit of heat supplied was determined solely by the temperatures of the hot source from which it drew heat and cold sink to which it discharged heat. That maximum capacity was realized by reversible operations, that is, ones that could run equally in the forward or the reverse directions. The key insights leading to the later formulation of the second law of thermodynamics are routinely and correctly attributed to Carnot's work.

Nonetheless, Carnot's successes are a Gettier-like case. For Carnot's analysis was carried out within the caloric theory of heat, using methods that routinely led to incorrect results. In the later thermodynamic theory, a heat engine *converts* the energy of heat into work energy. In Carnot's analysis, caloric, that is heat, is a conserved substance. There is no conversion of heat to work. It is solely the *transfer* of heat from hot to cold that engenders motive power. It is then inevitable that much of Carnot's analysis produced results that we now think incorrect. Most notably, because he assumed that caloric is conserved, the quantity of heat drawn from the heat source must match the quantity discharged to the heat sink.

How did Carnot achieve his successes using an understanding of heat engines that is so at variance with the modern view? Here are the saving factors he might not have foreseen. The general thermodynamic principles that were carried forward to modern thermodynamics prove to be insensitive to the particular assumptions that Carnot made erroneously about heat and heat engines. A quite specific one of his erroneous assumptions made his successes possible. For, if we assume that heat engines convert heat energy into work, it would be natural to conclude that an engine that converts heat fully into work is the most efficient. There would be no point in including a heat sink for waste heat in the analysis of the idealized, most efficient engine. Modern thermodynamics recognizes that no heat engine can operate without a colder heat sink to which to discharge waste heat. It is a necessity of its operation and, without it, an analysis of the style of Carnot cannot be carried out. Carnot's assumption about the conservation of caloric, however, requires the heat sink and made it a fixture in the analyses of later thermodynamicists.

[24] For an analysis of how he was able to achieve this feat, see Norton (2022).

4.4 Thomson's Discovery of the Electron

J. J. Thomson (1897) reported the results of his experiments on cathode rays. They were, he announced, beams of negatively charged particles. This announcement is now celebrated as the discovery of the electron, the first particle of the many to come in modern particle physics. In his paper, Thomson sought to determine whether these cathode rays are waves propagating in the ether, as Hertz, Lenard and others had supposed; or whether they are streams of charged particles. Thomson found that the rays are deflected by electric and magnetic fields exactly as would a stream of particles, if each particle carried a negative electric charge. The deflections allowed him to determine a unique mass to charge ratio that is characteristic of the particles that would soon come to be called "electrons."

The case against wave propagation specifically was laid out elsewhere in Thomson's writings. (For a more detailed narrative of the evidential arguments deciding between waves and particles, see Norton, 2021, Ch. 9.) The decisive consideration was that cathode rays are deflected by a homogeneous magnetic field. There was no mechanism in the theory of waves in the ether that would allow this. Charged particles, however, are deflected by such a homogenous field by the Lorentz force law.

Thomson's methods were those dictated by late 19th century electrodynamics. Using them, he made his case well. What he could not foresee was that the quantum theory of the early 20th century would find cathode rays to be the propagation of a quantum electron wave. His case for electrons as localized particles and not waves collapsed. A quantum electron wave can be deflected by a homogeneous magnetic field. Thomson's 19th century methods were ill-adapted to the quantum physics of the 20th century.

However, in spite of the failings of his methods, enough of his conclusions for his particles survived. The unanticipated saving factor is that the deflections he attributed to the behavior of particles were recovered in the quantum theory by Ehrenfest's theorem. It showed that these are the motions of a quantum wave, if the quantum wave is sufficiently localized in space. That then allowed the mass and charge Thomson had determined for his particles to become the corresponding parameters of the quantum electron.

4.5 Bohr's Discovery of the Quantum Character of Atoms

Niels Bohr's (1913) "On the Constitution of Atoms and Molecules" depicts hydrogen and other atoms as consisting of positively charged nuclei, surrounded by one or more electrons in stable, discretely spaced energy states. The electrons can transition up and down between these stable energy states by absorbing or emitting the energy difference between the states in electromagnetic waves. These results provided the basis for the quantum theory of atoms that built upon Bohr's findings.

Even though the results just described are foundational in the later quantum theory of atoms, Bohr's methods were quite ill-adapted to the new quantum theory. His analysis employed the methods and results of 19th century electrodynamics. Bohr's electrons were not the quantum waves that surround the nucleus of an atom in diffuse clouds. His were still the localized particles of Thomson and they orbited in circular or elliptical orbits. Bohr's electrons jumped instantaneously between their stable energy states. Their transitions were not governed by the Schrödinger evolution of the yet to be developed quantum electrodynamics. The electromagnetic energy absorbed and emitted by Bohr's electrons were in the form of classical waves, not as the light quanta Einstein had introduced in 1905 that became the foundation of the quantum treatment of electromagnetic systems.

Worse, Bohr's 19th century methods were employed inconsistently. The classical theory insisted that an orbiting charge could not adopt a stable state, but would spiral into the atomic nucleus. Bohr posited stability. The classical theory allowed that an electron could adopt a continuous set of energy states in its orbiting the atomic nucleus. On the strength of the evidence of discrete atomic spectra, Bohr inferred otherwise.

In sum, Bohr's methods were unreliable when judged by the later quantum theory. Most of his results were incorrect. But within the compendium of false results lay a few that became the foundation of the new quantum theory of atoms. These were the ones that reflected the characteristics of the atomic spectra that formed the empirical basis of his theory: the stability and energetic discreteness of his electron states reflected the stability and discreteness of the observed atomic spectra. The unanticipated saving factor is that just these were the results that would survive into the new science to come.

4.6 Einstein's Discovery of the General Theory of Relativity

For over seven years, from 1907 to 1915, Einstein expended at times exhaustive efforts to generalize his special theory of relativity. The result was his general theory of relativity. It was celebrated then and still today as our most successful and most profound relativistic theory of gravity. Private notes, correspondence and intermediate publications have provided us extensive documentation of the stages of Einstein's thinking and the methods he used. On his later summary, his methods were controlled by three principles: the idea that the relativity of motion was to be extended to acceleration; Mach's principle, that the inertia of a body results from an interaction with the other masses of the universe; and the principle of equivalence, which asserted the identity of gravity and acceleration.

What makes Einstein's discovery Gettier-like is that two of these ideas, so foundational to his discovery, prove to be incompatible with his final theory. His theory does not generalize the relativity of motion to acceleration; and his rewriting of the generalized principle as a principle of general covariance was shown by Kretschmann in 1917 to be vacuous. Mach's principle also failed in the theory. Einstein disowned it for its dependence on an ontology of interacting masses at variance with the notion of unified field theory that dominated his later thinking. In so far as Einstein's methods were controlled by these principles, we could not now expect them to lead him successfully to his final theory. Yet he did arrive at the theory.[25]

In retrospect, I believe his success depended at least in part on his insistence on further conditions that are sufficiently unremarkable as to elude inclusion in celebratory histories, but are, in retrospect, powerful guides toward the final theory. I have explored some of them in Norton (2020). During the crucial period of the discovery of the theory, Einstein could not have anticipated that these were the saving factors and that the principles he prized would fail in the theory.

4.7 Hubble's Discovery of the Hubble Law

Hubble's (1929) celebrated discovery of the recession of the galaxies was recounted in *small-e empiricism*, Chapter 10. Its main result is a linear relationship between the velocity of galactic recession and the distance to the galaxy. In retrospect, his results were a mix of successes and failures. The linearity of the relation survives and carries his name. Hubble's determination of the constant of proportionality, 500 km/sec/megaparsec, proved to be 6 or 7 times greater than the later value of roughly 70 km/sec/megaparsec.

[25] For my early recounting of Einstein's discovery, based on private calculations in his Zurich notebook, see Norton (1984). For more details of the incompatibility of these principles with his theory, see Norton (1993).

The analysis in *small-e empiricism*, Chapter 10, recounts where Hubble's methods failed in the determination of the constant. The difficulty was that he did not know that there is more than one type of Cepheid variable star that formed the basis of his estimates of the distanced to the galaxies. He had presumed the wrong type in estimating distances to nearby galaxies and this compromised all this remaining estimates. The saving factor that Hubble could not anticipate was a result of later investigations. There is only one other type of Cepheid to be considered. Accommodating it into the analysis merely requires a recalibration of Hubble's distance estimates by a constant factor. The correction altered the constant, but not the celebrated linearity of Hubble's law.

5. Conclusion

It is interesting to reflect on the different receptions received by Gettier cases in epistemology and by Gettier-like cases in science. In the first, Gettier cases are alarming. They are taken to reveal a profound flaw in a central and widely used definition of knowledge as justified, true belief. The challenge, which has been met poorly, is to find a way to defend this definition.

In science, Gettier-like cases are regarded as benign and even inevitable. They are recognized as cases of scientists proceeding responsibly with the best methods then available. What matters is the fact discovered, even if it was by methods that are in retrospect judged to be flawed. For such flaws can be identified and corrected and the discovery can be given a more secure foundation. What does not matter is whether the original discovery conformed with some epistemologist's definition of knowledge. Progress in science is, at any moment imperfect. Over time these imperfections can be corrected, so that our science rests on ever more secure foundations. This is science progressing responsibly within the bounds of fallibilism.

References

Berry, Arthur (1898) *A Short History of Astronomy*. London: John Murray.

Bohr, Niels (1913) "I. On the Constitution of Atoms and Molecules," *Philosophical Magazine* Series 6, **26**, No. 151, pp. 1-25.

Bradley, James (1729). "An account of a new discovered motion of the fixed stars." *Philosophical Transactions of the Royal Society*, **35**, pp. 637–661.

Carnot, Sadi (1824). *Réflexions sur la Puissance Motrice du Feu*. Paris: Bachelier. Trans., *Reflections on the Motive Power of Heat*. Wiley. 1897.

Dalton, John (1808) *A New System of Chemical Philosophy*. Part I. Manchester: R. Bickerstaff.

Dalton, John (1810) *A New System of Chemical Philosophy*. Part II. Manchester: R. Bickerstaff.

de Grefte, Job (2023) "Knowledge as Justified True Belief" *Erkenntnis*, **88**, pp. 531–549.

Hubble, Edwin (1929) "A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae." *Proceedings of the National Academy of Sciences*, **15**, pp.168–73.

Lavoisier, Antoine-Laurent (1790) *Elements of Chemistry.* Trans. Robert Kerr. Edinburgh: William Creech.

Norton, John D. (1984) "How Einstein Found His Field Equations: 1912-1915," *Historical Studies in the Physical Sciences*, **14,** pp. 253-315.

Norton, John D. (1993) "General Covariance and the Foundations of General Relativity: Eight Decades of Dispute," *Reports on Progress in Physics*, **56**, pp.791-858.

Norton, John D. (2020) "Einstein's Conflicting Heuristics: The Discovery of General Relativity," pp. 17-48 in *Thinking about Space and Time*: *100 Years of Applying and Interpreting General Relativity*. *Einstein Studies*, Volume 15. C. Beisbart, T. Sauer, C. Wüthrich (eds). Cham, Switzerland: Birhäuser/Springer Nature.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Norton, John D. (2022) "How Analogy Helped Create the New Science of Thermodynamics," *Synthese*. **200**, article 269, pp. 1-42.

Norton, John D. (2024) *The Large-Scale Structure of Inductive Inference.* BSPS*Open*/University of Calgary Press.

Thomson, Joseph J. (1897) "Cathode Rays," *Philosophical Magazine,* **44,** pp. 293–316.

7 Small-r realism

Chapter 7

Small-r realism[26]

1. Introduction

The debate between scientific realism and antirealism has held a central position in the philosophy of science since the middle of the twentieth century. It is driven by two competing intuitions. Realists are guided by a triumphalist sentiment that our best science has enjoyed extraordinary success in discerning contingent truths of the world, where those truths extend from the prosaic to the profound. Antirealists regard such triumphalism as overreaching and counsel a more skeptical attitude to science's putative successes.

Something like the scientific realist intuition is correct, in my view. The long-standing problem has been to find a precise and sustainable formulation of it. This chapter seeks to establish that small-e empiricism solves the problem. It is already the strongest formulation of scientific realism that is admissible within the bounds of the fallibility of inductive inference. Since it concedes this fallibility to antirealists, it amounts to a view that successfully mediates between the extremes of an excessively optimistic realism and excessively pessimistic antirealism.

What will sustain this mediating view, I shall argue in Section 2 below, is that there are two parts to the scientific realists' intuition. The first is that we have the epistemic means to learn core, contingent propositions in a science. The second is that individual sciences have successfully employed these means. This first part only is amenable to a formulation as a general doctrine of scientific realism. Determining the truth of the second part lies beyond what a general doctrine can achieve. It must be determined in a case-by-case analysis. The enduring failure of formulations of scientific realism lies in unsuccessful attempts to find a general formulation that implements this second part.

Section 3 reports that small-e empiricism already embodies a formulation of scientific realism that conforms with this first part. The formulation, here called "small-r realism,"

[26] My thanks for Howard Sankey for helpful comments on this chapter and the two following.

identifies inductive inference as the uniquely available epistemic means. Crucially, small-r realism makes no assertion that any particular science has attained a final truth. Rather it asserts that sciences have the means to do this through inductive inferences from experience. The extent to which any particular science has succeeded in applying the means is left to the individual science. The complications of the fallibility of inductive inference are thereby relocated to where they belong: in the case-by-case analysis of the inductive support of the individual sciences. This particularism extends to how strongly supported propositions are to be interpreted. Whether such a proposition should be interpreted literally, as some versions of scientific realism require, should also be left to a case-by-case analysis. The dangers of assuming otherwise are illustrated by the notorious complications of interpreting the meaning of entropy.

The resulting small-r realism is essentially coincident with small-e empiricism. It adds only a reminder of the fallibility of inductive inference. Thus, as Section 4 reviews, small-e empiricism/small-r realism moderates the extremes of empiricism and realism into a single, mediating position that preserves the best of both views.

Section 5 contains a short, summary conclusion.

2. How to Formulate Scientific Realism: Escaping Fallibilism

2.1 A Division of Labor

How can we take the triumphalist intuition that science has successfully discerned deep truths of nature and give it a precise formulation? A primary concern is that the formulation should accommodate the fallibility of inductive inference. My sense is that the literature in scientific realism has met only with limited success in this task. Common formulations tend to skirt around a successful formulation without getting it precisely right.

The key to a successful formulation is to recognize that the intuition underlying scientific realism has two parts:

Core Intuition of Scientific Realism

• Our epistemic resources are powerful enough to inform us of deeper truths of science.

• Our present sciences have used these resources effectively to learn these truths.

The first part only is amenable to a formulation of scientific realism as a general philosophical doctrine. For it has to express something general and robust through time that will apply across all the sciences. The second part depends on the specific state of individual sciences at specific times. Whether a particular science at a specific time meets the conditions of this second part must be determined on a case-by-case basis. It is not within the scope of a general philosophical doctrine. The case-by-case analyses must be carried out in the individual sciences themselves or in the history of science.

2.2 Inductive Support of Experience as the Means

Small-e empiricism supplies a single, unique means as the epistemic resource of the first part. It is inductive inference; and small-e empiricism affirms that this resource does have the capacity to inform us of these deeper truths. That inductive inference is fallible does not undo this capacity. Rather, it just means that the support might err, but that the chance of an error can be made very small, when the inductive support is very strong.

Whether the inductive support for some contingent proposition in science has failed is not something that can be determined by general principles. In the case of very strong inductive support, such failure is very unlikely, but can happen. If we have very many contingent propositions, all enjoying strong support, then we can still expect that one among the very many will fail. As Chapter 4 above noted, that is as much as can be said at the level of general principles and as much as can be incorporated into a general formulation of scientific realism. Anything further must be handled on a case-by-case basis by an examination of the particular circumstances of the contingent proposition at issue. Commonly, we find in history of science that such unlikely failures are revealed by the discovery of new evidence. As much as centuries of time may pass before this new evidence emerges.

Where existing formulations of scientific realism err is that they try to find some way of incorporating both parts of the initial intuition into their formulations. They are, in effect, engaged in a futile attempt to evade the fallibility of inductive inference. We shall review some of them below. Briefly, these formulations may attribute not truth but a nebulous "approximate truth" to well-supported contingent propositions. Or they may try to select out that part of a present theory that will endure, while being unable unambiguously to identify precisely which those parts may be. Or they may speculate on future successes. If a science has not secured the truth today, we are assured that it is on the way to the truth and that scientific progress will take us there eventually.

2.3 Early forms of Scientific Realism

Writing in philosophy of science in the mid-twentieth century started to turn away from antirealist philosophies, such as was reported in *small-e empiricism*, Chapter 3, Section 2.5, had been pressed by the Vienna Circle. Reichenbach's logical empiricism was an early phase of this rise. We saw in *small-e empiricism*, Chapter 3, Section 3.4, that it included a sharp repudiation of the logical positivists' rejection of both realism and antirealism, as then conceived. It derived from Reichenbach's assigning of probabilistic weights to propositions. They included propositions expressing sense impressions, as well as those expressing facts remote from experience, such as the temperature of the interior of the sun. These last propositions allowed Reichenbach to formulate his "realistic conception of the world," which asserts that the external world has an existence independent from our impressions. This is a version of the division of labor urged here. For Reichenbach's claim is that his probability theory of meaning gives us the means to affirm probabilistic facts in the external world, without specifying further which those facts are.

Other versions of this division of labor appear in formulations of scientific realism from the mid twentieth century, at the outset of the present debate. For example, Boyd (1973, pp.1, 3) gave a formulation marred only by its inclusion of an appeal to approximate truth:

> By scientific realism I mean the doctrine that the sort of evidence which ordinarily counts in favor of the acceptance of a scientific law or theory is, ordinarily, evidence for the (at least approximate) truth of the law or theory as an account of the causal relations obtaining between the entities quantified over in the law or theory in question.

and

> … experimental evidence for a theory is evidence for the truth of even its non-observational laws and, hence, for the truth of observational predictions deduced from them, since deductive inference preserves truth.

The idea of a division of labor between philosophers and scientists is clearly articulated in J. J. C. Smart's (1963, p.27, his emphasis) monograph on scientific realism. He defends what we would now call an "entity realism" and notes:

> I shall try to uphold a realistic view of these theoretical entities, and will try to establish that the elementary particles of physics are just as respectable entities as

tables and galvanometers. By 'just as respectable' here I mean '*philosophically* just as respectable'. Unicorns are philosophically just as respectable as cows: it is for the zoologist, not the philosopher, to decide that there are no unicorns. Similarly, the physicist may come to replace his present theory of elementary particles with some other theory of what goes on at the sub-atomic level. Nevertheless, he will presumably postulate some new set of theoretical entities to replace the old ones. It is not for the philosopher to decide between one physical theory and another any more than it is up to him to become a zoologist and decide for or against the existence of unicorns. My problem is concerned with the ontological status of the physicist's theoretical entities, whatever they are.

3. small-r realism

3.1 The Formulation

These last reflections delineate a version of scientific realism that lies within the existing commitments of small-e empiricism. In more compact form, it is:[27]

> *small-r realism*: the propositions of experience can provide inductive support for the truth of the contingent propositions of the sciences. This inductive support can be very strong, but can never rise to absolute certainty; and no other means can provide that certainty.

This formulation is largely a restatement of the core doctrine of small-e empiricism, as formulated in Chapter 2 above. The addition is an explicit acknowledgment of the fallibility of inductive inference, as explored in Chapter 4. It is merely a reminder of what is part of the definition of inductive inference, its fallibility.

Small-r realism simply affirms that inductive inferences are a sufficient means for securing, fallibly, contingent propositions of a science. The doctrine cannot and should not assert that any of our latest science has actually secured such propositions. For there are many sciences and some are in active evolution.

A case-by-case analysis of individual sciences is encouraging. Many have secured very strong evidential support for their core, contingent propositions. If we add this fact to small-r

[27] These contingent propositions include what are commonly called scientific laws, for, within small-e empiricism, they are just contingent propositions of very broad scope.

realism we arrive at the optimistic result amenable to scientific realists: the world is most likely truly as these sciences say. We have strong inductive support for the contingent proposition in cosmology of the expansion of the universe from a big bang. We infer that it is very likely that this expansion truly is the case. This optimistic result can be recovered for central propositions drawn from across the sciences. It cannot be recovered universally. Evidential support in some sciences is weak. No theory of quantum gravity enjoys evidential support comparable in strength to that of, say, ordinary chemistry. All such theories must be treated as speculative.

3.2 Interpretive Autonomy of Science

An important consequence of small-r realism/small-e empiricism is the interpretive autonomy of science. The autonomy requires the failure or superfluity of any general principles of interpretation, such as a realist's blanket injunction to read the propositions of mature sciences literally or an antirealist's blanket injunction that these propositions cannot be read literally. For how such propositions are to be interpreted is itself a factual matter for the science. Small-r realism/small-e empiricism asserts that our only means of assessing all such factual matters is by the inductive support of experience.

We may conceive a general principle of interpretation that goes beyond what the inductive support of experience supplies. Since the inductive support of experience gives us all the support we can have, the principle will fail, in so far as it goes beyond what inductive inference can support. If it does not go beyond what inductive inference can support, then it is superfluous. Anything it can contribute has already been contributed by the inductive support of the propositions from experience.

Nothing fancy is intended by the term "interpretation" in the interpretive autonomy of science. It is just what we take the propositions of a science to mean in a quite ordinary sense. For mature science, it is generally unproblematic to take the interpretation simply to be what the propositions assert literally. When astronomy tells us that our planetary system is heliocentric, it tells us just what is says, that the orbits of the planets are centered on the sun.

Although this is true generally, it is not universally true. We must always allow for the possibility that the literal interpretation is incorrect. A straightforward case arises with effective theories. An example is the theory that gases and liquids are continuous fluids whose behavior is governed by the Navier-Stokes equation. It is very well supported inductively by experience. However further experience tells us that the theory is only true of averaged behaviors, after the molecular granularity of gases and liquids has been smoothed out. The theory's assertion of continuity is not a literal truth, since a literal truth requires the infinite divisibility of the fluid. The inductively well-supported atomic theory of matter denies this infinite divisibility.

3.3 The Interpretation of Entropy

One presently intractable example is enough to illustrate the delicacy of problem. When water boils to steam at 100C and 1 atmosphere pressure (101.42 kPa), its internal energy per unit mass u changes by[28]

$$\Delta u = 2{,}087 \text{ kJ/kg};$$

and its entropy per unit mass s changes by

$$\Delta s = 6.0471 \text{ kJ/(kg.K)}.$$

Both of these quantities, u and s, are core theoretical terms in thermodynamics. The propositions asserting both changes are very strongly supported by experimental measurements on the more directly measurable properties of water and steam, such as their temperatures, pressures and densities, and by the measurable quantities of heat supplied to them. The entropy change Δs is computed[29] as the heat supplied in (reversibly) vaporizing the water to steam Δh divided by the absolute temperature T, $\Delta s = \Delta h\, /T$, where $T = 373.15$K.

If approached at a more general level, the similar propositions asserting the values of Δu and Δs would naturally both be interpreted literally. A literal interpretation of the value of Δu is strongly supported by all the evidence that has repeatedly affirmed that energy is a conserved quantity passed between materials undergoing changes. It is a commonplace of thermodynamics that the energy content of the water increases as indicated, whether the change is considered phenomenologically or statistically in light of the molecular constitution of water.

---

[28] Data from Parry *et al.* (2014) Table S-1. It gives the value of Δs directly and the enthalpy change per unit mass $\Delta h = 2{,}256.5$ kJ/kg. The internal energy change per unit mass Δu is calculated from $\Delta h = \Delta u + p\Delta v$, using the pressure $p = 101.42$ kPa and the volume change per unit mass $\Delta v = 1.6708$ m^3/kg.

[29] The heat of vaporization is given by the enthalpy change $\Delta h = \Delta u + p\Delta v$, where Δu is the internal energy increase of the steam and $p\Delta v$ is the energy lost to work in expanding the steam by a volume Δv against the resisting pressure p.

Tempting as it is, no similar literal interpretation of the value of the entropy change Δs is sustainable. The question of just how we are to think of entropy has remained troublesome since its introduction by Clausius in 1865. Is it a real physical quantity like internal energy? Is it a measure of accessible energy? Is it a merely instrumental tool for gauging the temporal direction of processes? Is it a purely effective magnitude resulting from our limited control of molecular-scale processes? Is it a parameter of the probability distribution for energy over the system's components? Is it a measure of the probability of a state or of the randomness of a state? Is it a measure of information in the spirit of Shannon's information theory? Or is it a quantity best understood by conceiving thermodynamic processes as computations? These are open issues over which debate endures.[30]

The proposition Δs = 6.0471 kJ/(kg.K) is as strongly supported by the evidence and as good a candidate for truth as any in thermodynamics. It has been an unchallengeable result since the nineteenth century. However, there is no agreement on what is its literal construal.

4. A Single Mediating Position

Small-r realism provides a view that mediates as successfully as possible between the strictest versions of scientific realism and antirealism. It preserves as much of the original intuitions behind scientific realism, while discarding those parts incompatible with the fallibility of inductive inference. In not endorsing the blanket truth of our mature sciences, it concedes to some of reservations of antirealism, while not succumbing to a pessimistic overreaction to inductive fallibilism. That the commitments of small-r realism coincide with those of small-e empiricism is a welcome result. For it tells us that together that they provide a single view that preserves what is valuable in all of the existing doctrines of empiricism, scientific realism and antirealism.

Here is how small-r realism relates to Big-R Realism and to antirealism.

4.1 Not Big-R Realism

By a Big-R Realism, I mean a version of scientific realism that proposes, as a general result, that our best science has in some way discovered the final truths of the science. It may be a proposal of structural realism that asserts an enduring truth in the identification of structure; or

[30] I have written extensively in the debate. See Norton (2025) for a recent contribution.

it may be a version that asserts an equally enduring “approximate truth” of some contingent proposition.

Small-r realism is close in spirit to Big-R Realism in its epistemic optimism. The methods of science are fully capable of providing strong support for the truth of the core propositions of a science. Small-r realism differs from a Big-R Realism in affirming only the capacity of inductive inference from experience to secure the truth of contingent propositions central to a scientific theory. Whether a science has succeeded in securing such strong support must be determined on a case-by-case basis. Even ardent Big-R Realists must accept that some sciences are in a rudimentary state and have not secured this strong support. To sustain itself, Big-R Realism must provide a way to separate out those sciences or those parts of the sciences that have secured the final truth from those that have not. This amounts to the need for some generally applicable principle that tells us when the truth of a contingent proposition has been determined with finality.

According to small-e empiricism, such a principle is unattainable. The best we can have is very strong, but fallible inductive support for contingent propositions. The failure of attempts to work around this limit is assured. It is just a matter of examining these attempts to discern how they fail. Claims of approximate truth founder on the ambiguity of the adjective “approximate” when applied to “truth.” Claims of enduring structure require something that has yet to be found: a precise definition of structure such that this structure of a present scientific theory can be affirmed to survive through all future discoveries.

These last remarks should not be understood as a form of skepticism. They are entirely compatible with the possibility of the contingent propositions of a science enjoying such strong inductive support that the falsity of their central propositions is well-nigh inconceivable. Such is the case, I believe, with many of our best sciences. Slender as it may be, the possibility of falsity remains as a matter of principle, even if not one of any practical import.

4.2 Not Antirealism

By antirealism I mean those positions that assert that science can secure truth only for contingent propositions that express experience directly or, commonly, possible experiences, but not for propositions that go beyond experience. These latter propositions may be portrayed as merely of instrumental or pragmatic utility. The view is not just that science happens to have failed to secure the truth of the latter. Rather, it asserts that to secure it lies beyond the methods available to science as a matter of foundational principle.

Small-r realism conforms with antirealism in not automatically declaring the truth of core propositions of mature sciences. It differs from antirealism in treating alike the import of inductive support for the contingent propositions of a science. Small-r realism has to treat them alike since it has no means to distinguish them. small-r realism/small-e empiricism has abandoned the long dubious bifurcation of experience and theory, as reported in Chapter 2.

Antirealism asserts a one-size-fits-all blanket view of skepticism towards any propositions beyond experience. In small-r realism, all that matters is the extent of the inductive support for a contingent proposition. It does not even have a general principle that accords specific observational propositions greater support than the propositions of the deeper theory to which the observations pertain. Decisions concerning comparative strengths of support must be made on a case-by-case basis.

Commonly, individual observation reports are better supported than their theoretical generalizations. Examples are easy to find. In the 1990s, Schilling and collaborators (1993) had observed superconductivity at the highest temperature (under normal pressures) yet found, 133K, in their sample of Hg-Ba-Ca-Cu-O. Given the continuing progress in finding higher temperature superconducting material, this observational report is better supported inductively than the generalization that all superconduction arises at normal pressures below 135K.

Inversions are equally easy to find, in which the strengths of support for particular experiential claims are weaker that those of the theoretical claims that govern them. For example, the evidential support for the truth of the conservation of energy is far stronger than for the veracity of any particular report of a perpetual motion machine. The evidential support for the standard chronology of species evolution is far stronger than for the veracity of a report of coincident fossilized human and dinosaur footprints.

5. Conclusion

The present literature in philosophy of science depicts empiricism and scientific realism as irreconcilable opposites, largely I believe through the influence of van Fraassen's *Scientific Image*. It is an unfortunate conceptualization since empiricism and realism each have something of value. Each, however, also harbors elements that should be discarded. Big-E Empiricism endorses an indefensible inductive skepticism. Big-R Realism underestimates the intractability of the inductive fallibility.

My goal in this chapter has been to show that, if we discard these troublesome elements in each, we arrive at a welcome consonance between the two views. Small-e empiricism discards inductive skepticism. Small-r realism accepts that inductive fallibilism is inescapable. What results in small-e empiricism and small-r realism is a single, optimal, mediating position that preserves the best of both extremes and discards the worst.

References

Boyd, Richard N. (1973) "Realism, Underdetermination, and a Causal Theory of Evidence," *Noûs*, **7**, pp. 1-12.

Norton, John D. (2025) "The Simply Uninformed Thermodynamics of Erasure," *Philosophy of Physics*, 3(1): 3, 1–21.

Parry, William T. *et al.* (2014) *ASME International Steam Tables for Industrial Use*. New York: The American Society of Mechanical Engineers.

Schilling, A.; Cantoni, M.; Guo, J. D.; and Ott, H. R. (1993). "Superconductivity above 130K in the Hg–Ba–Ca–Cu–O system," *Nature*. **363**, pp. 56–58.

Smart, J. J. C. (1963) *Philosophy and Scientific Realism*. London: Routledge and Kegan Paul.

8. The Proliferation of Scientific Realisms

Chapter 8

The Proliferation of Scientific Realisms

1. Introduction

The original motivating intuition of scientific realism was narrow. It was to seek a precise formulation of the optimistic view that our science has enjoyed an unparalleled success. A good starting point might just assert that the central propositions of our best sciences are true or likely so. Unfortunately, the literature that developed out of the renewed mid-twentieth century interest in scientific realism has not maintained a similar narrowness. Formulations of scientific realism have rapidly multiplied into the expanding disarray of competing senses now filling the present literature. This proliferation all but guarantees an enduring and unresolvable debate. We cannot agree on a doctrine, whether it is right or wrong, if we cannot first decide precisely what it is.

With this problem in view, small-r realism has been formulated austerely to capture as precisely as possible what is correct and what is viable in the original intuition, while still accommodating the fallibility of inductive inference. It is offered as a corrective to the present surplus of competing formulations. It reflects the need to contract the definition of scientific realism to something narrow and to abandon the permissiveness that allows everyone to have their own proprietary version of scientific realism. That is not to deny everyone a place to state their views. It is to deny a blanket license that allows the term "scientific realism" to have any meaning anyone may choose. It is to insist that the term "scientific realism" should have a simple, univocal meaning. If we are to identify as such a formulation of scientific realism, it is, I maintain, small-r realism.

Section 2 offers a short review of the bewildering proliferation of views in the present literature that carry the label "realism." Their great variety makes it almost impossible to give a simple, univocal answer to the question of what is scientific realism.

Section 3 introduces the project of the remainder of the chapter. It is to assess how well the realism literature accommodates the core concept of small-e empiricism/small-r realism: inductive inference. Strong forms of realism attempt to discount the import of inductive fallibility. This investigation reaffirms small-r realism as the formulation of realism that escapes this problem.

The review of optimistic attempts by scientific realists to diminish the import of inductive fallibilism begins in Section 4. One strategy employs the notion of "approximate truth." It fails through its opacity. To declare a theory approximately true tells us nothing about where the theory fails. Similarly, we derive only scant protection from inductive fallibilism by declaring exceptions as inconceivable monsters, since what is inconceivable today may become conceivable tomorrow, even if that happens only very rarely.

Section 5 reviews three versions of selective realism: structural, entity and explanationist (indispensabilist) realisms. Each asserts that a realistic construal should be reserved for selected elements of a theory. They risk improperly breaching the limits of inductive inference in two ways. First, they discount the import of inductive inference for those parts of a theory not selected; and second they exceed what inductive inference can provide with assurances of truth for those parts selected. In so far as these selective realisms breach the limits of inductive inference, they are unsustainable. In so far as they do not breach them, they add nothing of principle to small-r realism.

The further difficulty of selective realism lies in the details. There are multiple proposals for which parts of a theory should be selected and some disagreement among the proposals. Structural and entity realists directly conflict. A still further problem is that the identification of the selected elements can be vague. There are many—too many—proposals for structure. Although the idea that some elements of a theory are indispensable to the theory's empirical success, there seems to be no way to identify which they are, when the theory is proposed. It can even be troublesome retrospectively, as the example of Carnot's 1824 memoir on heat engines shows.

Section 6 contains a short, summary conclusion.

2. Scientific Realisms: Expansion, Fragmentation and Drift

This section offers a brief survey of the development of conceptions of scientific realism. It illustrates how its conception has drifted, expanded and fragmented. Some versions are expansive and seek to combine many elements into a single synoptic view. Others merely identify one element as realism in a literature that has become increasingly fragmented. This is a literature urgently in need of repair. To repeat the main claim of this chapter, small-r realism offers such a repair. It implements, in a compact and narrow form, the original intuition underlying scientific realism.

2.1 Aims and Beliefs

The proliferation of differing senses began almost immediately. van Fraassen's early definition of scientific realism provoked a flood of responses and provided a focus for the growing realism debate. He gave it as (1980, p.8, his emphasis):

> *Science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true.* This is the correct statement of scientific realism.

It differs in important ways from the original intuition. Since the avowed purpose of van Fraassen's text was to refute scientific realism, the broad acceptance of van Fraassen's formulation strikes me as an enduring error on the part of scientific realists. How wise is it for scientific realists to allow their opponent to define their view? Such a definition will surely be tailored to serve the purposes of the opponent. Accordingly, van Fraassen included two notions that are not present in the original intuition of scientific realism and should not have been admitted. They are an assertion of the aim of science and an assertion about mental states, that is, beliefs. They were not new ideas in this literature, but van Fraassen's formulation gave them prominence as part of a sharp and short definition of scientific realism.

These additions diffused and diluted the definition of scientific realism. There are no aims without aim seekers and no beliefs without believers. Neither are required by the motivating intuition of scientific realism. They confound the definition with superfluous encumbrances. That science *aims*, or more accurately, scientists aim for something is no assertion of its success. We all seek things that, alas, we never achieve. We may even lack the means to realize our aims. That scientists *believe* something is a similar distraction. That someone happens to believe something is no assertion of its truth or even likelihood of truth. To incorporate the goals and mental states of scientists into the definition of scientific realism yields a defective definition whose presence persists in the literature to its detriment.

2.2 Proliferation

Early work on scientific realism was advanced by a conference on scientific realism held at the University of North Carolina, Greensboro, on March 26-28, 1982. A subsequent conference volume, edited by the conference organizer, Jarrett Leplin, identified ten versions of scientific realism in the literature of the time. They were: (Leplin, 1984, pp. 1-2)

1. The best current scientific theories are at least approximately true.
2. The central terms of the best current theories are genuinely referential.
3. The approximate truth of a scientific theory is sufficient explanation of its predictive success.
4. The (approximate) truth of a scientific theory is the only possible explanation of its predictive success.
5. A scientific theory may be approximately true even if referentially unsuccessful.
6. The history of at least the mature sciences shows progressive approximation to a true account of the physical world.
7. The theoretical claims of scientific theories are to be read literally, and so read are definitively true or false.
8. Scientific theories make genuine, existential claims.
9. The predictive success of a theory is evidence for the referential success of its central terms.
10. Science aims at a literally true account of the physical world, and its success is to be reckoned by its progress toward achieving this aim.

This list is reproduced here not because any of these formulations is especially successful. Rather, it is to illustrate just how scattered and fragmented were formulations of scientific realism, even in this early time. Leplin remarked "…scientific realism is a majority position whose advocates are so divided as to appear a minority."

Several factors encouraged continuation of this proliferation of notions of scientific realism. First is a disagreement on precisely what the doctrine should assert. Does it assert the truth of claims in a science? Their approximate truth? That their terms are referentially successful? That science aims for truth? That successive stages of science progress towards the truth? And so on.

Next, claims about the means through which these possibilities are realized are also included in the definitions. Is the truth of propositions secured through their explanatory success? Through their predictive prowess?

Next, the characterizations of scientific realism include negations of alternative positions. To assert that propositions of a science can be true or false would surely be superfluous, were it not that instrumentalists might assert otherwise. Similarly, had instrumentalists not asserted otherwise, there would be little point in including in a definition of realism that theoretical terms do more than merely correlate observables.

Finally, the term "realism" itself has a long history in philosophy. Two usages were already well settled in the nineteenth century and persist in later writing.[31] One sense of realism asserts that genus and species (in the nineteenth century formulation), or universals (in the twentieth century formulation), have a real existence. It stands in opposition to nominalists, who assert otherwise. In another, realism asserts the veracity of our ordinary perceptions, in contrast with idealism that asserts a mind-dependency of these perceptions. Both of these ideas, we shall see, diffused into formulations of scientific realism.

2.3 Synthesis

One reaction to this proliferation of senses of scientific realism was the attempt to synthesize them into a single, all-embracing doctrine. Psillos (1999, p. xvii) sought to represent scientific realism as the conjunction of three "stances," to use an unfortunate[32] term favored by van Fraassen. They were the metaphysical, semantic and epistemic and were described as:

1 The metaphysical stance asserts that the world has a definite and mind independent natural-kind structure.

2 The semantic stance takes scientific theories at face-value, seeing them as truth-conditioned descriptions of their intended domain, both observable and unobservable. Hence, they are capable of being true or false. Theoretical

---

[31] Both senses appear in the entry for "realism" in the nineteenth century *Vocabulary of Philosophy* (Fleming, 1860, p. 422); and in the entry for "realism" in the authoritative twentieth century *Encyclopedia of Philosophy* (Edwards, 1967, pp. 77-83). The latter reports multiple versions of realism: "naïve realism," "new realism," "perspective realism," "representative realism" and "critical realism."

[32] It is unfortunate since it depicts a thesis as a mental state with some suggestion of arbitrariness. In a debate, for example, we may *choose* to adopt a combative or a conciliatory stance. We saw in Chapter 4 that this freedom of choice was explicit in van Fraassen's voluntarism.

assertions are not reducible to claims about the behaviour of observables, nor are they merely instrumental devices for establishing connections between observables. The theoretical terms featuring in theories have putative factual reference. So, if scientific theories are true, the unobservable entities they posit populate the world.

3 The epistemic stance regards mature and predictively successful scientific theories as well-confirmed and approximately true of the world. So, the entities posited by them, or, at any rate, entities very similar to those posited, do inhabit the world.

Even though they extend well beyond the original motivating intuition of scientific realism, both of the older, nineteenth century senses of realism have been incorporated into expanded statements of scientific realism. Psillos' first statement ("…the world has a definite and mind independent natural-kind structure…") is a brief assertion of both.

At the same time, Niiniluoto (1999, p. 2) identified what he called "six different problems of realism" (p. 2). He selected among responses to them to produce the "critical scientific realism" the work was to defend. It is the conjunction of six theses, R0-R5: (p. 10)

At least part of reality is ontologically independent of human minds. (R0)

Truth is a semantical relation between language and reality. Its meaning is given by a modern (Tarskian) version of the correspondence theory, and its best indicator is given by systematic enquiry using the methods of science. (R1)

The concepts of truth and falsity are in principle applicable to all linguistic products of scientific enquiry, including observation reports, laws, and theories. In particular, claims about the existence of theoretical entities have a truth value. (R2)

Truth (together with some other epistemic utilities) is an essential aim of science. (R3)

Truth is not easily accessible or recognizable, and even our best theories can fail to be true. Nevertheless, it is possible to approach the truth, and to make rational assessments of such cognitive progress. (R4)

The best explanation for the practical success of science is the assumption that scientific theories in fact are approximately true or sufficiently close to the truth in the relevant respects. Hence, it is rational to believe that the use of the self-corrective methods of science in the long run has been, and will be, progressive in the cognitive sense. (R5)

Chakravarrty's (2007), as its title *A Metaphysics for Scientific Realism* suggests, is another attempt at synthesis he calls "semirealism." The view seeks to synthesize popular versions of scientific realism, notably entity and structural realism, along with a realism in metaphysics. Concepts in metaphysics that are to play a central role in semirealism include causation, disposition, laws and kinds.

From the perspective of small-e empiricism, there is little scope for a synthesis of scientific realism and metaphysical realism to add anything substantial to scientific realism. Whether something is real is a contingent, factual matter that can only be established by inductive inference from experience. We can attribute a distinct, factual reality to a metaphysical conception only because we have sufficient empirical evidence for it. That means that it has *already* been accommodated into empirically well-supported science. If the reality of a metaphysical conception has no such empirical support, then it is unwarranted and can add nothing factual to the realism.[33] This latter case is likely given Chakravarrty's (2024, p. 279) observation "… that metaphysical inquiry is largely a priori: metaphysical arguments proceed substantially via reflection on non-empirical considerations…"

2.4 Fragmentation

In the more recent literature, these different components are separated and variously offered individually as scientific realism. We can find just about any one of the above ideas or some combination of them identified as scientific realism. Here are some examples. They have not been collected systematically, but at least illustrate the spread of definitions.

> (Scientific) Realism. The theoretical entities of natural science do actually exist essentially as scientific theorizing characterizes them. They are real items of the world's furniture and do indeed possess substantially the descriptive constitution ascribed to them by science. (Rescher, 1987, p. xii)

[33] Chakravarrty responds to concerns of this type in his (2024, p. 281, his emphasis): "…one of the central claims I make about scientific ontology is that all approaches to it that interpret science as yielding some sort of knowledge of the world involve metaphysics in *some* way."

I propose what I take to be a standard construal of the scientific realist position as a form of the traditional metaphysical realist doctrine that the world exists independently of the mental. The realist position is a position of epistemic optimism, which holds against the sceptic that humans are able to acquire knowledge of the world. (Sankey, 2008, p.3)

Scientific realism is the view that our scientific theories are approximately true. (Dicken, 2016, p. 1)

Scientific realism may be crudely characterised as insisting that science tells us how the world is, in both its observable and unobservable aspects… (French, 2018, p. 394)

I will understand scientific realism in what I take to be a classical sense. According to scientific realism, the aim of science is to discover the truth about the objective reality which we inhabit. Progress in science consists in advance toward this aim. (Sankey, 2018, p. 72)

Scientific realism postulates that science aims for truth in both the domains of the observable and the unobservable, and is capable of achieving this aim, at least approximately.
According to scientific realism successive theories of modern science bring us progressively closer to the truth, that is, a depiction of nature's processes and entities as they actually are. (Dieks, 2023, p. 1)

Traditionally, scientific realists have taken the perfectly sensible view that the dramatic empirical and practical achievements of our best scientific theories show that those theories must be at least probably, approximately true. (Stanford, 2018, p. 212)

> We take scientific realism to be the doctrine that most of the statements of the mature scientific theories that scientists accept are true, or approximately true, whether the statement is about observable or unobservable states of affairs. Here, “true” is to be understood in a straightforward correspondence sense, as given by classical referential semantics. (Azhar and Butterfield, 2018, p. 304)

> There are semantic and epistemic variants of scientific realism in the literature. The semantic version holds that most successful theories are (nearly) true or (approximately) so. The epistemic version holds that successful theories are warranted. Let me call the former “semantic realism” and the latter “epistemic realism" hereafter.” (Park, 2022, p.1)

We can sympathize with Anjan Chakravartty who undertook the daunting task of surveying the range of present formulations of scientific realism. He wrote: (2017, Section 1)

> It is perhaps only a slight exaggeration to say that scientific realism is characterized differently by every author who discusses it, and this presents a challenge to anyone hoping to learn what it is.

The narrative continues with an attempt to inventory the many versions of scientific realism in the present literature:

> … most people define scientific realism in terms of the truth or approximate truth of scientific theories or certain aspects of theories. Some define it in terms of the successful reference of theoretical terms to things in the world, both observable and unobservable.
>
> Others define scientific realism not in terms of truth or reference, but in terms of belief in the ontology of scientific theories.
>
> Another way to think about scientific realism is in terms of the epistemic aims of scientific inquiry … . science aims to produce true descriptions of things in the world (or approximately true descriptions, or ones whose central terms successfully refer, and so on).

2.5 Type-X Realism

These last formulations of scientific realism are intended to represent a generic conception of scientific realism. Much of the more recent activity in the literature arises from a sense that the generic notion of scientific realism needs to be amended in ways that are no longer compatible with the plain label "scientific realism." This sense is indicated by the addition of an adjective that I label as an addition of "X" to "realism." Each instance of an X flags that the new proposal is more than a simple attempt to capture the generic idea. It is in some way correcting it. The correction may be a narrowing of the scope of scientific realism. Or it may be a migration to a neighboring conception that may not be a realism after all, but is a form of antirealism. As a group they are a vote of no confidence in the generic formulations for, they do not allow those formulations to stand.

The most common form of "type-X realism" is "selective realism." It is a class of views in which the assertion of reality is restricted to some elements within a successful scientific theory. They include "entity realism," "structural realism" and "explanationist realism." I will return to them below in Section 10. There are, however, views carrying the label of type-X realism that are antirealisms. They are reviewed in Section 8 below.

3. Induction in the Realism-Antirealism Debate

We have just seen that there is a multiplicity of different views concerning realism and its possible modes of failure. There can be no simple diagnosis for why this multiplicity has come about. The issues raised by realism or its failure cut to the heart of our understanding of science as an epistemic enterprise. They touch on many views and approaches and thus will elicit many responses. My concern in this chapter and the next is to consider the realism-antirealism debate in the context of small-e empiricism, where inductive inference plays a central role. The issue that arises repeatedly in the realism-antirealism debate is the use of inductive inferences and, in particular, the accommodation of inductive fallibilism.

The main claim of the following sections is that there is a misconstrual of the import of inductive fallibilism in the scientific realism literature. (The next chapter will argue for a misconstrual again in the antirealism literature.) Here we see a repeat of the misconstrual of inductive fallibilism in philosophy that was explored in an earlier chapter. Big-R Realists seek to discount inductive fallibility, assuming that it can be controlled by careful adjustments to the formulation of scientific realism. They are mistaken.

4. The Search for Permanent Truths

Strong versions of scientific realism underestimate the import of the fallibility of inductive inference. They treat its fallibility as an obstacle that can be surmounted by more careful formulations of the doctrine of scientific realism. These versions seek an assurance in the correctness of aspects of scientific theories that transcend the limits arising from the fallibility of inductive inference.

All these attempts fail for a reason that is foundational for small-e empiricism: our only means of ascertaining the truth of a contingent proposition is through inductive inference. As a quite general matter, attempts to secure assurances of truth that go beyond the fallibility of inductive inference necessarily fail.

This section reviews two attempts to tame the fallibility of inductive inference by adjustments to the notion of truth or of what we can accept as true. The first, the idea of approximate truth, fails because of its opacity. To declare a proposition as an approximate truth tells us only that something is right and something is wrong with it, but not which part is right and which is not. The second "monster-barring" declares certain possibilities as so extremely unlikely as to be beyond serious consideration. The difficulty with it is that what is inconceivable today, may be quite conceivable tomorrow, even if that eventuality is rare.

4.1 Approximate Truth

One of the most common attempts to evade the fallibility of inductive inference is to replace a claim of truth for some favored theory with one of approximate truth. The approach reflects the intuition that a stable and successful scientific theory can only be so if it has gotten something right. However, in an accommodation of the pessimistic meta-induction, realists are loath to translate the "something right" into a flat declaration of truth. In its place, "approximate truth" is used.

The decisive difficulty with the notion is its opacity. There is an extensive literature, founded mostly by Karl Popper, that seeks to provide a serviceable definition using the notions of "verisimilitude" or "truthlikeness." The literature provides conditions that allow us to say when one false theory is closer to the truth than another. A simple formulation compares the sets of true consequences derived from each. The one with the larger set is closer to the truth. This imperfect proposal comprises early steps in an expanding repertoire of proposals. It became enmeshed in continuing technical difficulties that seem, at least to me, remote from practical applications.[34]

The overall project of defining truthlikeness does not serve to rescue scientific realism from the fallibility of inductive inference. The informal idea is that a successful theory has some parts that are correct and some not. A viable definition of truthlikeness would tell us *what it means* for such a theory to have parts that resemble the truth. The difficulty is that such definitions do not tell us in any particular case, *which parts* of the theory resemble the truth. Cases to which the definition would be applied arise only when all applicable evidence has been applied inductively. To go beyond and, in addition, to designate which parts of the theory are truthlike and which not is to do more than can be done with inductive inferences from the evidence. But what inductive inference does exhausts what can be done.

There is a further problem of opacity in the way the notion of approximate truth is used in definitions of scientific realism. The idea of approximation requires some measure of closeness. In many cases, such measures are readily at hand. For example:

> The Washington Monument, a celebratory granite column in the Mall in Washington DC, is approximately 560 feet tall.

The notion of approximation used here is imprecisely given but easily made more precise. The "approximately" is satisfied if it is taken to mean "within 10 feet of the true height," which happens to be 554 feet 7 11⁄32 inches. Since examples like this can be given precise meaning, it is tempting to migrate the modifier "approximately" from the height of the column to the proposition overall:

> It is approximately true that the Washington Monument, a celebratory granite column in the Mall in Washington DC, is 560 feet tall.

As long as this reformulation is understood to assert exactly what the first version asserted, there is no problem. It is just a clumsy way of asserting the approximation in the height. If, however, we do not know how to reverse the migration of "approximately," the proposition has become opaque. Might we mean that the monument is only approximately in the Mall, but is perhaps

[34] For a survey, see Oddie and Cevolani (2022).

somewhere in larger Washington D.C.? Or might we mean that the monument is only approximately a column, since it is actually an obelisk, which has more precise properties. Or is that the monument is only partly granite? All we know that something might be wrong in the proposition, but we cannot say precisely what.

This opacity afflicts the use of approximate truth in scientific realism. An excessively cautious nineteenth century Newtonian might have wanted to characterize Newtonian theory as approximately true. But that Newtonian would scarcely have been able to reposition the modifier “approximately” correctly. The Newtonian did not then know which newer theories were to replace Newtonian theory. Recall the example of stellar aberration in an earlier chapter. A Newtonian would not know that the Newtonian rule of composition of velocities fails for scalar speeds close to that of light, but gives approximately correct results for angular shifts in the case of the aberration of starlight. These issues do not even touch the more severe problem of making sense of how a Newtonian particle is “approximately” a quantum particle.

4.2 Monster-Barring

Small-r realism assures us that our epistemic resources are powerful enough to secure contingent truths of a science with very little chance of error. We have good reason to expect that most of our mature, well-supported science will persist. However, small-r realism enjoins us not to take the next step, tempting as it might be. We should not try to argue away the small but enduring chance of error that remains even with the strongest of inductive support. Here we should proceed as in ordinary life. We do not live our lives in fear of dying by a lightning or meteor strike, but we cannot assert its absolute impossibility.

Further, we should not seek to provide a general formula for those sciences or parts of a science that have achieved permanence and are absolutely immune from correction. As long as the support of experience is inductive, there will always be some ineliminable chance of error, even if very small. No formula, no matter how artful, can eliminate it entirely.

A few accounts of scientific realism appear, at least on superficial reading, to succumb to the temptation. However, my sense is that, on closer reading, they prudently retain the qualification that the presumed stability of the science is still subject to a failure, even if it is very unlikely. They identify sciences in which the chance of error has become very small. The possibility of error remains, but it is discounted by a strategy that can be called “monster-barring."[35] They are such that the chance of their realization is so small that, *for practical purposes,* they can be ignored.

Hoefer's (2020) handles the problem deftly. His "tautological scientific realism" ("TSR," p.22) explicitly excludes present quantum theories from realism, as his title "Scientific Realism without the Quantum" suggests. He is surely correct in this exclusion. Although the successes of present quantum theory are widespread and extraordinary, the theory is beset with enduring foundational problems.[36] He describes the distinction crucial to TSR as: (2020, p. 22)

> The claim of TSR is not that our current science is so very successful that we should *infer* that it is probably true. That would be IBE [inference to the best explanation], and would be wrong. The claim of TSR is rather that certain parts of our current scientific lore are such that we can't conceive of any way that they could be seriously mistaken, other than by appealing to scenarios involving *radical* skepticism.

Hoefer offers present chemistry as a science in which we cannot now conceive of serious mistakes without this radical skepticism. He identifies three versions of this radical skepticism. Past regularities may simply fail to persist; or some "super-genius" may find a brilliant alternative chemistry; or we may find ourselves in a fanciful scenario such as the classic brains-in-vats fantasy.

My addition is to emphasize that these sorts of failures for a well-established theory are rare and it is simply a misplaced caution to take them seriously. However, we have no way to rule them out absolutely. There can be no viable, general principle that precludes the actuality of these extreme cases. The classical physicists of the nineteenth century did find their certainties overturned by the radical pathologies of quantum theory. It posits a radical non-locality of states in space and some interpretations of quantum theory ask to us to accept the reality of an ever-growing multiplicity of parallel worlds in which all possibilities are realized.

---

[35] The term "monster-barring" is borrowed from Lakatos' (1976) *Proofs and Refutations.*

[36] The measurement problem has resisted all attempts to find a single, universally accepted resolution; and the need to renormalize in quantum field theory results from the failure of the theory at high energies.

Vickers (2023) shares Hoefer's and my optimism about the success of science. What he seeks is clearly stated in his title, *Identifying Future-Proof Science*. His goal is (p. 1) "identifying scientific claims we can be confident will last forever." Vickers is surely correct to dispute the inductive skepticism of antirealists. He quite properly calls into question Wray's blanket inductive skepticism on the future stability of scientific theories. Vickers (2023, p. 4) displays the sort of skepticism he wishes to dispute in his quoting of Wray's (2018, p.1) claim:

> I will argue that our current best [scientific] theories are quite likely going to be replaced in the future by theories that make significantly different ontological assumptions.

Vickers eventually introduces criteria for identifying future-proof science. He allows—as he must—that even future-proof science may err. He is eager to minimize the chance of this error. He sets it as very low through a comparison: (p. 217)

> The claim … is that we are as certain about the status of those scientific claims as we are about historical claims such as the claim that the Second World War really happened, the Hiroshima bomb really was dropped on 6 August 1945, and so on.

My sense, however, is that Vickers is a too optimistic in two ways. First, although he is right to expect most of his future-proof science to remain, he underestimates the chance of error. It is equated with the chance of an error in ordinary experience, such as in his example that he is (at the moment of writing) striking the keys of his laptop. If we are to take this possibility of error seriously, Vickers suggests a dire consequence follows: (p. 4)

> … if it is insisted that we aren't sure about scientific ideas on these grounds, then we have to accept that we are never sure about anything. In which case words such as 'fact', 'sure', and 'certain' are never applicable, and might as well be eradicated from the dictionary: 'Knowledge is impossible!'

What Vickers regards as an unacceptable outcome is, for me, just the right way to understand the fallibility of inductive inference, if the terms Vickers lists are to connote absolute certainty, as I have made clear in the last chapter.

The second problem is that Vickers provides sociological criteria for identifying future-proof science. It is that at least 95% of a large scientific community that entertains a diversity of perspectives accepts the claim as (using Vickers' quote marks) "established scientific fact" (p. 217). From the perspective of small-e empiricism, the only way that such claims can be established is by the inductive support of experience. They are not established by what any community, no matter how prestigious, is willing to assert. The best that can be said is that the endorsement of the scientific community is a strong indicator that the requisite inductive support is in place. However, that endorsement is only a strong indicator but is not constitutive of that support.

5. Selective Realisms

Invocations of approximate truth indicate that something, but not everything, is correct in a theory. Selective realisms are a family of more narrowly targeted attempts to identify which parts of the theory are true. This section recounts three leading versions. If all they seek to do is to extract as much as can be learned from the evidence of experience inductively, they are just a form of scientific realism with no special powers that conforms with small-r realism. Selective realisms fail when they try to outmaneuver the powers inductive inference and the restrictions of inductive fallibility. This can happen in two ways.

In the first, they seek more for the elements of a theory that they have selected as the real than inductive inference can supply. That is, they seek a level of certainty for them that goes beyond the limits of inductive fallibility. These efforts must fail. For, according to small-e empiricism, inductive inferences from experience are our only means of learning the truth of contingent propositions.

In the second, they seek to discount the inductive bearing of the evidence of experience. This discounting applies to the elements of the theory that are not distinguished as selected for the formulation of selective realism. They are not accorded reality even if there is evidence supporting their reality, which is usually the case. These efforts must also fail. For, once again, inductive inferences from experience are our sole means of ascertaining the truth of contingent propositions. We have no further means of deciding when their assessments are cogent and when they are to be ignored.

There is another problem, internal to the selective realism literature. Different versions of selective realism contradict each other on which elements of a theory are to be taken as real and which are not. Entity realism accords a fixed reality to some entities, but not the relations that govern them. Structural realism takes the reversed view. It is only these relations that persist, not the entities supposed to bear them. That there can be such marked differences in the selection of selective realism indicated the fragility of the selection.

5.1 Structural Realism

One of the most popular amendments to scientific realism is structural realism. The overall idea is simple and carries an initial appeal. It came to prominence in the work of John Worrall (1989), where the approach is introduced as a response to the pessimistic meta-induction. Drawing inspiration from earlier analyses by Henri Poincaré, Worrall accepted that only some parts of our current best scientific theories can be correct. The proposal is that we can distinguish those parts as the "structure." It is the factual or real component of the theory that survives changes in theory. Worrall's enduring example is of Augustin Fresnel's early nineteenth century recovery of various results in optics. Fresnel derived his results on the assumption that light propagates as waves in an elastic medium, the ether. As an illustration, I and R are the intensities of the components of the incident and reflected beams of light, polarized in the plane of the reflecting surface. i and r are the angles of incidence and reflection. These quantities are related by the equation:

$$R/I = \tan(i-r) / \tan(i+r)$$

The key point for Worrall's analysis is that this equation and others like it were preserved when Fresnel's theory was replaced by Maxwell's electrodynamics of light. He noted: (p. 117, his emphasis)

> Roughly speaking, it seems right to say that Fresnel completely misidentified the *nature* of light, but nonetheless it is no miracle that his theory enjoyed the empirical predictive success that it did; it is no miracle because Fresnel's theory, as science later saw it, attributed to light the right *structure*.

and then (p. 118, his emphasis)

> [Maxwell's] field in no clear sense approximates the ether, but disturbances in it do obey *formally* similar laws to those obeyed by elastic disturbances in a mechanical medium.

Worrall called this view "structural or syntactic realism" (p. 112). He recognized that this illustration from Fresnel-Maxwell was exceptional in having Fresnel's equation carried over unchanged to Maxwell's theory. He noted: (p. 120, his emphasis)

> The much more common pattern is that the old equations reappear as *limiting cases* of the new - that is, the old and new equations are strictly inconsistent, but the new tend to the old as some quantity tends to some limit.

An illustration is provided in Ladyman and Ross' (2007, pp. 94-95) later development of structural realism. The Galilean transformation between inertial frames of reference in relative motion at speed v is defined as[37]

$$x' = x-vt \quad y' = y \quad z' = z \quad t' = t$$

It is replaced in special relativity by the Lorentz transformation[38]

$$x' = \gamma(v)(x-vt) \quad y' = y \quad z' = z \quad t' = \gamma(v)(t-vx/c^2)$$

Whereas we may justly worry that these transformations differ greatly in structure, Ladyman and Ross reassure us that: (p. 95)

> However, when ... v/c tends to 0, the mathematical structure of the latter [Lorentz transformation] clearly increasingly approximates that of the Galilean transformations.

This original conception of structural realism was soon refined and expanded into a range of views. A major division was emphasized by Ladyman (1998) and Ladyman and Ross (2007). They identified "epistemic structural realism" in which the structure of a theory is all that we can know of the theory. It was contrasted with "ontic structural realism" characterized by Ladyman and Ross (2007, p. 130) as:

> Ontic Structural Realism (OSR) is the view that the world has an objective modal structure that is ontologically fundamental, in the sense of not supervening on the intrinsic properties of a set of individuals. According to OSR, even the identity and individuality of objects depends on the relational structure of the world. Hence, a first approximation to our metaphysics is: 'There are no things. Structure is all there is.'

Worrall's original proposal was already sufficient to provoke considerable response in the literature. The ontologically adventurous proposal of ontic structural realism accelerated the response. For surveys, see Frigg and Votsis (2011) and Ladyman (2023). Among the many issues debated, one drew special attention. If in the ontic version "structure is all there is," what are we to make of the non-structural elements of a theory? Ladyman (2023, Section 4) identifies

---

[37] x, y, z are Cartesian spatial coordinates and t is a time coordinate. Primes indicate coordinates in a frame of reference moving at speed v along the unprimed frame's x axis in the $+x$ direction.

[38] $\gamma(v) = 1/(1-v^2/c^2)^{1/2}$

seven different answers. They span from "(1) … there are no individuals…" to "(7) Individual objects are constructs."

The issue I wish to raise here is the great difficulty in practice in identifying structure, in both epistemic and ontic versions. The case of the Galilean-Lorentz transformations is one of the simplest, but it already delicate. There is no simple preservation of formulae in the transition from the Galilean to the Lorentz transformation. We need to consider the limit of very small v/c to find it. In so doing, we eradicate the relativity of simultaneity, which is the distinctive novelty of special relativity and at the heart of the theory's novel effects. Recovering a continuity of structure between classical, Newtonian physics and quantum theory is even more delicate since the mathematical structure of the quantum theory differs greatly from that of classical physics.

The result is that, at least in my reading of the structural realism literature, is an enduring and bewildering vagueness in exactly what counts as structure. Where Worrall identified structure syntactically with formulae, Frigg and Votsis (2011, p.231) report the identification of structure using set theory, category theory, graph theory, information theory and group theory. Ladyman (2023, Section 1) notes "Clearly, there [are] many different notions of structure, and many different kinds of structural realism."

My view is that we can identify which structure is preserved through theory change and, if the selection is well-supported inductively, we can have some confidence that this structure will persist into future science. The fallibility of inductive inference, however, precludes an unqualified assertion that we have found with finality the structure that will persist into future science. This judgment is instantiated in a study of the evolution of theories of the electron that I co-authored with Jon Bain (Bain and Norton, 2001). This study is on a short list of illustrations of structural realism offered in Ladyman and Ross (2007, p. 94).

We sought to identify the characterization of the electron that persisted through an extraordinary era of growth in physics. It started with Thomson's (1897) identification of the particle in cathode rays, through the transformations of relativity and quantum theory and then on to the quantum field theory of the standard model of particle physics. We found that the Hamiltonian or Lagrangian for the electron remained remarkably stable throughout this era of tumultuous change. This, we argued, contradicted the claim of the pessimistic meta-induction that the history of science is a catalog of failure. Rather the persistence of this one structure gives an alternative picture of the sequence of theories accumulating and sequentially improving our understanding of the electron.

5.2 Structural Realism Contravenes Inductive Inference

Our ambitions in Bain and Norton (2001) were modest. We wrote: (p. 456)

> … the history of the electron shows us that our theories are always corrigible. Although we cannot display the structure [of some imagined final theory], we can certainly display our best candidate for that structure, recognizing that its form and content are likely to change as understanding grows. At any one time in the development of theories of the electron, we can read our best candidate from the latest theory. It is simply the smallest part of the latest theory that is able to explain the successes of earlier theories.

There are two elements in this modesty, one explicit and one implicit:

- Our *assured* identification of the structure of the electron was purely retrospective. It was the best account we could provide of the stable elements in the past evolution of theories of the electron. Inductive fallibility precludes projecting that assurance into future science.
- Our analysis did not deny that, at each stage, we had good evidence for those parts of the theory of the electron that extended beyond its structure.

A version of structural realism that conforms with these two elements, limited to this particular case, would, in my view be both serviceable and defensible. However, it would offer nothing that is not already in small-r realism.[39] In this case, it simply reports the evidence that electrons *have been* characterized stably by their Hamiltonians and that it gives inductively good, but fallible reasons for expecting the characterization to persist. The import of evidence for other parts of the various, evolving theories of the electron, is not discounted.

In so far as versions of structural realism in the literature breach either of these conditions, they are in conflict with what inductive inference does and can provide.

First, where Bain and I were careful to limit the future prospects of our results, I find no corresponding caution in the structural realism literature. There, the structure identified is given a

[39] Here I join a tradition of criticism reported in Ladyman (2023, Section 5) that a defensible form of structural realism collapses into generic scientific realism.

timeless character without an obligatory qualification of its corrigibility. In so far as this amounts to a declaration that the final truth has been discovered, it contradicts the fallibility of inductive inference that precludes any such final declaration.

That a timeless character is sought for structure seems to be the original motivation for structural realism and at least a part of its continuing motivation. French and Ladyman (2003, p. 33) report:

> For Worral the motivation for adopting structural realism is just the need to respond to the pessimistic meta-induction. While we share the desire to articulate a form of realism that can account for theory-change, our understanding of structural realism is also informed by the need to provide an ontology that can dissolve some of the metaphysical conundrums of modern physics, and also the need to have a conception of how theories represent the world that is compatible with the role of models and idealisations in physics.

To be clear, since the pessimistic meta-induction seeks to establish future failures of science, a response must concern itself with the prospects of future science.

Of course, I may be reading too much into the lack of a disclaimer that the timeless import of structural realism is intended only to be secure retrospectively but prospectively corrigible. However, if such a disclaimer is intended tacitly, then what results is a doctrine that fails to extend beyond small-r realism. For there is no *assured* way to identify *prospectively* the structure that will be preserved through all future theory change. Thomson in 1897 might well have imagined that future theorizing would alter his account of the electron. However, he would have had no basis for identifying the Hamiltonian specifically as that component that would endure. Small-r realism allows us to assert that, on our best present evidence, electrons are characterized by their Hamiltonians and, precisely because of this strong inductive support, we would expect this characterization to persist, but we can give no absolute assurance of it. More generally, small-r realism allows the expectation that something of the present theory will endure, if it does happen that future theories contradict present theory. But it can give no absolute assurance of what the preserved element will be.

The result is a diluted structural realism. It merely gives the anodyne assurance that something will remain unchanged and something will not, if a new theory happens to replace the present theory. What it cannot do is to provide any final characterization of what will assuredly remain unchanged. That is, it cannot provide an independent theory of structure with substantive content. For that would be to anticipate the specific character of all future theories, in contradiction with the fallibility of inductive inference.

The second breach with small-r realism is in the *denial* of the import of the inductive support for the non-structural parts of a theory. Epistemic structural realism denies that these non-structural parts can be known; and ontic structural realism denies that they can be real. What may tempt structural realists to this extraordinary denial is that, when present sciences replace past sciences, the denial minimizes the corrections to the inventory of what is real. Thomson may have thought that he had identified electrons as charged, classical particles. However, structural realists tell us that he was not authorized to this belief. He was only authorized to accept the structure, even though he could not then have known what it was.

This denial by structural realists is irrational in the direct sense that it contravenes the dictates of rationality as encoded in an inductive logic. As a general matter, in any well-developed science, the scientists have good evidence from experience for all parts of their theory. Thomson's (1897) "Cathode Rays" is celebrated as the definitive identification of electrons as charged particles, as opposed to Lenard and Hertz's identification of them as ethereal waves. Norton (2021, Ch.9, Section 6) reviews in some detail how Thomson's arguments in that and associated papers established his identification to the satisfaction of the then present physics community. The bearing of the evidence on the non-structural parts of a theory in such cases cannot be discounted because it is inconvenient for structural realism.

5.3 Clarifying Structural Realism[40]

These last remarks constitute my best efforts to circumscribe and assess the state of structural realism. What has made these efforts difficult is a lack of coherence in the structural realism literature itself, especially after the introduction of ontic structural realism. One manifestation is the enduring proliferation of different accounts of the nature of structure. Another is an evident tension between claims made for structural realism. For example, French and Ladyman (2011, p. 32) write:

> The job of predicting what will be preserved and what abandoned by future science belongs to science itself not to philosophy, …

[40] I thank Aiden Meyer for helpful discussion that led to the writing of this subsection.

This is a most prudent assertion, quite compatible with small-r realism. However, the sentence continues with a "but" that undermines the prudence:

> … but our claim is that from a structuralist point of view it is possible to explicate the continuity in scientific theories that often does not hold at the level of objects and properties.

What they had in mind had already been laid out in text earlier in the paragraph. They first noted that some structure in Newtonian physics had not been preserved in the transition to special relativity. They continued: (pp. 31-32, my emphasis)

> The advocate of OSR is not claiming that the structure of our current theories will be preserved simpliciter, but rather that the well-confirmed relations between the phenomena will be preserved in at least *approximate* form and that the modal structure of the theories that underlies them, and plays the appropriate *explanatory* role, will also be preserved in *approximate* form.

Ontic structural realism, they clearly assert, predicts the preservation of something in future science. What precisely that will be is now hedged into vagueness by the terms I have highlighted. Section 4.1 above recalled the troublesome ambiguity associated with the modifier "approximate"; and the next chapter will recount the corresponding problem for "explanatory."

Ladyman, apparently, was quite comfortable with this formulation. In his later (2021, p. 246), he quotes it as an expression of his and French's joint view. The quoting of this passage is then immediately followed by an assertion of the claim foreshadowed by the title of the 2021 paper, "Structure not Selection," that ontic structural realism is not a selective realism.

> However, this does not enable any kind of selective realism because it says nothing at all about how the laws of the theory are only approximate, and so does not enable the selection of what will be retained in advance.

The passage Ladyman quotes gives a prescription, albeit ambiguous, for which parts of a science will be preserved in future science; and that is a just what a selective realism can do.

5.4 Entity Realism

Entity realism asserts that we should be cautious asserting the truth of a theory. However, it asserts, the prospects are much better when it comes to asserting the reality of entities within a theory. We find the same misalignment with inductive inference as with structural realism. On the one hand it seeks to go beyond what inductive support can provide in assuring us of the reality of certain, selected entities. On the other hand, it contravenes the import of inductive support in denying reality to anything else.

Since Hacking (1982, 1983, 1989) is the most prominent proponent of entity realism, we can explore his specific arguments in greater detail. Hacking employs Putnam's referential model of meaning to provide theoretical entities like electrons with some independence from the theories in which they appear. Hacking (1982, pp. 73-75; 1983, Ch. 6) employs the resources of Putnam's theory to argue that the many theories of the electron, posited over decades, were all referring to the same entity, while attributing different properties to it.

More is needed if we are to affirm that the entity is real. Many authors have written about the same character, Sherlock Holmes, and have given him different roles in their fictional writing. He is still not real. To establish the reality of his core example, the theoretical entity electrons, Hacking reports the experiences of experimenters who use electrons to create further effects. We, in his remarks, are to identify with the experimenters. He gives the same characterization in his two accounts. Hacking sets the entire passage in italics to emphasize its importance: (1982, p. 77; 1983, p. 265)

> *We are completely convinced of the reality of electrons when we regularly set out to build and often enough succeed in building—new kinds of device that use various well-understood causal properties of electrons to interfere in other more hypothetical parts of nature*.

This is the repeated theme of Hacking's narrative. Experimenters are convinced of the reality of electrons and the experience of using electrons in experiments is convincing:

> The experimenter is convinced of the reality of entities some of whose causal properties are sufficiently well understood that they can be used to interfere *elsewhere* in nature. (1982, p. 75, his emphasis)
>
> The experimentalist does not believe in electrons because, in the words retrieved from medieval science by Duhem, they "save the phenomena." On the contrary, we believe in them because we use them to *create* new phenomena, such as the phenomenon of parity violation in weak neutral current interactions. (1982, p. 84, his emphasis)

> Experimenting on an entity does not commit you to believing that it exists. Only manipulating an entity, in order to experiment on something else, need do that. (1983, p. 263)

What is affirmed here, again and again, is a powerful psychological effect. Experimenters cannot be anything other than realists about electrons in order to design and carry out their experiments.[41]

What is missing from Hacking's analysis is a sustained analysis of whether this experimentalists' conviction is well-founded evidentially. Mere reports of the firmness of their conviction is an indicator that there might be such a case to be made. But the reports do not make the case.[42]

Hacking's narrative aims to privilege the reality of entities used in interventions over other realities. He wrote (1982, p. 71) "The experimentalist need only be a realist about the entities used as tools." The difficulty is that the experiences Hacking reports occur in many other contexts. When one learns special relativity, it is astonishing how often its two simple postulates generate effect after effect in physics. Particle physicists have found again and again that the standard model of particle physics accommodates a steady stream of new accelerator results. General relativists find again and again that when any competitor to general relativity can be put to the test, general relativity wins. The same is true for engineers who successfully apply Newton's mechanics to system after system. In all these cases, the theorists develop a realist attachment to the core posits of each of these theory.

We do not need to pause here to assess whether these convictions are sustained by a viable inductive argument from the evidence. For present purposes we need only this. IF there is a viable argument form here, then there is nothing special about the entities used in interventions.

[41] Hacking's repeated reports do not fit with his remark: (1982, p. 71) "But experimenters are realists about the entities that they use in order to investigate other hypotheses or hypothetical entities. That is not a sociological fact."

[42] In efforts to reconstruct Hacking's argument, Miller (2016) found four possible arguments and one non-argument.

Hacking has identified a quite general inductive argument form that extends well beyond entity realism. It has returned us to a generic scientific realism that authorizes a wide range of realities.

Hacking does insist that what makes the experimental intervention case special is that it depends on the appreciation of the causal powers of the entity and their use in interventions. He wrote: (1982, p. 86)

> We can, however, call something real, in the sense in which it matters to scientific realism, only when we understand quite well what its causal properties are.

And again

> Experimenting on an entity does not commit you to believing that it exists. Only *manipulating* an entity, in order to experiment on something else, need do that. (1983, p. 263, his emphasis)

I did not see how Hacking made the case for this beyond merely repeatedly asserting it. The experiences of special relativists, general relativists, particle physicists and classical physicists are just as compelling.

5.5 Entity Realism versus Structural Realism

It is a most curious aspect of the selective realism literature that two of its most prominent versions stand in direct contradiction with each other. Structural realism elevates not just the general theory, but a particular element in it, the structure. Epistemic structural realism precludes knowledge of the unobservable entities that bear the theoretical structure. Frigg and Votsis' (2011, p. 232, their emphasis) give this definition: "*Epistemic structural realism* (ESR) is the view that all we can know about the unobservable world is its structure." Ontic structural realism is more severe. Ladyman and Ross offer the memorable slogan above "There are no things. Structure is all there is." For structural realists, the entities are either opaque to us or fictional.

Hacking, in his entity realism inverts this position. The equations are variable, but the entities are fixed. They are the true realities. Hacking noted: (1982, p. 72)

> The experimenter is convinced of the existence of plenty of "inferred" and "unobservable" entities. But no one in the lab believes in the literal truth of present theories about those entities. Although various properties are confidently ascribed to electrons, most of these properties can be embedded in plenty of different inconsistent theories about which the experimenter is agnostic.

He adds in the later publication: (1983, p. 264)

> There are a lot of theories, models, approximations, pictures, formalisms, methods and so forth involving electrons, but there is no reason to suppose that the intersection of these is a theory at all. Nor is there any reason to think that there is such a thing as 'the most powerful non-trivial theory contained in the intersection of all the theories in which this or that member of a team has been trained to believe'. Even if there are a lot of shared beliefs, there is no reason to suppose they form anything worth calling a theory

5.6 Explanationist Realism (Indispensabilism)

A third variety of selective realism has come to be called "explanationism." Here is how it is introduced in Chakravartty's survey: (2017, Section 1.3, his emphasis)

> These species of generic realism can be viewed as falling into three families or camps: explanationist realism; entity realism; and structural realism. There is a shared principle of speciation here, in that all three approaches are attempts to identify more specifically the component parts of scientific theories that are most worthy of epistemic commitment. *Explanationism* recommends realist commitment with respect to those parts of our best theories—regarding (unobservable) entities, laws, etc. that are in some sense indispensable or otherwise important to explaining their empirical success—for instance, components of theories that are crucial in order to derive successful, novel predictions.

The label "explanationism" is misleading. The key thesis is of an *indispensability* of the components of a theory that are responsible for its success. These components, loosely speaking, explain why some theory has been successful in spite of it harboring other components that are false. Proponents of this indispensabilism are not advocating an explanation-based abductive account of inductive inference.

The core of the view is that we can identify which elements of an otherwise false theory are responsible for its successes; and that we are entitled to attribute reality to them. In this form, the view is a more relaxed version of structural realism. Structural realism identifies these real elements as the structure of a theory and seeks to give structure a more specific characterization. Indispensabilism merely asserts that there are such elements and they are to be identified by their indispensable role in generating empirical successes. In this formulation, the thesis is one of inductive inference: we are authorized to infer to these elements because of their indispensability.

Because of its close affinity with structural realism, its assessment here is similar. Indispensabilism provides no viable addition to small-r realism. This is the benign reading of indispendabilism. In its narrowest form it may merely identify its indispensable elements *retrospectively* as truths about past and present theories. It may also enjoin us to infer to the likely persistence *prospectively* of the indispensable elements in future science. As long as this inference proceeds within the limits of inductive fallibility, it conforms with small-r realism. What it adds is a detail in the specific form the inductive inference is to take as one founded in indispensability.

There is a less benign reading that may be suggested but is left somewhat open. It is the proposal of an enduring truth that the presently indispensable elements will assuredly persist in future theories. If that stronger claim is intended, it contradicts the fallibility of inductive inference and is unsustainable.

A version of this indispenabilism is found distributed through Kitcher (1993, Ch.5). He remarks, for example, "No sensible realist should ever want to assert that the *idle* parts of an individual practice, past or present, are justified by the success of the whole." (p. 142, his emphasis). A more fully elaborated and best-known version of indispensabilism was provided in papers by Psillos (1994, 1996), whose essential content was reproduced and developed in his monograph (1999). He describes the view as: (1999, p. 103)[43]

> Put positively, it is enough to show that the theoretical laws and mechanisms which generated the successes of past theories have been retained in our current scientific image. I shall call this the *divide et impera* move. It is based on the claim that when a theory is abandoned, its theoretical constituents, i.e. the theoretical mechanisms and laws it posited, should not be rejected *en bloc*. Some of those theoretical constituents are inconsistent with what we now accept, and therefore they have to

[43] In the strategy of *divide et impera* ("divide and conquer"), famously used by Julius Caesar in his subjugation of Gaul, victory is secured by fomenting and exploiting division in one's enemies. The connection to indispensabilism seems slender as does the desirability of employing a label that routinely applies to the invasion and subjugation of one's neighbors.

> be rejected. But not all are. Some of them have been retained as essential constituents of subsequent theories. The *divide et impera* move suggests that if it turns out that the theoretical constituents that were responsible for the empirical success of otherwise abandoned theories are those that have been retained in our current scientific image, then a substantive version of scientific realism can still be defended.

The upshot of this move is a response to the pessimistic meta-induction: (p. 104, his emphasis)

> The right assertion seems to be that the genuine empirical success of a theory does make it reasonable to believe that the theory has *truth-like constituent theoretical claims*.

Where Kitcher's analysis is concerned with the referential success of terms, Psillos' emphasized that his account concerns empirical support for certain parts of a theory: (2022, Section 2.5)

> The underlying thought is that the empirical successes of a theory do not indiscriminably support all theoretical claims of the theory, but rather the empirical support is differentially distributed among the various claims of the theory according to the contribution they make to the generation of the successes.

The account is, at its core, a contribution to inductive logic.

5.7 Indispensabilist Inductive Logic

Psillos' version of indispensabilism posits that the empirical evidence can selectively favor different parts of a theory. That such favoring can be identified on a case by cases basis is plausible. It is, in general, difficult and delicate. Judgments of which parts of a theory deserve credit for its success tend to be fragile and are commonly revised when the theory is superseded. In the case of stellar aberration from the last chapter, Bradley could very plausibly have attributed the success of his aberration formula to the rule of Newtonian addition of velocities. After Einstein's development of special relativity, the success is judged as due to something narrower: that Newton's rule happens to coincide with the relativistic rule only in the direction of composed velocities for the particular case of aberration.

With this delicacy recognized, it would be an important advance for realism if an indispensabilist analysis could give a general rule for identifying such favored elements. Psillos does offer such a rule: (1999, p. 105)

> When does a theoretical constituent H indispensably contribute to the generation of, say, a successful prediction? Suppose that H together with another set of hypotheses H' (and some auxiliaries A) entail a prediction P. H indispensably contributes to the generation of P if H' and A alone cannot yield P and no other available hypothesis H^* which is consistent with H' and A can replace H without loss in the relevant derivation of P.

This rule has some initial plausibility. It is, however, precisely the sort of general inductive rule that is precluded by the material theory of induction. Within a much more extensive critical appraisal of Psillos' proposal, Lyons (2006, p. 540) points out the immediate weakness of the rule. It generally proves easy, in every case, to find many alternative hypotheses that can serve as the H^* of the rule. It follows that no H is indispensable. Psillos had already recognized the threat in his original narrative and had proposed a solution: (1999, p. 105)

> … if we impose some natural epistemic constraints on the potential replacement—if, for instance, we require that the replacement be independently motivated, non ad hoc, potentially explanatory, etc.—then it is not certain at all that a suitable replacement can always be found.

This is the first step in a mode of failure for proposals for universally applicable rules of inductive inference that, through my work on the material theory of induction, I have found to be quite common. Counterexamples are found to an initially plausible rule. Extra clauses are added to protect the rule from the counterexamples. Psillos' extra clauses are only vaguely indicated. Were they specified more carefully, we would find further counterexamples for the modified rule; and the process would continue indefinitely.

Psillos' rule is an example of what I have called "hypothetical induction" in my survey (Norton, 2005). Simple versions of hypothetical induction suffer the same problem as Psillos'. There are too many hypotheses that entail the evidence. An entire industry of additions to the original rule grew up in unsuccessful attempts to preserve the universal applicability of a modified version of the rule.

5.8 Indispensabilism and Sadi Carnot's *Réflexions*

That identifying the indispensable part of a theory is far from easy is illustrated by the first of Psillos' two major illustrations. Sadi Carnot's (1824) *Réflexions sur la Puissance Motrice du Feu* is one of the most extraordinary works of science of all eras. In a few short pages it lays the foundations for a new science, thermodynamics, along with novel methods of analysis of great subtlety. It succeeds in doing so while working with an incorrect theory of heat, caloric, as a conserved substance and without positing an independent second law of thermodynamics.[44]

Psillos (1994; 1999, Ch.6) seeks to show that Carnot's success did not depend on any indispensable elements that were not passed over to the subsequent formulations of thermodynamics in which heat is merely a form of energy that can be converted into work. Psillos' investigation of the history of Carnot's work is thorough and commendable. However, his logical analysis of which parts of Carnot's work were indispensable for the success of Carnot's analysis and were carried forward to later theories is mistaken. Here is how Psillos summarizes these efforts: (1996, p. 310)

> It turns out, for instance, that Carnot's explanation of the fact that maximum work is produced in a Carnot-cycle employed only the principle of the impossibility of perpetual motion and not the assumption that heat is a material substance

Psillos emphasizes that the assumption of the conservation of heat is dispensable: (1999, p. 119)

> Hence, Carnot did not appeal to any assumptions about the conservation of heat in order to establish his law.

This appraisal is incorrect on all points. The decisive element in Carnot's determination of the conditions for the recovery of the maximum work is his general depiction of a heat engine as operating by passing heat from a hot source to a cold reservoir. Without it, Carnot's analysis cannot be set up. Carnot used this conception since he assumed heat to be conserved. It followed that the heat in a heat engine had to go somewhere and it was only the motion of heat from hot to cold that facilitated the generation of motive power.

This foundational posit was derived from the assumption of the conservation of heat. If we conceive of heat as transformable to work, there is no need for this posit. The obvious default is that the maximum work energy that can be recovered is just the energy of the heat itself. That amount greatly exceeds the maxima recoverable with a Carnot-style analysis. In this sense, Carnot's (false) assumption of the conservation of heat was indispensable to his analysis.

[44] For my analysis of what is extraordinary in Carnot's work, see Norton (2022). The concept of a thermodynamically reversible process is a proposal of the most profound foundational importance to thermodynamics.

Psillos correctly notes that Carnot's argument employs the impossibility of a perpetual motion machine. However, he is mistaken in suggesting that it is indispensable for the later recovery of Carnot's results, such as in William Thomson's (1852) celebrated papers. In them, the assumption of the conservation of heat is discarded and it is allowed that heat can be converted to and from work energy. The impossibility of a perpetual motion machine does appear as what becomes the first law of thermodynamics. It appears as Thomson's (1852, p.11) Prop. I. The first law alone is an insufficient basis for deriving Carnot's results on the maximum of work recoverable from a heat engine. The second law of thermodynamics must be posited independently if we are to arrive at these results. Accordingly, Thomson (1852, pp. 12-14) needed to introduce extra posits, inspired by Carnot's analysis, as assumptions so that he could recover the limits Carnot deduced in his earlier work.

6. Conclusion

The previous chapter argued that we need to reduce the formulation of scientific realism to accommodate inductive fallibility; and that small-r realism provides a sustainable reduction that cannot be expanded. This chapter has sought to support this claim in light of the present proliferations of writing on scientific realism. It reviews the range of stronger formulations of scientific realisms in the present literature. It finds that formulation stronger than small-r realism fail when they go beyond small-r realism.

The enormity of the present literature in scientific realism and the great variety of proposals has made this a daunting task. My hope is that the survey has collected enough of the variety of formulations do that contact, even if weak, has been made with most serious proposals. I apologize in advance to those who find their favorite formulation slighted. My sense, however, is that the survey has captured the major strands of the present literature. Some views that carry the name "realism"—"perspectival realism" and "realistic realism"—will be discussed in the next chapter, since they are, in my view, versions of antirealism.

Finally, I have ample anecdotal evidence of a common reaction to small-r realism. Almost everyone who favors scientific realism has a version they find appealing. The reaction is initial disbelief that this favorite version overreaches and that we need to retract from it. This chapter may not convince everyone, but at least it conveys why I see there to be overreaches.

References

Azhar, Feraz and Butterfield, Jeremy (2018) “Scientific Realism and Primordial Cosmology,” Ch. 24 in Saatsi (2018)

Bain, Jonathan and Norton, John D. (2001) “What Should Philosophers of Science Learn from the History of the Electron?” Ch. 15 in Jed Z. Buchwald and Andrew Warwick, eds., *Histories of the Electron: The Birth of Microphysics*. Cambridge, MA: MIT Press.

Carnot, Sadi (1824) *Réflexions sur la Puissance Motrice du Feu*. Paris: Bachelier. Trans., *Reflections on the Motive Power of Heat*. New York: John Wiley & Sons, 1897.

Chakravartty, Anjan (2007) *A Metaphysics for Scientific Realism: Knowing the Unobservable.* Cambridge: Cambridge University Press.

Chakravartty, Anjan (2017) “Scientific Realism,” *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/sum2017/entries/scientific-realism/ .

Chakravartty, Anjan (2024) “On Semirealism, Realism More Generally, and Underlying Epistemic Stances,” *Journal for General Philosophy of Science*, **55**, pp. 269–288.

Dicken Paul (2016) *A Critical Introduction to Scientific Realism*. London: Bloomsbury.

Dieks, Dennis (2023) “Emergence, Continuity, and Scientific Realism,” *Global Philosophy,* (2023) 33:44.

Edwards, Paul, ed., (1967) *The Encyclopedia of Philosophy*. Volume 7. New York: The Macmillan Co. and the Free Press.

Fleming, William (1860) *The Vocabulary of Philosophy*. Philadelphia: Smith, English & Co.

French, Steven (2000) “Identity and Individuality in Quantum Theory” *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/spr2000/entries/qt-idind/

French, Steven (2018) “Realism and Metaphysics,” Ch. 31 in Saatsi (2018).

French, Steven and Ladyman, James (2003) “Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure,” *Synthese*, **136**, pp. 31-56.

French, Steven and Ladyman, James (2011) “In Defense of Ontic Structural Realism,” Ch. 2 in P. and A. Bokulich, eds., *Scientific Structuralism*. Dordrecht: Springer.

Frigg, Roman and Votsis, Ioannis (2011) “Everything you always wanted to know about structural realism but were afraid to ask,” *European Journal for Philosophy of Science* **1**, pp. 227–276.

van Fraassen, Bas C. (1980) *The Scientific Image*. Oxford: Clarendon Press.

Hacking, Ian (1982) "Experimentation and Scientific Realism," *Philosophical Topics*. **13**, pp. 71-87.

Hacking, Ian (1983) *Representing and Intervening.* Cambridge: Cambridge University Press.

Hacking, Ian (1989) "Extragalactic Reality: The Case of Gravitational Lensing," *Philosophy of Science*, **56**, pp. 555-581.

Hoefer, Carl (2020) "Scientific Realism without the Quantum" Ch. 2, pp. 19-34 in, Steven French and Juha Saatsi, eds., *Scientific Realism and the Quantum.* New York: Oxford University Press.

Kitcher, Philip (1993) *The Advancement of Science*. Oxford: Oxford University Press.

Ladyman, James (1998) "What is Structural Realism?" *Studies in History and Philosophy of Science*, **29**, pp. 409-424.

Ladyman, James (2021) "Structure not Selection," in T. D. Lyons and P. Vickers, eds., *Contemporary Scientific Realism*. Oxford University Press.

Ladyman, James (2023) "Structural Realism," *The Stanford Encyclopedia of Philosophy* (Summer 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), <https://plato.stanford.edu/archives/sum2023/entries/structural-realism/>.

Ladyman, James and Ross, Don (2007) *Everything Must Go: Metaphysics Naturalized*. Oxford: Oxford University Press.

Lakatos, Imre (1976) *Proofs and Refutations*. Cambridge: Cambridge University Press.

Leplin, Jarrett (1984) "Introduction" pp. 1-7 in, J. Leplin, ed., *Scientific Realism*. Berkeley: University of California Press.

Lyons, Timothy D. (2006) "Scientific Realism and the *Stratagema de Divide et Impera*," *The British Journal for the Philosophy of Science*, Sep., **57**, pp. 537-560.

Miller, Boaz (2016) "What is Hacking's argument for entity realism?" *Synthese*. **193**, pp. 991–1006.

Niiniluoto, Ilkka (1999) *Critical Scientific Realism*. Oxford: Oxford University Press.

Norton, John D. (2005) "A Little Survey of Induction," in P. Achinstein, ed., *Scientific Evidence: Philosophical Theories and Applications*. Johns Hopkins University Press, pp. 9-34.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Norton, John D. (2022) “How Analogy Helped Create the New Science of Thermodynamics,” *Synthese*. **200,** article 269, pp. 1-42.

Oddie, Graham and Cevolani, Gustavo (2022) “Truthlikeness,” *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), <https://plato.stanford.edu/archives/win2022/entries/truthlikeness/>.

Park, Seungai (2022) *Embracing Scientific Realism*. *Synthese Library 445: Studies in Epistemology, Logic, Methodology and Philosophy of Science.* Cham: Springer.

Psillos, Stathis (1994) “A Philosophical Study of the Transition from the Caloric Theory of Heat to Thermodynamics,” *Studies in History and Philosophy of Science*. **25**, pp. 159-90.

Psillos, Stathis (1996) “Scientific Realism and the ‘Pessimistic Induction’,” *Philosophy of Science*, **63**, pp. S306-S314.

Psillos, Stathis (1999) *Scientific Realism: How science tracks truth.* London and New York: Routledge.

Psillos, Stathis (2022) “Realism and Theory Change in Science,” *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), <https://plato.stanford.edu/archives/fall2022/entries/realism-theory-change/>.

Rescher, Nicholas (1987) *Scientific Realism*: *A Critical Reappraisal.* Dordrecht: Reidel.

Saatsi, Juha, ed. (2018) *The Routledge Handbook of Scientific Realism*. London and New York: Routledge.

Sankey, Howard (2008) *Scientific Realism and the Rationality of Science*. Aldershot/Burlington: Ashgate.

Sankey, Howard (2018) “Kuhn, Relativism and Realism,” Ch. 6 in Saatsi (2018).

Stanford, Kyle (2018) “Unconceived Alternatives and the Strategy of Historical Ostention,” Ch. 17 in Saatsi (2018).

Thomson, Joseph J. (1897) “Cathode Rays.” *Philosophical Magazine* **44**, pp. 293–316.

Thomson, William (1852) “On the Dynamical Theory of Heat, with Numerical Results Deduced From Mr. Joule’s Equivalent of a Thermal Unit, and M. Regnault’s Observations on Steam,” *Philosophical Magazine*, **4** (Ser. 4), pp. 8-21, 105-117, 168-176, 256-260.

Vickers, Peter (2023) *Identifying Future-Proof Science*. Oxford: Oxford University Press.

Worrall, John (1989) “Structural Realism: The Best of Both Worlds?” *Dialectica*, **43**, pp. 99-124.

Wray, K. Brad (2018) *Resisting Scientific Realism*. Cambridge: Cambridge University Press.

9. Inductive Skepticism in Scientific Antirealism

Chapter 9

Inductive Skepticism in Scientific Antirealism

1. Introduction

Antirealism in philosophy of science is not a single doctrine. It is a collective term for accounts that deny scientific realism. It follows that the term is inherently vague; and it is so for two reasons. First, since there is at present no single definition of scientific realism, the particular sense of antirealism invoked will depend on which version of scientific realism is disputed. Second, each instance of antirealism arises within a viewpoint that contradicts some version of scientific realism. Thus, each instance will in turn reflect the particular orientation of the positive viewpoint that hosts it.

This inherent vagueness makes general claims about antirealism difficult to sustain. The same is not true of the *arguments* used to support antirealism. A persistent feature of these arguments is a misappraisal of the cogency of inductive inference and inductive support. An inductive skepticism that just denies the inductive import of evidence is common. Some versions arrive at it by amplifying the fallibility of inductive inference into the assured failure of inductive inference. Examples below are the "all or nothing" argument and the pessimistic meta-induction.

Antirealists must deny this inductive import. For the inductive import of the massive compilation of evidence for our best sciences strongly supports their truth, in contradiction with antirealist claims.

This chapter explores the role and variety of anti-inductive views in scientific antirealism. Section 2 reviews the widely noted inductive skepticism of the undetermination thesis. It depends on a naïve form of hypothetico-deductive confirmation. Section 3 recounts the anti-explanationism van Fraassen developed to defend his version of antirealism. It employs a correct but misdirected concern that explanatory prowess need not indicate truth. Section 4 recounts the simpler "all or nothing" argument for inductive skepticism. It depends on a false dilemma: either inductive support is inerrant or a failure.

Sections 5 and 6 review the most enduring argument for antirealism, the pessimistic meta-induction. It tries to establish that present and future sciences must be forever poorly grounded since there have been notable past failures. On closer examination, the pessimistic meta-induction itself fails in every way that inductive inferences can be assessed. Most importantly, it pre-emptively discounts the strength of present evidence for all sciences without even considering that evidence.

Section 7 completes the analysis of antirealism by reviewing two views that carry the label "realism" but prove, on closer inspection, to be forms of antirealism: perspectival realism and Chang's "realistic realism." Perspectival realism, reviewed in Subsections 7.1 to 7.4, seeks to compartmentalize the import of evidence into separate perspectives that are divided according to the source of the evidence. This compartmentalization is inconsistent with how evidence from different sources is used routinely. Evidence from different sources can mutually support each other, so that each are seen to support the same facts. Realistic realism, discussed briefly in Section 7.5, seeks to replace a realist conception of theories with an antirealist, pragmatic conception. Section 8 is a short conclusion.

2. The Underdetermination Thesis

A persistent strand of anti-inductivism arises in the context of the widely reported thesis of the underdetermination of theory by evidence.[45] A weaker version of the thesis merely asserts that sometimes we may not have sufficient evidence available to establish a particular theory. This circumstance is common when new sciences or scientific theories are emerging. It is not a form of anti-inductivism.

The anti-inductive version of the thesis is its strong form. It asserts that no body of evidence, no matter how extensive, can determine a particular theory. By "determine," I understand that the body of evidence is to supply strong inductive support for one specific theory, such that it is preferred over all its rivals.

The thesis enjoys continuing respect outside philosophy of science. Within philosophy of science, even though some support persists, it has been widely derided. My critique in Norton (2008) finds the strong form of the thesis to be little more than a speculation that some find plausible and even compelling and others, like me, find unsustainable.

The thesis, I have argued, depends on a naïve form of hypothetico-deductive confirmation, in which two hypotheses are mistakenly judged equally well confirmed if they

[45] See Stanford (2023) for a recent survey.

deductively entail the same evidence. We saw in Chapter 4 above that Bridgman was unable to separate the standard theory of the history of the universe from a fantasy theory in which the universe was created five minutes ago with all memory traces of an earlier time as part of the creation. Bridgman was correct that each entail the same evidence. However, as I noted in Chapter 4, he was wrong to conclude that this evidence cannot separate the two, for the evidence favors the standard theory strongly.

The arguments for the strong form of the underdetermination thesis routinely display theories that are empirically or observationally indistinguishable. In addition to employing the naïve form of hypothetico-deductive confirmation, these examples are self-defeating. For their indistinguishability is established by displaying a mapping between their observational consequences. The fact that such a mapping can be displayed leaves open the possibility that the theories are merely variant formulations of the same facts.

Finally, there is a striking consequence for our mature sciences. The results of these sciences have acquired an exclusive stability and security that is sustained by a rich body of evidence. Its strength of support is judged by the scientists to be overwhelming. The strong form of the underdetermination thesis requires that these scientists, in science after science, are simply mistaken in their assessment of the evidence. The more plausible alternative is that a few skeptical philosophers have been hasty is adopting an oversimplified and faulty account of inductive evidence.

3. Anti-explanationism

Chapter 4 of *small-e empiricism* reported that a broadly focused anti-inductivism became explicit in van Fraassen's developing skepticism. One element of van Fraassen's growing anti-inductivism proved especially troublesome for scientific realists. It assailed directly what is still the most popular argument for scientific realism, the "no miracles argument." This argument depends on the notion of explanation in science. According to it, the immense success of science would be inexplicable—miraculous—if some version of scientific realism were not true. A celebrated version was given by Putnam, when contemplating the reality of electrons, DNA and the curved spacetime of general relativity. He wrote (1975-76, p. 178, his emphasis)[46]

[46] Putnam's version is distinctive in speaking of "miracles." The same argument, without "miracle" verbiage, is already in Smart (1963, p. 47).

> But if these objects don't really exist at all, then it is a *miracle* that a theory which speaks of gravitational action at a distance successfully predicts phenomena; it is a *miracle* that a theory which speaks of curved space-time successfully predicts phenomena; and the fact that the laws of the former theory are derivable "in the limit" from the laws of the latter theory has no explained methodological significance.

Van Fraassen (1980, Ch.2) denied that the success of science requires adopting scientific realism. It is not needed, he insisted, to explain the success of science. Famously, he offered an alternative, modeled on evolutionary biology: (p. 40, his emphasis)[47]

> … I claim that the success of current scientific theories is no miracle. It is not even surprising to the scientific (Darwinist) mind. For any scientific theory is born into a life of fierce competition, a jungle red in tooth and claw. Only the successful theories survive—the ones which *in fact* latched on to actual regularities in nature.

Van Fraassen's analysis is routinely mentioned as the source of a rejoinder to the no miracles argument in the form of a riposte I have often heard in discussions: "What has explanation to do with truth?!" No doubt this riposte appears somewhere in print, but I could not find it. A less colorful version appears in Fine (1986, p. 115):

> … to argue for realism one must employ methods more stringent than those in ordinary scientific practice. In particular, we must not beg the question as to the significance of explanatory hypotheses by assuming that they carry truth as well as explanatory efficacy.

To begin, I do think that the anti-explanationist argument has merit. It derives from the overall weakness of the inductive argument form known as "inference to the best explanation." The cogency of the rule depends on explanations carrying inductive force. The difficulty is that notions of explanation are so varied and are even so poorly articulated that a tight connection between explanatory prowess and truth is precluded.

---

[47] Appealing as is the rhetoric of this response for its vivid imagery, its weakness has been apparent to me from the outset. It fails if the survival of a scientific hypothesis is enhanced, even only in some probability, if it latches onto truths deeper than superficial regularities. Nothing in van Fraassen's arguments precludes it, beyond his voluntary decision to deny it.

Consider, for example, that some processes are so complicated that only an opaque theory can adequately describe them in full, quantitative detail. This full account is intractable as an explanation. In such cases, we can still provide serviceable explanations in intentionally simplified narratives. An example is that changes in greatly complex weather patterns are often explained with folk-meteorological conceptions like a "cold front." That is, the explanations depend on the introduction of falsities into the narrative.

This critique of the import of explanation in the no miracles argument has merit. The mistake, however, is to stop with it. For it is common for scientists to claim support for their theories on the basis of the theory's explanatory prowess. In Norton (2021, Ch.8-9), I gave an extended analysis of canonical examples of inferences to the best explanation in science. I found that they did not depend on any special notion of explanation. Rather, they had a structure that succeeded without it. Each example consisted of a favored theory that proved adequate to the evidence. It was compared with a foil that failed, either through refutation by the evidence or by the accumulation of an unsustainable inductive debt.

The outcome is that the no miracles argument cannot be employed as a general argument for scientific realism. It proceeds at too general a level. The notion of success is too broad and the notion of explanation enjoined too vague. However, the argument can work if it is implemented in individual cases, in which the cogency of the real argument, independent of explanatory flourishes, it recoverable. My narrative in (2021, Ch. 8-9) gives many examples. Darwin, for example, invoked natural selection as the explanation for the diversity and adaptation of species. His case for natural selection can be reconstructed without needing any special notion of explanation.

4. The All or Nothing Argument

These last sections catalog skepticism about specific forms of inductive inference. There is a simpler dismissal of inductive inference in the antirealism literature: it is just to ignore it. This response is invited if Big-R Realism is understood to seek an unattainable goal: the assured truth of at least some parts of science. The easy response correctly reports the fallibility of inductive support: *assured* truth is unattainable. That response can then be amplified, incorrectly, into the fallacy of a false dilemma: either truth is assuredly attainable or we cannot attain truth at all. Since the first horn is precluded, we arrive at the antirealism of the second.

This false dilemma seems to be how Chang built the case for his antirealist "realism for realistic people." (See Subsection 7.5 below.) He wrote: (2022, p. 2)[48]

> Seeking absolute truth is not an operational ideal–there is nothing we can actually do in order to approach that ideal. According to the common picture of scientific knowledge, science should give us the true picture of the reality that exists well-formed 'out there' completely independently of our conceptions and our experiences. But such 'reality' is not accessible to us and there are no actual methods by which we could attain assured knowledge about it.

And then: (pp. 6-7, his emphasis)

> I wish to take philosophers' attention away from standard scientific realist attempts to show that the impossible is *somehow* possible, namely that empirical science can attain assured truths about what truly goes beyond experience. As I will argue in Chapter 2, it is time to accept the fact that we cannot know whether we have got the Truth about the World (and that such thoughts are perhaps not even meaningful). Scientific realists go astray by persisting in trying to find a way around this fact, while antirealists make the mistake of engaging unproductively with that realist persistence.

We pass between the horns of the dilemma by accepting that assured truth is unattainable, but that we can secure very strong inductive support for the truth of contingent propositions in science. It suffices.

5. Pessimistic Meta-induction

The most common argument for antirealism is the pessimistic induction or, more accurately, the pessimistic meta-induction. In its barest form, it is an overreaction to the fallibility of inductive inference. We can find many occasions in the history of science in which our favored theories failed in some way. Therefore, we are to expect our present theories to suffer the same fate. That is, we are to adopt an inductive skepticism in which we no longer have confidence in the inductive inferences supporting, for example, the truth of propositions or the reality of entities in our present science. From some failures historically of inductive support for

---

[48] Rescher (1987, p. 41) accuses instrumentalists of arriving at their view by a similar misappraisal of the import of what he calls the "fallibility and corrigibility of science."

some aspects of scientific theories, we are to leap imprudently to a complete lack of confidence in our present scientific theories and perhaps even to the stronger expectation of their ultimate and enduring failure. Such is the summary of the argument given by Stanford (2006), who identified the argument in the writings of Henri Poincaré: (p. 7, his emphasis)

> Contemporary philosophers of science call this argument the *pessimistic induction*: its central idea is that the scientific theories of the past have turned out to be false despite exhibiting just the same impressive sorts of virtues that present theories do, so we should expect our own successful theories to ultimately suffer the same fate.

The paper most often credited with advancing this pessimistic argument is Laudan's (1981),[49] "A Confutation of Convergent Realism." Laudan's paper is rich in examples of failed theories. After his discussion of other failures, he gives a list of further failures: (p. 33)

> –the crystalline spheres of ancient and medieval astronomy;
> –the humoral theory of medicine;
> –the effluvial theory of static electricity;
> –"catastrophist" geology, with its commitment to a universal (Noachian) deluge;
> –the phlogiston theory of chemistry;
> –the caloric theory of heat;
> –the vibratory theory of heat;
> –the vital force theories of physiology;
> –the electromagnetic aether;
> –the optical aether;
> –the theory of circular inertia;
> –theories of spontaneous generation.

It is followed by a confident boast:

> This list, which could be extended *ad nauseam*, involves in every case a theory which was once successful and well confirmed, but which contained central terms which (we now believe) were non-referring.

[49] The term "confutation" is uncommon and means "disproof" according to the 1912 Websterian dictionary.

This historical tale of woe provided the basis for Laudan's disputing of scientific realism. He argued that this history shows that the success of a scientific theory provides no basis for scientific realism in two specific forms. A theory's success is independent from whether the theory's terms successfully refer and from whether its propositions are approximately true. He concluded: (p. 47)

> The fact that a theory's central terms refer does not entail that it will be successful; and a theory's success is no warrant for the claim that all or most of its central terms refer.
>
> The notion of approximate truth is presently too vague to permit one to judge whether a theory consisting entirely of approximately true laws would be empirically successful; what is clear is that a theory may be empirically successful even if it is not approximately true.

It is notable that Laudan's "Confutation" does not characterize his argument as an induction over the history of science. The term "pessimistic" also does not appear. Both terms appeared in contemporary characterizations of the argument. Newton-Smith (1981, p. 14) described the "pessimistic induction." Putnam (1978, p. 25) called it a "meta-induction." They were soon combined into the "pessimistic meta-induction," which became a familiar label. French (2000) mentions in passing "the so-called pessimistic meta-induction." This latter term is, in my view, the better name for it. The inference does not dispute any particular result in present science. Rather it disputes our methods for arriving at these results. Since these methods are inductive, it is a skeptical induction about induction, that is, a meta-induction.

Although Laudan does not make this meta-inductive character explicit, it is present in his assertion above: "a theory's success is no warrant for the claim that all or most of its central terms refer." The "warrant" refers to the inference from empirical success to referential success, which is, prudently, inductive. Laudan disputes the generality of this inference by collecting instances of its failure. That is, he is using an induction over these instances to dispute the further inference indicated by "warrant."

Laudan's formulation of the pessimistic meta-induction is just one of many variants in a literature that grew large. A notable variant is Stanford's "problem of unconceived alternatives." The insecurity of inductive support for our present science lies, according to the argument, in our historically affirmed inability to conceive of the unrecognized alternatives that would be better supported. Stanford introduces his argument as: (2006, pp. 17-18, his emphasis)

> … any real threat from the problem of underdetermination comes not from the sorts of philosophically inspired theoretical alternatives that we can construct parasitically so as to perfectly mimic the predictive and explanatory achievements of our own theories, but instead from ordinary theoretical alternatives of the garden variety scientific sort that we have nonetheless simply not yet managed to conceive of in the first place. I will call this worry *the problem of unconceived alternatives*, and although it has historically received far less attention than either the search for empirical equivalents or the traditional pessimistic induction, I will suggest that it ultimately deserves far more.

This sharply formulated objection deserves an immediate response, before further responses are developed in the subsection that follows. Stanford is correct to note that we *might* be mistaken in our present assessment of the import of empirical evidence because of unconceived alternatives. That is, inductive support is fallible. It does not follow that we *should* infer that we are mistaken. Fallibility is not failure.

6. The Inductive Failure of the Pessimistic Meta-induction

The literature on the pessimistic meta-induction is large. It contains further variant forms of the argument and many analyses of its evident shortcomings.[50] My concern here is the extensive strand of criticism that assesses the pessimistic meta-induction as an inductive inference. In short, the pessimistic meta-induction fails in every aspect in which an inductive inference can be assessed. It is one that fails badly.

Mizrahi (2013; 2016; 2020, Section 5.1) identified numerous failings in the pessimistic meta-induction. If we are to conceive of the induction as an inference from a sample to a population, then it is elementary that the sample must representative. Here the sample fails to be representative of the population of present theories to which we are to infer. For most of the sample comes from earlier eras in science which do not resemble the present era. Further the sample must not be cherry-picked. That is, the individuals in the sample must not be selected because they manifest the characteristic sought. Developments of the pessimistic meta-induction

[50] See Wray (2015) for a brief survey.

are explicitly biased. To establish the history of science as a "graveyard"[51] of theories, they selectively report instances of failure.

To bring an induction from the history of science closer to proper inductive methods, Mizrahi (2013) assembled an inventory of past science labeled as "theory" or "law." He used a random number generator to pick a sample of 40 from each. He found 15% of the theories and 12.5% of the laws to be now abandoned. That is most have not been abandoned. Mizrahi derived an optimism and not pessimism from the associated success rates. In a susbsequent analysis (Mizrahi, 2016), drew a sample of 40 from an inventory of sciences in which the terms "theory," "entity" and "posit" appear. The results were similarly optimistic. Mizrahi (2013, p. 3225) concluded "…the pessimistic induction is a fallacious argument and the counterexamples to which it points miss their intended target."

Mizrahi left without sustained challenge the assumption that it is appropriate to treat the historical induction as an inference from a simple frequency to a population frequence. He gave one example as: (2013, p. 3222)

> 15% of sampled theories are abandoned theories (i.e. considered false).
>
> Therefore, 15 % of all theories are abandoned theories (i.e. considered false).

It is quite doubtful that this simple argument form is properly applied here. Indeed, the material theory of induction requires that it can only be applied if there are background facts, specific to the domain, that authorize it.

Shech (2019) examined whether suitable background facts can be identified that would warrant the inference. Those facts would have to prevail over many sciences and over a great expanse of time, from the historical instances taken as premises to the present sciences whose cogency is to be concluded. It is not so surprising that Shech's careful examination found no such warranting facts. They did not reside in a uniformity of methods used by the sciences over time and field, since there is no such uniformity. They did not reside in some uniformity of facts about the investigating scientists themselves; and again not in some uniformity of explanatory modes used. The outcome is a failure of the pessmistic meta-induction, not just in the details Mizrahi catalogued, but simply as inductive argument. There is no warrant for its use. Shech concluded: (2019, p. 928)

[51] A label reported by Psillos (1999, p. xxi).

> I have argued that historical inductions such as the (global) PMI and the problem of unconceived alternatives dissolve if we work with the material theory of induction. The reason is that we lack the material postulates needed to license the pessimistic inference: the great heterogeneity of case studies from the history of science of conceiving, constructing, and discovering (explanatory and predictively successful) theories along with the abundant variety of context that scientists find themselves in and traits that they exhibit make it unlikely that any commonality will be found strong enough to authorize the induction.

These problems still have not exhausted the deficiencies of the pessimistic meta-induction. What might well be the most serious failure has been identified by Psillos (2022, §2.4). The pessimistic meta-induction is supposed to establish that we cannot infer reliably to the truth of all present theories. The grave weakness is that the pessimistic meta-induction just ignores whatever may be the present evidence for the truth of our best present science. Any assessment of the truth of present science *must* take into account that present evidence. It is, by any reasonable standard, the strongest and most reliable evidence.

The observed expansion of the galaxies and the cosmic microwave background radiation provide strong evidence for big-bang cosmology. The extensive repertoire of atomic spectra provides strong evidence for quantum theory. Paleontological and embryological observations provide strong evidence for evolutionary theory. If we are to accept the conclusion of the pessimistic meta-induction, we would have to overrule the inductive support provided by all these specific experiences, tailored as they are to each specific theory. We would have to accept that the pre-emptive generality of the pessimistic meta-induction is the stronger factor.

If we contemplate the practicalities of such acceptance, its absurdity becomes apparent. Are we to press a modern astronomer to abandon confidence in the truth of planetary heliocentrism because Ptolemy was a geocentrist; or that a modern chemist should abandon the truth of the oxygen theory of combustion because the failed phlogiston theory enjoyed a brief period of celebrity? At best, that past efforts have failed can serve only in a secondary role, evidentially. It counsels caution in the assessment of the primary evidence for a present science.

There is a final irony in the pessimistic meta-induction. The whole burden of the pessimistic meta-induction is to assail inductive inference itself. That is, its goal is not just to cast doubt on *some particular* inductive inference, in science. It is to impugn the cogency of *all* *present* applications of inductive inference in science. If we are to identify one place where inductive inferences are done well, it is in our best science. The pessimistic induction is supposed to convince us that even here in this best possible case, we cannot rely on inductive inferences. If that is so, then we should have no confidence at all in the weaker inductive inferences of the pessimistic meta-induction induction. The extent to which it succeeds is the extent to which it fails.

7. Antirealist "Realisms"

Some views that carry the label "type-X realism" are, on closer inspection, versions of antirealism. They are offered as repairs to realism but, on closer examination, prove to be proposals that negate it. An essential part of that negation are forms of inductive skepticism. Two versions are sketched here.

7.1 Perspectival Realism

So-called "perspectival realism" is a clear case of type-X realism in which the conception has migrated sufficiently far from the original intuition that the label "realism" is a misnomer. According to it, objective truths in science are beyond the reach of scientists. It is not that scientists need to try harder. It is, in this view, an unbreachable barrier of principle. The best scientists can have, in Giere's (2006) formulation, are models that provide perspectives. The models cannot be said to be true of some target system, since we cannot expect a perfect fit. Talk of "similarity," Giere (2006, p. 66) finds, is more appropriate than talk of "truth."

According to perspectival realism, this assessment is conveyed in a quote that Giere imagines being given by a scientist, enlightened by perspectival realism: (p.6)

> "According to this highly confirmed theory (or reliable instrument), the world seems to be roughly such and such." There is no way legitimately to take the further objectivist step and declare unconditionally: "This theory (or instrument) provides us with a complete and literally correct picture of the world itself."

To articulate the view further, Giere (p.16) tells us that "Color vision provides the best exemplar I know for the kind of perspectivism that characterizes modern science." The human visual system provides the perspective in the exemplar: (p.27)

> Why should a surface with a given surface spectral reflectance be called "yellow"? Without reference to the particular characteristics of the human visual system, there is no physical basis whatsoever for this identification.

Presumably we are to treat the dependence of our *perception* of color on our visual system as a good analog for how all *scientific claims* are dependent on a perspective, and ineliminably so. If that is the analogy, it is obviously a poor mapping between the exemplar and the science. An analysis of human visual experience allows us to infer from our perception of yellow to an objective property of the surface: it reflects the light in the frequencies corresponding to yellow. The identification of this objective property by inferences is the better analog for claims in science.

This last remark is no mere quibble about a misused analogy. It reveals the mismatch between perspectival realism and the inferential practices of science. The specific measurements that manifest in different measured values will vary according to the instruments used. In that sense they are perspectival. It is routine in science to free those measurements from their perspectives, that is, from their attachment to this instrumentation by inductive inferences.

7.2 Perspectival Inductive Skepticism

Giere is right that our representation of some—but not all—systems come in different, limited perspectives. Each in some cases captures some elements of the system but omit or oversimplifies or misrepresents others. They cannot be combined tractably into a single, complete representation. A familiar example is maps of a city. Each city is represented by many different maps. There are street maps for road traffic; maps that highlight biking and walking paths; and maps that display public transit routes. Other maps display public utilities, such as electric power, water supply, drainage and sewage lines. Climate science provides another example closer to our science. Models of the Earth's climate identify one aspect of the processes in the climate and idealize or simplify away other elements. These models, as reported in Jeevanjee et al. (2017), cover atmospheric fluid, Earth's rotational, oceanic, surface, convection and radiative aspects of the climate.

Where Giere errs, however, is in generalizing what applies to some systems investigated in science to all scientific investigation. This expansion is based on a mistaken, general claim about inductive inference. It is that the import of evidence from different sources cannot be combined. Their import is always restricted narrowly to the perspective associated specifically with that form of evidence.

The arbitrariness of this perspectival compartmentalization is a form of inductive skepticism. It is contradicted by small-e empiricism and familiar inductive inferences. Take the case in which we have different experiential processes emanating from the same source. Consider, for example, an electric discharge from the clouds in stormy weather. We learn of it through the electromagnetic process of a lightning flash; through the vibratory propagations of the sound of thunder; and, if the wind blows towards us, the material propagation of pungent ozone gas. Each item of evidence combines to affirm that the others were the result of the same electric discharge. For example, the pungent smell of ozone is evidence for an ozone-producing electric discharge; and we know from electrochemistry that this same discharge produces light, which affirms that the lightning results from the same discharge.

7.3 How LIGO Merges Perspective

We can combine perspective and routinely do so. An example is the recent triumph of LIGO's gravitational wave detector. On August 17, 2017, it detected a binary neutron star merger in the event labeled GW170817. It was a distinctive detection insofar as the detection arose in both gravitational and electromagnetic wave channels. Here is the collaborative's summary of the result: (Abbott et al., p.2)

> Here, we report on the global effort that led to the first joint detection of gravitational and electromagnetic radiation from a single source. An ∼ 100 s long gravitational-wave signal (GW170817) was followed by an sGRB (GRB 170817A) and an optical transient (SSS17a/AT 2017gfo) found in the host galaxy NGC 4993. The source was detected across the electromagnetic spectrum—in the X-ray, ultraviolet, optical, infrared, and radio bands—over hours, days, and weeks. These observations support the hypothesis that GW170817 was produced by the merger of two neutron stars in NGC4993, followed by an sGRB and a kilonova powered by the radioactive decay of r-process nuclei synthesized in the ejecta."

Compare this summary with Giere's (2006, p. 92) assessment of a similar case:

> … the same object can often be observed from several different perspectives, such as a nearby galaxy observed by both optical and radio telescopes. This is indeed good evidence that there is "something" there, but that is scarcely knowledge in the

objectivist sense. The knowledge we get comes from one perspective or another, not from no perspective at all. Multiplying perspectives does not eliminate perspectives."

Giere's perspectival realism prohibits the collaborative from affirming anything beyond the following collection of perspectival claims:

We have detected a binary neutron star merger in the gravitational wave perspective.

We have detected a binary neutron star merger in the X-ray perspective.

We have detected a binary neutron star merger in the ultraviolet wave perspective.

We have detected a binary neutron star merger in the optical wave perspective.

We have detected a binary neutron star merger in the infrared wave perspective.

We have detected a binary neutron star merger in the radio wave perspective.

The major claim of the collaborative is a breach of this perspectival prohibition. The collaborative uses the measurements in each perspective as observational support for the objective fact of a binary neutron star merger in the host galaxy NGC 4993; and properly so. This well-founded objective declaration is precisely what perspectival realism prohibits.

7.4 A Situated Perspectival Realism.

Massimi (2018, 2022) has developed Giere's perspectivalism by emphasizing that perspectives are also culturally and historically situated: (2022, pp. 5-6, Massimi's emphasis)

Scientific perspective (sp): A scientific perspective sp is the actual—historically and culturally situated—scientific practice of a real scientific community at a given historical time. Scientific practice should here be understood to include: (i) the body of *scientific knowledge claims* advanced; (ii) the experimental, theoretical, and technological resources available to *reliably* make those scientific knowledge claims; and (iii) second-order (methodological epistemic) principles that can *justify* the *reliability* of the scientific knowledge claims so advanced.

The earlier account affirmed Giere's denial of objectivity: (Massimi, 2018, p. 165)

There cannot be an objective, unique, true description of the way the world is as soon as we acknowledge that our scientific knowledge is always from a specific vantage point – either in the sense of (1) [historically situated] or (2) [culturally situated] or both.

By expanding perspectives to include cultural and historical factors, this form of perspectival realism increases the compartmentalization of the import of evidence and the inductive skepticism associated with it. It would, for example, block the inference reported in *small-e empiricism,* Chapter 11, from historical reports drawn from different cultures in China, Japan, the Arab world and Europe to the simple fact of the supernova of 1006AD.

Overall, perspectivalism cannot be called a "realism," if we are to adhere to the original intuition of scientific realism. Because of its insistence that objectivity is impossible as a matter of principle, perspectival realism is better conceived as a form of antirealism.[52]

7.5 Realistic Realism

Chang's (2022) *Realism for Realistic People* is, superficially another instance of the migration of the notion of scientific realism. "Realism" is the first word in the title. It becomes apparent almost immediately in the text that the policy it advocates for "realistic people" is simply to abandon the prevalent conception of scientific realism for a form of antirealism. On its first page (p.1), scientific realists are disparaged as those who have "succumb[ed] to an authoritarian or even autocratic impulse." The text continues to renounce as a "widespread conceit that modern science has *basically* the right answers, or at least the right methods for getting the right answers." This renunciation of scientific realism is thorough. It disallows even the methods of inductive inquiry that, in small-r realism, comprise the foundation of scientific realism. This conceit of realists, we are told, is "maintain[ed] … like an article of faith."

After this introduction, readers can have little doubt about Chang's distain for scientific realists and that the term "realistic" in his title must be read ironically. A new sense is intended. That Chang is trying to redefine "realism" is quite explicit. He tells us so by quoting from his own, earlier remarks: (p.7, his emphasis)

> I am trying to carry out a task that I already advocated a decade ago: 'I think *realistic* people (including most empiricists and pragmatists) should re-claim the label of "realism"!'

The direction of the redefinition is given in the work's subtitle: *A New Pragmatist Philosophy of Science*. Its details go beyond our present concerns. Very briefly, this new pragmatist realism is

[52] Massimi's (2018, p. 171; 2022, p. 179) response is to introduce a conception of "perspectival truth," which is truth contextualized to different perspectives.

centered on the notion of (p.4) "operational coherence," which is "[v]ery roughly … a matter of making elements of our activities fit together harmoniously so that our aims may be achieved." And that:

> …what we should mean by something being 'real' is that it can be employed in coherent activities that rely on its existence and its basic properties. Truth and reality conceived in such a way are attributes grounded in our activities.

8. Conclusion

Antirealism thinking is correct to resist Big-R Realism if it is understood to include claim that science has arrives *assuredly* at final truths. Here it aligns with small-r realism/small-e empiricism Where it errs is in going beyond small-e empiricism and denying that science has had any success in attaining fallibly truths beyond experience. It is an overreaction that compels it to adopt a skeptical anti-inductivism. For otherwise, the inductive import of the extensive bodies of evidence for our mature science would contradict the antirealism.

References

Abbott BP, et al. (LIGO, Virgo and other collaborations) (2017). "Multi-messenger Observations of a Binary Neutron Star Merger," *The Astrophysical Journal*. 848: L12.

Chang, Hasok (2022) *Realism for Realistic People*. Cambridge: Cambridge University Press.

Fine, Arthur (1986) "The Natural Ontological Attitude," Ch. 7 in *The Shaky Game: Einstein, Realism and the Quantum Theory.* Chicago: University Chicago Press.

van Fraassen, Bas C. (1980) *The Scientific Image*. Oxford: Clarendon Press.

French, Steven (2000) "Identity and Individuality in Quantum Theory" *Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/archives/spr2000/entries/qt-idind/

Giere, Ronald N. (2006) *Scientific Perspectivism*. Chicago: University of Chicago Press.

Jeevanjee, Nadir; Hassanzadeh, Petram; Hill, Spencer; and Sheshadri, Aditi (2017), "A perspective on climate model hierarchies," *Journal of Advances in Modeling Earth Systems,* **9**, pp. 1760–1771, doi:10.1002/2017MS001038.

Laudan, Larry (1981) "A Confutation of Convergent Realism," *Philosophy of Science*, **48**, pp. 19-49.

Massimi, Michela (2018) "Perspectivism," Ch. 13 in Saatsi (2018).

Massimi, Michela (2022) *Perspectival Realism*. New York: Oxford University Press.

Mizrahi, Moti (2013) “The pessimistic induction: a bad argument gone too far,” *Synthese*, **190**, pp. 3209-3226.

Mizrahi, Moti (2016) “The History of Science as a Graveyard of Theories: A Philosophers’ Myth?” *International Studies in the Philosophy of Science*, **30**, pp. 263-278.

Mizrahi, Moti (2020) *The Relativity of Theory: Key Positions and Arguments in the Contemporary Scientific Realism/Antirealism Debate*. Springer.

Newton-Smith, William (1981) *The Rationality of Science.* London and New York: Routledge.

Norton, John D. (2008) “Must Evidence Underdetermine Theory?” in *The Challenge of the Social and the Pressure of Practice: Science and Values Revisited*, M. Carrier, D. Howard and J. Kourany, eds., Pittsburgh: University of Pittsburgh Press, pp. 17-44.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Psillos, Stathis (1999) *Scientific Realism: How science tracks truth.* London and New York: Routledge.

Psillos, Stathis (2022) “Realism and Theory Change in Science,” *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), <https://plato.stanford.edu/archives/fall2022/entries/realism-theory-change/>.

Putnam. Hilary (1975-76) “What Is ‘Realism’?” *Proceedings of the Aristotelian Society* , New Series, **76**, pp. 177-194.

Putnam, Hilary (1978) *Meaning and the Moral Sciences*. London: Routledge and Kegan Paul.

Rescher, Nicholas (1987) *Scientific Realism*: *A Critical Reappraisal.* Dordrecht: Reidel.

Shech, Elay (2019) “Historical Inductions Meet the Material Theory,” *Philosophy of Science*, **86,** pp. 918–929.

Smart, J. J. C. (1963) *Philosophy and Scientific Realism*. London: Routledge and Kegan Paul.

Stanford, Kyle (2006) *Exceeding Our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: Oxford University Press.

Stanford, Kyle, "Underdetermination of Scientific Theory", *The Stanford Encyclopedia of Philosophy* (Summer 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), <https://plato.stanford.edu/archives/sum2023/entries/scientific-underdetermination/>.

Wray, K. Brad (2015) “Pessimistic Inductions: Four Varieties,” *International Studies in the Philosophy of Science*, **29**, pp. 61-73.

10. Against Metaphysics

Chapter 10
Against Metaphysics

1. Introduction

The rejection of metaphysics has been a fixture in empiricist thought since at least the positivist precursors of modern empiricism in the nineteenth century. It became a rallying cry for the logical positivists of the early twentieth century. Small-e empiricism retains this rejection but without the extremism of the logical positivists' denial of meaning to propositions in metaphysics. The primary basis of this rejection is the familiar empiricist complaint. There is no way to learn contingent truths other than through experience. In so far as metaphysics seeks such truths prior to the consideration of experience, it is ill-fated from the outset.

This principal empiricist objection, developed in Section 2, is to non-empirical metaphysics specifically, which holds that it can recover foundational truths of the world without drawing on experience. Small-e empiricism entails the inevitable failure of this project. For once experience has been abandoned, the only resource left to non-empirical metaphysics is the logical consistency of its posits. Mere logical consistency is too permissive and allows many competing narratives. Non-empirical metaphysics has no further resources that can distinguish among these competing narratives. The inevitable result is that debates in metaphysics persist indefinitely. Criteria such as simplicity and explanatory power fail to resolve these debates non-empirically, for, as it is maintained in the material theory of induction, such criteria are only potent if they are surrogates for specific background facts that are in turn supported empirically. Section 3 illustrates such languishing debates in cases in which the metaphysical claims are so structured as to escape empirical scrutiny. They concern claims for the reality of possible worlds, the reality of the past and future and Platonic methods in thought experiments.

Section 4 addresses the emergence of so-called "scientific metaphysics." It seeks to derive results in metaphysics empirically after the manner of empirical science, in contradiction with the traditional conception of metaphysics as prior to empirical science. For an empiricist, its project is most welcome, in so far as it seeks to employ empirical methods. They are the best and only way to secure contingent results of foundational importance. What is unwelcome is the label “metaphysics.” For what results is the identification of the empirically derived, foundational content of individual sciences as metaphysics. The one term, metaphysics, now embraces two different projects: the methodologically ill-fated, non-empirical project of traditional metaphysics and the fertile, empirical project of the sciences. The first failed project accrues some unearned legitimacy by its association with the second, while the second is compromised by its association with the first.

Section 5 presents a second objection to metaphysics of comparable importance that has a lesser presence in the literature. A foundational assumption of traditional metaphysics is that its subject matter is common to all the sciences. The claim that routinely follows is that we can abstract away all the details of the concepts that distinguish their roles in particular sciences and be left with a substantial general metaphysics of these concepts that applies universally. That this abstraction returns substantial results cannot be legislated in advance. It must be established on a case-by-case basis. The assumption fails for two significant cases discussed in Sections 6 and 7: the metaphysics of possibility and the metaphysics of causation. For them, the abstraction is fatal to the ambitions of traditional metaphysics.

Section 8 continues the scrutiny of the metaphysics of causation by reporting a striking example of the misplaced ambitions of investigations carried out prior to empirical research. We are supposed to learn from these investigations of the impossibility of closed causal loops, of causation backward in time and of time travel backwards in time. Empirical inquiry into these questions is declared pointless since they have been decided a priori.

The following Section 9 responds to a common objection that empiricism cannot refute metaphysics because empiricism is itself a metaphysics. This *tu quoque* objection fails. Merely labeling the empirically founded arguments of this chapter as metaphysics reveals no flaw in the arguments. The concluding Section 10 reaffirms that non-empirical metaphysics must fail in its efforts to recover foundational truths prior to experience precisely because it forgoes the guidance of experience.

2. The Arbitrariness of Non-Empirical Metaphysics

2.1. Non-Empirical Metaphysics

The difficulty facing an empiricist critique of metaphysics is to identify the target of criticism. Accounts of metaphysics routinely begin with a lament that it is difficult and perhaps even impossible to define metaphysics.[53] My impression is that this reflects the fluid scope of metaphysics. Its subject matter and methodology have varied historically and still varies today from metaphysician to metaphysician.

We can, however, find sufficiently serviceable formulations. Lowe (2002, p.v) offers "a fairly traditional and still very widely shared [conception]" which is:

> … that metaphysics deals with the most profound questions that can be raised concerning the fundamental structure of reality. According to this conception, metaphysics goes deeper than any merely empirical science, even physics, because it provides the very framework within which such sciences are conceived and related to one another.

What is important for present purposes is that, at least a sizeable community of metaphysicians, hold that metaphysics is inherently non-empirical. Lowe's survey (2002) develops the characterization:

> But metaphysics is not at heart an empirical science—it does not typically appeal to experimental or observational data in support of its claims. (p.2)
>
> … metaphysical knowledge-claims are akin to mathematical knowledge-claims, which are likewise held to be non-empirical, in the sense that they are not answerable to empirical evidence for their support or confirmation. (p. 10)

The final assessment is that there is a division of labor and methodology between metaphysics and science. Metaphysics proceeds non-empirically to provide the framework of possibilities. After metaphysics has laid this essential groundwork, science proceeds empirically with the task of picking out which possibility is actualized: (p. 11, his emphasis)

> Thus metaphysics, like mathematics, does have a non-empirical subject-matter, to the extent that it is the intellectual discipline whose concern it is to chart the *possibilities* of real existence. Metaphysics is concerned to discover what the totality of existence *could* embrace: that is to say, what categories of entities could exist and which of them could co-exist. Having charted the possibilities, the question will remain as to which of many mutually incompatible possibilities for the fundamental structure of reality *actually* obtains-and this question can only be

[53] The opening line of van Inwagen et al. (2025) is "It is not easy to say what metaphysics is."

> answered, if at all, with the aid of empirical evidence, and then only tentatively and provisionally.

This distinction between metaphysics as non-empirical and science as empirical is not uncommon. Kit Fine for example states it as: (2011, p.9)

> It [metaphysics] is distinguished, in part, from physics and other branches of science by the a priori character of its methods. The claims of science rest on observation; the claims of metaphysics do not, except perhaps incidentally. Its findings issue from the study rather than from the laboratory.

Section 6 below provides a striking example in the metaphysics of causation. Its results are recovered antecedent to empirical investigations, yet they are supposed to restrict what empirical investigations can recover. None, we are assured, can affirm empirically closed causal loops, causation backward in time and time travel backward in time, for prior metaphysics has precluded them.

2.2 The Empiricist Critique: An Inevitable Failure

In so far as metaphysics seeks deep truths of the world non-empirically, then the empiricist critique is immediate and automatic. The *only* way we can learn of these deeper truths of nature is empirically, through experience. This restriction is the defining characteristic of small-e empiricism as developed above in Chapter 2. If non-empirical metaphysics seeks to inform us about contingent facts in the world that are not accessible to experience, it necessarily fails.

What happens if we persist with non-empirical investigations that seek to establish contingent facts that are not accessible to empirical investigation? Once the constraints of experience are removed, according to small-e empiricism, all that remains is the constraint of logical consistency. (There have been proposals for other constraints. I will discount them in Sections 2.5 below.) Mere logical consistency is too weak a constraint to produce definite results. The best outcome that can be achieved is an internally consistent system of putative facts. Logical consistency admits very many such systems that differ widely in their factual content. We have no means to decide which among them is actual. One system may prevail in the literature through the rhetorical prowess of its promoters. It may remain so until their efforts are eclipsed by the prowess of their detractors. From this perspective, it is not at all surprising that results in non-empirical metaphysics are unstable and that endless, inconsequential debates persist in non-empirical metaphysics. The field has no enduring means to decide them.

This instability applies to debates over contingent propositions that are inaccessible to experience. If the propositions are accessible to experience, then empirical investigations can override the a priori considerations. A striking example is recalled in the discussion of causal metaphysics below, in Section 5.5. No lesser figure than the great Isaac Newton declared on a priori grounds that gravitation could not be an action at a distance. Extensive empirical investigations in the nineteenth century gave strongest support to gravitation as an action at a distance. These empirical investigations simply eclipsed a priori hesitations about action at a distance. It was established, at least temporarily, as empirical science.

This power of empirical evidence to decide questions is routine in the empirical sciences. In areas in which sufficient empirical evidence is available, a unique science emerges and remains stable. That stability is not lost because someone finds a more convincing way to reframe the science. Mature sciences with rich empirical foundations easily resist such efforts. It takes novel empirical evidence to overturn them.

2.3 Conceptual Analysis

Might it be that non-empirical metaphysics can escape this empiricist objection when the analysis is limited to conceptual analysis? Here Jackson has defined conceptual analysis as: (2021, p. 396)

> … conceptual analysis is in the business of clarifying and explicating what it takes, or if you like, how something has to be, to fall under some concept or other.

He has defended conceptual analysis as a part of metaphysics.

Conceptual analysis escapes the empiricist critique as long as it is limited to *defining* the conditions under which some concept would be instantiated in reality. The cost of the escape is high. In this limited capacity, conceptual analysis is an a priori activity. Offering definitions does not take us to the goal of metaphysics, "the fundamental structure of reality." The same difficulty facing non-empirical metaphysics remains. The a priori activity of defining concepts is constrained only by logical consistency; and logical consistency is a weak constraint that admits vastly too many concepts. For any concept that may suggest a factual constraint on reality, there will be counter-concepts that negate them; and many more that mediate between them. The sheer enormity of choices is just too great for tractable exploration by an armchair theorist. Logical consistency gives us no way to decide among them.[54]

Conceptual analysis can illuminate the fundamental structure of reality if, in addition, we have some confidence that a select few of these many concepts are or may be instantiated in reality. That confidence cannot come from reflecting on concepts a priori. That confidence can only come, empiricism insists, from engagement with reality through experience. In this empirically informed practice, conceptual analysis no longer belongs within non-empirical metaphysics, but is an instance of an investigation, informed empirically in essential aspects.[55]

In practice, conceptual analysis has never been rigidly and exclusively a priori. Concepts are conceived and introduced in response to real puzzles in experience. How are we to think of the present, past and future? What can we say about the stuff around us? Things happen? Which? May they? Or must they? All these scenarios provide a posteriori guidance as to which few concepts of the many logically admissible may draw experiences into a unified account. As a result, these factual conditions drawn from experience are reflected in the concepts themselves so that the overall analysis has an ineliminable, a posteriori component.

In short, in so far as conceptual analysis is purely a priori, it escapes the empiricist objection, but is unable by itself to deliver the goal of metaphysics: to discern the fundamental structure of reality. In so far as conceptual analysis can contribute to this goal, it relies ineliminably on a posteriori elements and thus extends beyond non-empirical metaphysics.

---

[54] For example, we may conceive, as they did in the 19th century, that causation is just determinism, so that a causal world must be factually deterministic. A counter-concept is that of a completely chaotic world in which any past regularity is a coincidence that may not be sustained in the future. A mediating concept proposes some partial dependency of the future on the past. We can eventually decide among the vast array of possibilities, but only by finding which is picked out by experience.

[55] Kölbel (2023) has articulated a version of the tension between non-empirical and empirical conceptual analysis as a dilemma and has explored its consequences, along with a helpful inventory of relevant literature.

2.4 What this Critique is Not

The empiricist critique of non-empirical metaphysics developed here falls within a long-established tradition of empiricist critiques of metaphysics. That tradition inclines towards extremes that I find untenable. My goal has been to present a narrowly confined critique that retracts from the extremes of these earlier critiques. It depends only on the empiricist commitment to experience as the sole basis for learning about the factual world. It is inductively optimistic. We can learn inductively truths of the world beyond experience. We need not disparage propositions in non-empirical metaphysics as meaningless. They fail because of their arbitrariness.

Chapter 2 of *small-e empiricism* reported the repudiation of metaphysics in the nineteenth century positivist tradition of Comte and the sensationalism of Mach. Theirs was a blanket renunciation of everything in science that went beyond simple relations among observables. It is a form of anti-inductivism that contradicts small-e empiricism. Comte renounced abstract forces and Mach decried Absolute Space. Small-e empiricism is inductively optimistic. Abstract forces and Absolute Space are quite admissible, as long as there is good inductive support for them in experience. If an empirical investigation includes them, there is no objection of principle, for the empirical investigation is simply drawing here on what is responsible science.

Chapter 3 of *small-e empiricism* reported that the logical positivists of the early twentieth century introduced a still more extreme version of the empiricist critique in the form of the verifiability criterion of meaning. Initially, a proposition was judged meaningless if it could not be verified, that is proven true, by experience. Later versions moderated this extreme view. A proposition could be meaningful in Reichenbach's probability theory of meaning if some non-zero probability could be assigned to it. To assail the meaning of propositions is an excusable response when applied to the excesses of Carnap's idealistic contemporaries. Otherwise, it paints with too broad a brush. Consider propositions in ordinary discourse that assert the existence of God or that living beings are animated by a life force. We may differ on their truth, but they are clear enough to have a meaning. An account of meaning that cannot return this has failed ordinary discourse.

As already sketched in *small-e empiricism,* Chapter 12, Section 8, the small-e empiricist critique of non-empirical metaphysics takes no general position on the meaningfulness of the metaphysical propositions. Some may well be meaningless and other meaningful. They fail, not for lack of meaning, but for their arbitrariness.

2.5 Naturalized Methodology

The small-e empiricist critique of non-empirical metaphysics is based on the assumption that there are only two resources available to a scientifically responsible metaphysics: inductive inference from experience and logical consistency. This assumption can be challenged. It is routine for scientists to resort to considerations of simplicity or explanatory power and in ways that may not seem to draw directly on the support of experience.

Might non-empirical metaphysics remain scientifically respectable if it relies on apparently non-empirical criteria such as simplicity and explanatory power? Emery (2023) has explored this possibility in what she calls “methodological naturalism.” (p.4) Its guiding principle is a laudable “whenever possible, use the same methodology that scientists use.” (p. 4) Emery found within this methodology “extra-empirical principles”: (p. 5)

> Examples of potential extra-empirical principles are principles like “Choose the simplest theory that is consistent with the data” or “Choose the theory that provides the best explanation of the data.” These principles allow us to choose between multiple theories—all of which are consistent with the data—on the basis of other features.

The proposal is that such extra-empirical principles may allow us to decide metaphysical questions in a scientifically responsible manner, when the decision lies beyond the reach of empirical data: (p. 5)

> For while it is unclear how the empirical data that we collect would ultimately decide debates about composition or personal identity or possible worlds, a principle like “Choose the simplest theory that is consistent with the data” might very well do so.

This is a view, Emery continues, that can be found already in nascent forms in the writing of established metaphysicians such as Quine, Sider, Zimmerman, Hawthorne and Paul.

The viability of such a naturalized methodology has already been challenged by earlier writers. A notable example is Bryant’s (2020) critique of what she calls “free range metaphysics.” It is metaphysics in which science has a merely nominal or minimal role. Its analysis is constrained by extra-empirical principles comparable to those of Emery. Bryant considers: (p. 197)

consistency, simplicity, intuitive plausibility, and explanatory power

Bryant (Section 2) then argues for the failure of each of these vehicles. She argues, as I do above, that logical consistency is too permissive. Her objections to the remaining vehicles are substantial.

Here I can amplify them through a single consideration of principle. The material theory of induction, defended in my (2021, 2024), is founded on the claim that all such vehicles are epistemically potent only if they are surrogates for particular background facts in the domain of the specific inductive inference. These vehicles may appear to operate independently of empirical support from experience, but that is an illusion. Contingent assertions about the background are only epistemically potent if they are truths; and our only means to ascertain their truth is empirical. This consideration is quite general. It applies to the vehicles of Bryant's list and to any addition we may consider. They will turn out to be epistemically potent only if they depend indirectly on the empirical support of experience.

A brief examination of Bryant's list illustrates how this works. My (2021, Ch.6-7) reviews our failure to find any factual principle of simplicity or parsimony. A celebrated example of the application of simplicity is in curve fitting. The standard approach is to present a hierarchy of functions, such as the constant, linear, quadratic, cubic, etc. and to proceed stepwise up the hierarchy until one is found that fits the data well enough. The method depends critically on the particular hierarchy of functions chosen. Its selection, if done properly, must reflect the factual conditions of the particular problem. My (2021, Ch.8-9) closely examines the argument form of inference to the best explanation. The difficulty facing any general formulation of such an argument form is the range of different accounts of explanation, none of which is entirely satisfactory. In examining a list of notable examples of inference to the best explanation in the history of science, I found no case in which the result depended on an epistemically potent concept of explanation. Rather all the cases fitted a simpler model in which no philosophical account of explanation actually played a role. Finally, we may allow that intuitive plausibility can be epistemically potent, but only if the intuitions are informed by suitable experiences, for which they are surrogates. Otherwise, intuitive plausibility just reflects unguided hunches. The frailty of intuition has been explored further here already in *small-e empiricism,* Chapter 12, Section 7.

3. Languishing Debates

The ways in which non-empirical metaphysics outstrip empirical assessment is quite familiar to readers of the metaphysics literature. This section gives examples of a familiar, indecisive dynamic. Claims are made for which there is no empirical or experiential manifestation. Hence there is no way empirically to determine their truth. Cases can be made for and against each view. However, we lack the resources to mount a decisive case either way. The debate languishes. The mere logical consistency of a case is too weak; and the extra-empirical principles are too elastic to give univocal guidance.

3.1 Possible World Realism

David Lewis (1986) has proposed a realism about possible worlds. Imagine some fantasy world of the wildest fiction. It is real, Lewis insists, for (p. 2, his emphasis) "absolutely *every* way that a world could possibly be *is* a way that some world is." One can of course construct elaborate narratives which favor the reality of a fantasy world with talking donkeys. Lewis defends this possible world realism for its parsimony and explanatory prowess in unification: (p. 4)

> We have only to believe in the vast realm of *possibilia*, and there we find what we need to advance our endeavours. We find the wherewithal to reduce the diversity of notions we must accept as primitive, and thereby to improve the unity and economy of the theory that is our professional concern–total theory, the whole of what we take to be true. … The benefits are worth their ontological cost. Modal realism is fruitful; that gives us good reason to believe that it is true.

A critic can take these same criteria and come to the opposite conclusion. Stalnaker (1976) makes the obvious point that the considerations of parsimony can be found to speak against Lewis' view: (p. 68)

> Even a philosopher who had no qualms about abstract objects like numbers, properties, states and kinds might balk at this proliferation of fullblooded universes which seem less real to us than our own only because we have never been there.

Correspondingly, in a direct response to Lewis' claim above, Stalnaker (1988) inverts Lewis' assessment of the costs and benefits: (p. 121, his emphasis)

> But I think the cost tends to undercut the benefit. Just as the nominalist may say to the fundamentalist Platonist, 'how can propositions about *those* [mathematical] things be relevant to the practice and applications that are why we value mathematics?' so we may say to the modal realist, 'how can the existence of other worlds be relevant to the modal, probabilistic, counterfactual, mental and semantic concepts we use to describe ourselves and our actual environment?'

The decision reduces to differing prescriptions how the predicate "real" can be applied. Competing narratives give competing prescriptions. There is no way to use experience to decide between them, for all the possible worlds, excluding the actual, are by supposition outside our experience.[56]

3.2 Presentism and Eternalism

The same problem of the arbitrariness of the attachment of the predicate "real" arises in the philosophy of time. The simplest form of the debate is between presentism and eternalism. Presentism holds that all that is real is the present moment of time. Eternalism holds that all moments, past, present and future, are equally real. There are further variations. The growing block model attributes reality just to the past and present, but not the future. The controlling difficulty is that all these models agree on the empirical physics of space and time. Whether we judge a past or a future event to be real or not makes no difference to how Einstein's field equations of general relativity are applied and to the mathematical spacetimes that result.

The outcome is familiar. The debate between presentism and eternalism endures because neither side can produce decisive arguments. There is a frequent sense of matched forces. For example, both presentism and eternalism struggle with the problem of the truth-makers of propositions. What is it in the world that makes propositions true? Presentism is challenged to

[56] A complication is that the many worlds interpretation of quantum mechanics authorizes the reality of multiple worlds corresponding to a decomposition into components of the world's Hilbert state vector. I leave open whether their reality enjoys empirical support. Lewis' possible world realism extends well beyond, since it includes possible worlds in which all the fundamental constants of physics have different values.

display the reality that makes propositions about the past true, since past events have no reality. Eternalism is challenged to explain why future events are not presently determined, since their truth-makers in future events are eternally true. In his introductory review of the *Oxford Handbook of the Philosophy of Time*, Callender (2011, p. 3) summarized the opposing positions as:

> … the presentist's temporally impoverished resources are insufficient to allow him or her to say that statements about the past or future are truth-evaluable. The claim that "there were dinosaurs in the Mesozoic Age" is not made true or false by anything on the present time slice, or at least, not obviously so. So the presentist faces a quandary: the claim is true, yet there is nothing that makes it so.

Correspondingly, the eternalist faces the complications of a surfeit of real moments:

> If the future is real, and our choices are "already" made, in what sense can we do otherwise? If I was "always" a philosopher in 2010, did I really have a choice about professions in 1990? Similarly, if the future is real, and the outcomes of our chancy processes are "already" occurrent, then in what sense is the chancy process genuinely chancy? If the coin "always" landed heads at a specific toss, in what sense did that roll have a 50 per cent chance of landing tails?

We need not linger over whether attributing reality to past and future events is meaningful. What matters here is that the attribution manifests in no way in experience; and that leaves presentists and eternalists without adequate means to make a decisive case either way. The debate will continue. Each side will develop an account of truth-making that favors their view, but the mere consistency of each is inadequate to preclude the other.

Many further considerations have been raised in this inconclusive debate. Eternalism finds comfort in the fact that our best theory of space and time, general relativity, employs a four-dimensional spacetime in which all events, past, present and future are treated equally.[57]

---

[57] A complication is that relations of simultaneity in the special theory of relativity differ according to the inertial motion of the observer. This fact can be developed to argue against presentism and for eternalism. I have reviewed why the argument is tendentious in https://sites.pitt.edu/~jdnorton/teaching/HPS_0410/chapters/significance_3/index.html#L5178

Presentism is favored by our subjective experience that the moment of the "now" is distinct from all past and future moments. Each view responds. Presentism objects that eternalism fails to distinguish any moment as this special "now." Eternalism portrays the now awkwardly as a subjective experience of we sentient beings as our brains mete out the moments serially to our consciousnesses.

In a reaction to the arbitrariness of the attachment of "real" to events in spacetime, I proposed (2015) the "burning fuse" model of time. In it, future events are real, but lose their reality when they pass through the present to become past. I wrote it as a spoof and offered what I thought were parodies of the sort of arguments used in the presentism-eternalism debate. I have since found that the model appears in inventories of views about time as the "shrinking block" model, although described (in Ney, 2014, p. 143) as "a rarely held position."

These last considerations are just a small part of the enduring but inconclusive debate over the non-empirical metaphysics of time. The *Oxford Handbook* cited above has 23 chapters and the question of the reality of the past and future appears in one way or another in almost all of them. Neighboring debates concern McTaggart's A and B series for time. The A series attributes the predicates "past," "present" and "future" to events. The B series attributes only the relations "earlier than" and "later than" to pairs of events. Another is the debate over endurantism and perdurantism. Endurantism asserts that objects exist wholly at each moment of time, whereas perdurantism asserts that it is only a temporal part of each object that exists at each moment. The debates over these positions persist because nothing in experience can decide them.

3.3 Thought experiments

A third example can be reviewed briefly. Brown (1991) has argued that we can learn non-empirically of the laws of nature through a thought experiment. The vehicle is Platonic vision. Through certain sorts of thought experiments, we are able to intuit, that is, see in our minds, the laws of nature as they reside in the Platonic world of forms. As a counterpoint, I have defended an empiricist account of thought experiments. They are merely picturesque arguments that can

---

In any case, the Minkowski spacetime of special relativity is exceptional. The sorts of spacetimes that are closer to that of our actual universe are ones like the Schwarzschild spacetime of a sun-like object and the Friedman-Lemaitre-Robertson-Walker spacetimes of cosmology. They all admit a unique relation of simultaneity.

inform us of the world only because they rely on empirically founded premises. In (2004), I challenged the Platonic view with pairs of thought experiments, such that each member of the pair arrives at a conclusion that negates the other. My empiricist account can determine which, if either, succeeds, most commonly by identifying fallacious reasoning. The Platonic account has no intrinsic means to separate the successful from the unsuccessful thought experiments. We are to assume that Platonic vision in at least one case was clouded, but how we are to separate clouded from clear vision is left open.

The difference between presentism and eternalism has no empirical expression in our best physical theories of space and time. Correspondingly, whether the Platonic account of thought experiments is true or not makes no difference to the practical analysis of thought experiments. All that can usefully be said about the epistemology of time and of thought experiments is already captured in their empiricist analyses.

4. Scientific Metaphysics

4.1 Scientific Metaphysics is not Traditional Metaphysics

That metaphysics does or even should proceed non-empirically is not a view universally held. The problem raised above for non-empirical metaphysics provides good reasons for the hesitation. The response is work in a field called variously “scientific metaphysics” or “empirical metaphysics” or “naturalized metaphysics” or “experimental metaphysics.” It is distinguished by its use of the sort of empirical evidence common in the sciences and by efforts to mimic the methods routinely used in science.

Kincaid (2013, p. 3), in the introduction to a collection of papers on scientific metaphysics, gave this characterization:

> … 1) 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 2) the belief that scientific results and scientific methods can be successfully applied to some problems that could be called metaphysical.

In so far as these scientific results and scientific methods are committed to proceeding empirically from experience, they are most welcome to an empiricist like me. Empirical methods are the best and only way to secure contingent facts at the foundations of science.

What is concerning is the continuation of the use of the term "metaphysics." There is a fundamental tension between the worthy project of proceeding empirically in the manner of empirical science and the traditional conception of metaphysics. For metaphysics is traditionally defined in a manner that contradicts an empirically founded conception of scientific metaphysics. The methods of metaphysics are traditionally understood as non-empirical. Its analyses are prior to that of any empirical science and its results are at a level of generality that transcends any one science, but extends to all of them.

This ambitious conception is stated routinely when the scope of metaphysics is defined. Here are some instances of it. Lowe (2002, p.2) allowed that metaphysics deals with topics also dealt with in empirical science, such as the nature of space, time and causation. However, he insisted that its methods are not empirical: (p. 2)

> But metaphysics is not at heart an empirical science—it does not typically appeal to experimental or observational data in support of its claims.

Lowe stressed that the subject matter of metaphysics transcends any specific science: (pp. 2-3, his emphasis)

> … its central concern is with *the fundamental structure of reality as a whole*. No special science—not even physics—can have that concern, because the subject-matter of every special science is identified more narrowly than this: for instance, biology is the science of living things, psychology is the science of mental states, and physics—as I have already indicated—is the science of those states and processes (energetic states and dynamic processes, for example) which are apparently common to all things existing in space and time.

Earlier in the text, Lowe had already described how metaphysics precedes empirical investigation since it provides a framework within which empirical investigations can proceed: (p. v, his emphasis)

> According to this conception, metaphysics goes deeper than any merely empirical science, even physics, because it provides the very framework within which such sciences are conceived and related to one another. A core text in metaphysics written from this point of view must aim, first and foremost, to elucidate certain universally applicable concepts–for example, those of *identity, necessity, causation, space,* and *time*—and then go on to examine some important doctrines which

> involve these concepts, such as the thesis that truths of identity are necessary and the claim that temporally backward causation is impossible.

Le Poidevin (2009, p.xx) gave a characterization of metaphysics in terms of its subject matter:

> [It is] concerned with what it is to be or be real, with what things there are, with the way that they are, and with the connection between the way things are and what things there are. And all this is pursued at a higher level of abstraction than typifies any of the special sciences like physics, geology or chemistry.

Williamson (2013, p. 3) repeats the conception of metaphysics as antecedent to those of empirical science:

> What is metaphysically necessary is what could not have been otherwise, what would have obtained whatever had obtained; not even the laws of physics can be assumed to be metaphysically necessary in this sense.

4.2 Individual Sciences as Metaphysics

These last characterizations of metaphysics characterize the content of metaphysics as antecedent to empirical science and more general than any particular science. In the later part of the twentieth century, an alternative conception became common in philosophy of science. The mere study of the foundations of a particular science came to be called "metaphysics." This alternative conception directly contradicts those just reported in the last subsection. It is the most important manifestation of what is now called scientific metaphysics.

How this variant usage arose in philosophy of science requires more detailed investigations than have been carried out here. However, we can track major moments in the transition. To begin, the logical positivists' disparaging of metaphysics in the earlier twentieth century left an enduring cloud over metaphysics in philosophy of science. Metaphysics continued to be discussed in many places, but commonly with hesitant qualifications. Consider its treatment in Neurath's *International Encyclopedia of Unified Sciences*. Morris (1938-1955) in his contribution "Scientific Empiricism" to Volume 1, Part 1, dismissed metaphysics with: (p. 73)

> All knowledge forms in principle one unified whole, and there exists no system of knowledge (such as metaphysics, aesthetics, ethics, religion) alongside of or superior to unified sciences.

Metaphysics then still had a small enough presence in philosophy of science for Ernest Nagel's 1961 *Structure of Science* to include the word just once in an incidental role (p. 187) is all its 618 pages.

The concept of metaphysics retained a comparably dubious status in philosophy of physics. Philipp Frank's (1938-1955) article, "Foundations of Physics," in Volume 1, Part 2, of the Neurath encyclopedia identified dubious understandings of quantum mechanics as "metaphysical interpretations." He wrote, under that heading: (p. 481)

> 49 Metaphysical Interpretations of Wave Mechanics
>
> Quite a few authors have maintained that by the new mechanics an "irrational element is introduced into physics" or that "physics is now supporting an idealistic word [*sic,* "world"?] picture" or that "physics is now in agreement with the doctrine of free will" or even that "by the wave mechanics for the first time in the history of human thought the conflict between religion and science has been settled." If we pursue precisely the analysis of the logical structure of the new mechanics, we will understand that there is no foundation for all these philosophical interpretations of the new mechanics.

The fortunes of the term "metaphysics" soon reversed. In his 1978 "Metaphysical Problems in the Foundations of Quantum Mechanics," a leading philosopher of physics, Abner Shimony, identified key novelties in the foundations of quantum mechanics as "metaphysical."[58] A decade later, Shimony (1989) amplified this casual use of the term "metaphysics" into something more programmatic. The principal application remained the foundations of quantum mechanics. However, Shimony now expanded the applications to other sciences. They included (pp. 63-64) general relativity for its demonstration that spacetime has substance-like properties through its dynamical interactions with matter; and molecular biology for its elimination of the need for vitalistic conceptions of life. Shimony summarized his view as: (1989, p.64, his emphasis)

> It is legitimate, in view of all these rich results, to speak of the enterprise of *experimental metaphysics*.

---

[58] They were that the superposition principle entails that quantum systems may lack definite values for certain properties; and that the state of entangled systems may have no complete, local specification.

The idea that the foundational results of quantum mechanics simply are results in metaphysics took hold, as did the idea that foundational work in quantum theory pursues "experimental metaphysics." The term was the title of a 1997 collection of papers[59] in quantum mechanics, written in Shimony's honor. Subsequently, it is became routine for foundational work in quantum mechanics to be identified as "metaphysical."

It is interesting to speculate on the reasons for this reversal. My suspicion is that it stems from the enduring antipathy among physicists in the 1970s towards the sort of foundational work Shimony was pursuing in quantum theory. Those pursing this sort of work found a more congenial intellectual reception among philosophers. That relieved them of the need to establish the value of their work among physicists. More metaphysically minded philosophers would, for their part, welcome a rehabilitation of metaphysics from novel work in the foundations of physics. Or so I surmise.

The idea that foundational work in a specific science can count as metaphysics has become commonplace in the decades after Shimony's work. Here are two illustrations among many. Jackson (2021) has identified principal foundational results of the special theory of relativity (STR) as metaphysics: (p. 398)

> STR is a thesis in metaphysics. This is because it is a thesis about the nature of the world we occupy. Among other things, STR holds that there is no absolute simultaneity at a distance. A victory salute in a soccer stadium in Brazil cannot be simultaneous with one in Italy in any absolute sense. The salutes can be simultaneous relative to some given inertial frame, but when they are, they will fail to be simultaneous relative to some other inertial frame.

In a short *Companion* entry, Leslie (2009) equated cosmology and metaphysics: (p. 195, his emphasis)

> Cosmology. In a wide sense, *cosmology equals metaphysics*: reality studied philosophically. This entry will instead discuss the study of the cosmos at large scales, using data from astronomy and physics.

[59] Cohen, Horne and Stachel (1997)

The idea of "experimental metaphysics" has spread beyond physics to work in so-called experimental philosophy, which is a study that draws heavily on empirical work in psychology. A recent collection of papers (Rose, 2017) has the title, *Experimental Metaphysics,* and specifies its goal as: (p. 1)

> The aim of this volume is to take up the question of how experimental philosophy can contribute to metaphysics.

4.3 Empirical and non-Empirical Metaphysics: a Misleading Unity

That empirically based, foundational work in specific sciences is called "metaphysics" may seem benign. It does contradict the traditional conception. Should we just conceive of it as an expansion of the traditional conception? The worst harm, we may suspect, is just that the meaning of metaphysics is broadened and thus diluted into vagueness. Everything is metaphysics, as long as it is not too specific in its content.

From a small-e empiricist perspective, matters are otherwise. It is dangerous to expand the scope of metaphysics to encompass the empirically derived foundations of individual sciences. For it combines activities of quite different legitimacy. One, non-empirical metaphysics, is so methodologically flawed as not to be a reliable way to learn contingent facts of the world. The other, empirical investigations, is, according to small-e empiricism, the only reliable way to learn contingent facts of the world.

Combining them risks the misunderstanding that both non-empirical and empirical methods are capable of achieving the same end, a unified body of foundational results called metaphysics. Under this misunderstanding, we would be free to use non-empirical or empirical methods to the same end. The choice is a matter of expedience and even personal preference. The combination masks the ineffectiveness of non-empirical methods and detracts from the effectiveness of empirical methods. It encourages the idea that an artful enough application of non-empirical methods might replace the vastly more onerous empirical methods. The proximity of well-supported results in empirically well-founded science lends an unearned legitimacy to the otherwise unfounded results claimed in non-empirical metaphysics. That same proximity detracts from the legitimacy of empirically well-founded results in a science, since they are grouped with results that do not enjoy the same legitimacy.

An analogy might capture the concern of small-e empiricism. Both astrology and astronomy deal with celestial phenomena. It would surely be most unwelcome to empirical astronomers for the two fields to be merged into a single field, called either astrology or astronomy. For astrology is an enterprise that has failed in its methods and results, where empirical astronomy has enjoyed immense success in both. They do not belong together.

Are the risks sketched realistic? We have an example is the last subsection. I quoted Jackson's (2021, p.398) assertion that "STR is a thesis in metaphysics." The passage quoted reported Einstein's empirically well-founded result of the relativity of simultaneity. The passage quoted is followed immediately by a claim routinely made in writing on the metaphysics of time: (pp. 398-99)

> A plausible thesis about the concept of existence implies that presentism and STR cannot be true together. … There is no sense to something existing relative to such and such, but not existing relative to so and so. … We have, therefore, a plausible example of a conceptual claim (about existence) telling us that one thesis in metaphysics (STR) implies the falsity of another thesis in metaphysics (presentism).

I have already argued in the last section that a decision between presentism and eternalism depends on an arbitrary attachment of the predicate "real," which is associated with no empirical manifestation. The illusion now presented is that a non-empirical assertion, the failure of presentism, is somehow supported by the empirically well-founded results of another branch of metaphysics, the empirically well-founded results of special theory of relativity.

5 The Problem of Fatal Abstraction[60]

5.1 The Assumption

The problem of fatal abstraction arises from an assumption that is fundamental to traditional metaphysics. Each of the concepts investigated in this metaphysics will have many manifestations. The assumption is that we can abstract away all the specifics that distinguish the different manifestations and be left with something non-trivial of universal applicability. That commonality provides the substance of the metaphysics. We saw an example of this assumption

[60] I failed for once to resist momentary cultural references in titles. "Fatal Attraction" is a 1987 psychological thriller movie starring Michael Douglas and Glenn Close. The problem was first sketched in Norton (2022, Section 12.4).

in the characterizations of traditional metaphysics sketched in the last section. Le Poidevin (2009, p.xx) remarked that metaphysics is "… pursued at a higher level of abstraction than typifies any of the special sciences…" Here is another formulation, given by Craig: (2005, p. 656)

> Metaphysics is a broad area of philosophy marked out by two types of inquiry. The first aims to be the most general investigation possible into the nature of reality: are there principles applying to everything that is real, to all that is? – if we abstract from the particular nature of existing things that which distinguishes them from each other, what can we know about them merely in virtue of the fact that they exist?

Loux (2006, p.x) endorses a similar scope for metaphysics in a tradition that he traces back to Aristotle:

> On this conception, metaphysics is the most general of all the disciplines; its aim is to identify the nature and structure of all that there is.

5.2 The Problem

The difficulty with this assumption is that it cannot be legislated in advance. Whether an abstraction produces something non-trivial has to be determined on a case-by-case basis. I will argue in the following sections that the assumption fails in two important cases, the metaphysics of possibility and the metaphysics of causation.

There are likely many other cases that cannot be recounted more fully here. The metaphysical fascination for laws is one. There are many regularities of great scope found empirically by the sciences. The goal of a metaphysics of laws is to abstract away the particulars of the regularities that locate them within specific sciences and thereby to recover a universally applicable account of laws. It was already noted in *small-e empiricism,* Chapter 12, Section 8.4, that such efforts have failed, for we have been unable to find exceptionless, well-supported regularities. Efforts to embellish an account of laws by introducing further metaphysics only deepen the difficulties. We might seek an account of laws in terms of natural kinds or universals. For Armstrong (1983, Ch.6), for example, laws are relations of necessitation between universals. We now have the added burden of defending these further metaphysical conceptions against the objections laid out in this chapter.

Another likely case of fatal abstraction is a newer topic in metaphysics, "grounding." It offers a general characterization of how one thing can depend on another. Although I will not pursue it further here, it does appear to depend on a fatal abstraction. Wilson (2014) has argued that the general conception is too weak to be illuminating, but that particular cases of grounding are better understood by looking at the particular way that the dependencies have arisen.[61]

These abstractions are fatal for the metaphysics, for once they are carried out, nothing of substance remains. The project of traditional metaphysics fails. There may be a deeper reason behind these failures. The assumption is that behind the intricate complications we find in many aspects of the world, there is a unifying simplicity of a type that can be captured by the generalizations of metaphysics. It would be a great kindness if the world was organized around the sorts of simplicities that we find comprehensible and form the substance of traditional metaphysics. The examples of the failures of the assumption are striking enough to preclude optimism for this kindness. To stretch the anthropomorphism, it appears that the world feels no obligation to organize itself such that it conforms with the simple generalities of metaphysics.

The problem of fatal abstraction arises for traditional metaphysics quite generally. It is, however, most troublesome for empirical metaphysics that seeks results of broad generality. For the problem challenges the methodology that such an empirical metaphysics would employ. That is, it would first find instances of the metaphysical conception and then seek to discern the generality common to all the instances. If the problem of fatal abstraction applies, there is no common generality to be discerned.

6. Metaphysics of Possibility

6.1 The Empiricist Conception of Possibility

The problem of fatal abstraction is acute for the metaphysics of possibility if we seek an empirical foundation for it. In Norton (2022), I explored the prospects for an empirically based account of possibility. The scope of such an account is quite limited. It must base its judgments of possibility and necessity on the evidence of experience. Thus, all judgements of possibility and necessity are relational. A proposition can only be judged possible or necessary with respect

[61] See Bliss and Trogden (2021, Section 3.2) for discussion of this and further concerns about grounding.

to some specific body of evidence. The formulation developed in Norton (2022, p. 129) is given in the paper's abstract as:

> What is possible, according to the empiricist conception, is what our evidence positively allows; and what is necessary is what it compels. These notions, along with logical possibility, are the only defensible notions of possibility and necessity.

The paper teases out, at some length, the consequences of this empiricist analysis. They conflict, to a great extent, with familiar results in the literature on the metaphysics of possibility that purport to have universal applicability. Here are a few of these conflicts.

A repeated source of the conflicts is that strong inductive support, no matter how strong, is fallible. That means that no necessity, in this account, can be assured absolutely as a truth. A simple example is the law of the conservation of energy. It is one for which we have the most massive support in our science and is very strongly supported on all our evidence. Thus, if any proposition in science meets the empiricist criterion of a necessity, then it is the law of the conservation of energy. Precisely because even the strongest evidence leaves open the chance of failure, there have been, from time to time, cogent proposals for its failure. They do have some evidence in their support, even if very weak. But that is all that is needed for a judgment of possibility. A simple example arises in the steady state cosmology of Bondi and Gold from the mid twentieth century. They proposed that an expanding universe can maintain a constant average matter density as long as a small amount of mass-energy was popping randomly and repeatedly into existence throughout the cosmos.

It follows that the familiar possible world semantics fails. According to it, a necessity is a truth in all possible worlds. The necessity of the conservation of energy is not a truth in all possible worlds, for it fails in worlds conforming with Bondi and Gold's cosmology. Repairs to the semantics are possible, but they are not simple and the difficulties can only be hinted at here. The semantics must now incorporate the body of evidence used for judgments of possibility and necessity. If we translate the strong inductive support of a proposition into truth in most possible worlds, we need to develop an account of this "most."

6.2 Modal Logic Misapplied as a Metaphysics of Possibility

It seems to be fairly widely accepted that the modal logic S5 captures the metaphysical notions of possibility and necessity.[62] For an empiricist, the idea is quite implausible. The notions of possibility and necessity are rich and complicated. The axioms of S5 are few and simple, yet they are supposed to apply across all domains in which possibility and necessity arise. The axioms are readily seen to fail for the empiricist conception.

To begin, consider the axiom "$\Box A \rightarrow A$", that is, informally, for a proposition A, if A is necessary, then it is true. Under the empiricist conception, A is necessary if it has strong inductive support, but, since such support is fallible, its truth is not guaranteed.

The duality of the operators of possibility $\Diamond$ and necessity $\Box$ are expressed as "$\Box A \rightarrow \sim\Diamond\sim A$", if A is necessary then it is not possible that A is false. This duality is violated by the example of the conservation of energy. If, $\Box A$, energy is necessarily conserved, then it is possible that energy conservation is false, $\Diamond\sim A$.

Finally, we have "$\Box A \rightarrow \Box\Box A$" that if a proposition is necessary, then it is necessary that it is necessary. Iterating, we have that, if the proposition is necessary, then it is necessary that it is necessary …and so on finitely often, "$\Box A \rightarrow \Box\Box\Box ... \Box A$". If necessity is strong inductive support, then this iteration fails. That is easiest to see if we represent strong inductive support by high probability, such as 0.99. Repeated iterations correspond to repeated multiplications 0.99 x 0.99x 0.99 x … x 0.99, which will eventually yield an arbitrarily small probability.

It may be helpful to consider an instantiation of this iteration. A signal is to be sent from a source location along a noisy telegraph wire. The signal is to be picked up at a sequence of repeater stations. Each reports to the next station that a signal was received; and each report is correct with probability 0.99 and incorrect with probability 0.01, due to its mistaking of noise for a signal. A positive report from the first repeater station amounts to a probability of 0.99 that the signal was sent. A positive report from the second repeater station amount to a probability of 0.99 x 0.99 that the signal was sent; and so on.

The general conclusion from this empiricist analysis is that possibility and necessity are intrinsically inductive notions. One might harbor a hope that something in inductive logic itself

[62] See for example Williamson (2013, 2016).

might allow us to see how the notions are governed by simple regularities. They may not be the axioms of S5 but they may be something of comparable simplicity. The material theory of induction dashes such hopes. Its core assertion is that there are no universally applicable rules of inductive inference. Thus, there can be no simple rules governing strong inductive support that can be sustained universally, that is, across all the manifestations of judgements of necessity.

In short, under this empiricist analysis, the material theory of induction entails that the metaphysics of possibility will suffer the problem of fatal abstraction.

7. The Metaphysics of Causation

7.1 The Expanding Variety of Causes

The metaphysics of causation also fails because of the problem of fatal abstraction. It is an established empirical fact that things in the world connect together in all sorts of interesting ways. Forces accelerated bodies. Voltages drive currents. Chemical potentials advance chemical reactions. Selective pressures lead to the evolution of biological species. Aspirin reduces fever and inhibits blood clotting. And so on, in innumerably more cases across all sciences. The goal of a metaphysics of causation is to abstract away all the particulars from this expanding collection of varied connections to recover a simple, unifying narrative. It is to encompass this great variety and thereby capture the fundamentally causal nature of reality. In a series of papers (Norton, 2003, 2007, 2009, 2016) and most recently Norton (2024a) I have argued that causal metaphysics has failed to find a unifying narrative that has any definite factual content.

To preclude confusion, this conclusion derives from my empiricism, but it is *not* a Humean skepticism about causation. For Hume's skepticism is based on Hume's inductive skepticism. Where Hume was an inductive pessimist, small-e empiricism is based on an inductive optimism.

Small-e empiricism allows that we can learn contingent truths inductively from experience that extend well beyond experience. In so doing, we find a great variety of factual dependencies in the world. Their range is ever growing. Opening each new domain of inquiry brings dependencies of a new character. At any stage of our learning, we may fit a simple, overarching narrative to those dependencies such as is sought by causal metaphysics. No such narrative survives when new dependencies are learned. This failure is no basis for epistemic pessimism. We do possess an ever-growing and rich repertoire of dependencies in the world. As we shall see in the brief survey that follows in the next subsection, what remains elusive is a unifying, simple narrative. If the long history of failed efforts to find one teaches us something, it is that the variety of dependencies in the world is inexhaustible and not subsumable under such a simple narrative. That is, attempts to abstract away the particulars of the many cases in this variety is fatal. It leaves nothing non-trivial behind.

This failure is most troubling for a non-empirical metaphysics of causation. Each failed attempt at a unifying metaphysics of causation is a failure of the non-empirical mode used to arrive at it. We saw in Chapter 2 that these modes are unable to explain their failure. All we learn from the repeated failures is that these modes are unproductive. Matters are only slightly better for an empirical metaphysics which tries to discern its content from experience. We do learn that things in the world connect in repeatable and reliable ways. These are the regularities that form the content of science. What an empirically informed metaphysics should learn from the endless variety of the connections found through experience, is that they do not conform with a simple, overarching metaphysics of causation.

7.2 No Factual Principle of Causation

The goal of a metaphysics of causation is to identify how causation governs processes in the world at a foundational level. If that identification is factual, it will be expressed in a principle of causality that circumscribes universally the sorts of processes that can occur in the world. As my papers cited above recount in more detail, efforts to find such a factual principle have failed repeatedly.

We may start a brief review in the seventeenth century of the continuing vacillations over how causation factually restricts processes in the world. A unifying theme in the new thinking of the seventeenth century was a repudiation of Aristotle's four component conception of causation: efficient, material, final and formal causes. Three of Aristotle's causes were to be discarded and only one, efficient cause, was viable. Even the great Isaac Newton insisted upon a severe restriction on which causal processes were admissible. He dismissed the idea of action at a distance as a novice scientist's mistake: (Newton, 1761, pp. 25-26.)

> [T]hat one body may act upon another at a distance through a vacuum, without the mediation of anything else, by and through which their action and force may be conveyed from one to another, is to me so great an absurdity, that I believe no man,

> who has in philosophical matters a competent faculty of thinking, can ever fall into it.

The irony was that by the nineteenth century, the gravitational forces of Newton's physics came to be begrudgingly accepted as just this action at a distance. An authoritative Teubner Encyclopedia article summarized the best empirical results on gravitation at the end of the nineteenth century. They favored gravitation as action at a distance. Empirical investigations, Zenneck (1901) reported, had failed to find any empirical manifestation of an intervening medium that could carry gravitational action (§17, pp. 42-43) and had also failed to place a finite limit on the speed of propagation of gravitation (§20, pp. 44-46). A priori speculation on the possibility of action at a distance was overridden and eclipsed by empirical investigations.

In the nineteenth century a new, purified conception of causation prevailed. It was purged of all elements except one. The principle of causality asserted only the determinism of all processes. The present state of the world fixes its future state. When quantum mechanics was discovered in the early twentieth century, this nineteenth century dogma of determinism was found empirically to fail. This failure was characterized in apocalyptic terms as the failure of causality.

Bridgman, writing later in 1959, recalled his resistance to the chanciness in newly emerged quantum theory, when he wrote his 1927 *Logic of Modern Physics*. It derived from a prior commitment to causality as requiring determinism that controlled his thinking as well as that of others like Einstein. Bridgman wrote: (1959, p. 519)

> This I said on page 117: "We cannot permanently be satisfied with a picture of radioactive phenomena which represents radioactive disintegration as a matter of chance." I suppose I said this because I thought that the requirement for determinism and causality, perhaps in the guise of a principle of sufficient reason, was an inescapable requirement of any thought that could be called rational. In this I was in good company, for I suspect that such an opinion is back of the intransigent refusal of Einstein and others to accept "chance" as in any sense an ultimate component of the world. The reason I have changed is that I have seen it demonstrated that it is not an inescapable requirement of our thought; I have seen others learn, and I have myself learned, how to think without this requirement.

Causal metaphysics soon recovered from this loss of determinism. The nineteenth century dogma of determinism was replaced by the new twentieth century dogma of probabilism. In the new conception, causes do not determine their effects. Rather, they raise the chances of the effect. Here chances are identified with probabilities. Mellor has asserted his complete confidence in the dogma: (1995, p. 21, his emphasis)

> The first point I must make about chances is that they are *probabilities*. My other claims about chance may be contentious, but not this one: everyone measures chances by numbers which satisfy the standard calculus of probabilities.

This new dogma fares no better than the nineteenth century dogma of determinism. Of course, there are many cases of processes that are governed by probabilities and that do conform with the conception of causes as probability raisers. The problem is that this conception does not apply universally to all processes. Chances and indefinitenesses in general are not always well represented by probabilities. For an entry into the literature that recounts this failure of probabilism, see Norton (2021, Ch. 10-16; 2024a).

Efforts to find a factual principle of causality persist, but without success. There is continuing interest in "Curie's principle." It derives from Pierre Curie's nineteenth century investigations in crystallography and asserts that a symmetry of a cause must be replicated as a symmetry of its effect. Norton (2016) argued that the principle has no independent factual content. In applying it, we can render it as true or false according to how we choose to map its terms onto the cases at hand. In another context, Norton (2009) argues that this same vagueness about exactly what counts as a cause and an effect compromises a recent attempt to formulate a principle of causality as the requirement that an effect cannot precede its cause in time. A recent literature has proposed various ways that causation places factual restrictions on the prospects for downward causation between different levels in scientific theories. Norton (2024a, Section 8-9) recounts their failure.

7.3 The Pragmatic Value of Causal Labeling

We may have hoped that identifying a process as causal would assure us that it conforms with some foundational, factual restriction on things in the world. The above discussion shows that such a factual restriction has proved elusive. We can still find a pragmatic value in causal labelling, without the labels imposing factual restrictions on processes.

Causal labels can assist our psychological understanding in the sense of allowing us to form more serviceable mental images. The voltage difference and amperage of an electric conductor are linearly related. We understand the relationship better if we are told that the voltage difference *causes* the current flow that is measured by the amperage. We can picture it as similar to how water pressure pushes water to flow in a pipe and even just how our muscular action sets ordinary objects into motion. Norton (2024a, Section 4) recounts a more substantial example in Einstein's "*A* and *B* coefficients paper" of 1917 in which the foundations of LASER technology are laid down. One of the processes is described causally as "stimulated emission," even though it is fully specified as a relationship between the ambient electromagnetic field and the probability of an electric charge state transition.

Causal labels can also enhance communication. Woodward's (2003) interventionist account of causation identifies a process as causal if an intervention on one variable is associated with changes in another variable. There is no factual principle of causality presumed in this definition. It is an empirical matter as to whether there are any processes that conform with the many clauses of Woodward's full definition. However, being told that a process is identified as causal in the interventionist sense, tells us a lot in a few words. For then we know that we can use the process to manipulate what is identified as the effect by intervening on the cause.

8. Causal Metaphysics as Failed a priori Science

This last use of causal language is benign as long as it is restricted to the introduction of pragmatically useful definitions, without any prior requirement that the world must instantiate them. A neighboring conception is not benign. According to it, causal metaphysics provides us causal concepts that are to be instantiated by subsequent empirical investigations. If that prior causal metaphysics is used to restrict factually what can happen prior to empirical investigations, then it is a form of ill-fated, a priori science.

Such a priori science arises in Mellor's metaphysics of causation. He remarked: (1995, p. 5, his emphasis)

> So while I will accommodate the relevant results of modern physics, I will not for example leave it to quantum physics to tell me whether causation can act immediately across spacelike intervals. On the contrary: only when our metaphysics as told us what causation is can we see if physics *could* reveal unmediated action at a distance (it couldn't).

Mellor continued, in an adventurous spirit, to insist on how spacetime *must* be on the strength of his causal metaphysics. He concluded (pp. 234-35) that there can be no causal loops and thus there must be a single linear dimension in spacetime whose direction is given by the direction of causation. He concluded: (his emphasis)

> So if that dimension *is* time, this fact shows why any possible spacetime, however many spatial dimensions it has, can have only one time dimension. More: it shows that any possible world with laws, and hence chances and causability, must have at least one dimension, that of time, even if it has no space.

It then follows that: (p. 235, his emphasis)

> … causes *must* precede their effects.

A further restriction on physics is derived from his strict causal metaphysics in Mellor's (1998) later monograph on time. He concluded in (1998, Ch.12) the impossibility of time travel, backward in time, such as would require a closed causal loop. This contradicts a modern literature that finds such time travel to be possible in an array of spacetimes arising in Einstein's general theory of relativity.[63] Mellor dismissed one of the first, Kurt Gödel's 1949 spacetime, and, presumably, the huge literature that followed with: (1998, p. 127)

> … all it shows is that general relativity does not rule them [timelike loops] out, not that nothing does. (Gödel's argument illustrates the common and pernicious fallacy that nothing can show what physics cannot show, so that anything which is physically possible, i.e. consistent with physics, must be possible in fact.)

I do not see a pernicious fallacy at work. Rather, I see physics telling us how such loops would fit with present physics and giving us prescriptions for what we would need to do to detect them empirically. In the Gödel case, we would need to find the bodies of the universe in rotation *en masse* about us, which is not the case. If there is a tendentious idea present, it is Mellor's folly in believing that a priori reflections on what causation must be can pre-empt what empirical investigations can yield.

[63] As an entertaining exercise, Norton (2025) describes a spacetime that is everywhere locally like that of special relativity excepting for a singularity. In it, a freely moving spaceship can end up encountering its past self and also find itself aging in the opposite time sense.

Lowe (2002, p. 332) has a similar view on the possibility of backward causation and time travel.

> … empirical science alone cannot definitively determine the answer to a question concerning metaphysical possibility, such as the one with which we are now concerned, although I do not wish to deny that scientific theorizing and empirical data can throw useful light on a question like this. The key point is that empirical data can only be taken to be good evidence for the truth of a scientific theory if its hypotheses are already entitled to be viewed as expressing metaphysical possibilities: that they can be so viewed is not, thus, something that can be determined on merely empirical grounds.

Once again, non-empirical, a priori reflections are to pre-empt what can be found empirically.[64]

9. Empiricism as Metaphysics: the *tu quoque* objection

I expect that, at some point, someone will object that my critique is itself a part of metaphysics and thus it is self-refuting. This *tu quoque* ("you too") defense of metaphysics is common. For example, Lowe (2002, p. 2) considered the view that metaphysics extends beyond the reach of the physical sciences and responded:

> Of course, physicalist philosophers may want to challenge this view: but then they are engaging precisely in a metaphysical debate, not one which belongs to the province of physics itself.

A *tu quoque* defense is generally weak and, in this case, clearly fails. Whether the present critique is labeled as a part of metaphysics or not does not undo any of the specifics of the criticism. The field designated by metaphysics is so diffuse that the label of metaphysics carries little discerning power. My principal objection is specifically to results claimed in metaphysics that concern the nature of reality that are not sustained empirically. My analysis is based on the proposal of small-e empiricism, which is a doctrine itself derived empirically. That empirical derivation includes the account of inductive inferences recommended for its implementation, the material theory of induction.

[64] Lowe (2002, p. 334) "… as metaphysicians we are not absolutely committed to the truth of Einstein's theory."

Someone pursing this response might still want to press me on my allegiance to experience as the means of learning about nature. Is not that allegiance nothing more than a private metaphysics? Again, that the allegiance is labeled as metaphysics does not compromise its use against non-empirical metaphysics. The basis for my allegiance is not in some abstract, covert attachment to metaphysics. It has been laid out Chapter 2 above. There the argument is that relying on experience has been fertile, where other modes have not. If other modes are tried and come into conflict with experience, they will abandon any of their results that conflict with experience.

We may choose to doubt or abandon experience as a means of learning about the world. The prospects for such a choice are poor. To implement it, we may decide just to abandon the entire project of learning about the world. That is to choose needlessly a total ignorance. Or we may be more modest and adopt other methods in the place of experience. That policy is epistemically self-destructive. We choose to ignore the mode that has a proven history of correcting all other modes. It commits us willfully to another form of ignorance.

10. Conclusion: Metaphysics, Empiricists and Its Critics

A recent summary of unfavorable attitudes towards metaphysics was given by Manley (2009, p. 4). He distinguished "strong deflationists" (who find the disputes entirely misguided), "mild deflationists" (who believe the disputes can be resolved) and "reformers" (who find the disputes substantive but in need of novel approaches). My view aligns best with the strong deflationists. In the same volume on "metametaphysics," the view that comes closest to the one developed here is Bennett's (2009). It is restricted to debates over composition and material constitution. She concludes: (p.71, her emphasis)

> There do not appear to be any real grounds for choosing between the competing positions about either composition or constitution. We are not justified in believing either side. These are basically cases of *underdetermination of theory by evidence*.

To metaphysicians who struggle with what they judge to be the deepest and most profound problems concerning the nature of things, my empiricist critique may well appear hasty, superficial and even lazy, for not engaging more deeply into the details of the metaphysical analysis. I am a "physicalist philosopher" whose critique of metaphysics appears to Lowe as misguided.

In response, may I report how elaborations of non-empirical metaphysics appear to me as a small-e empiricist? I do appreciate the intellectual creativity of work in non-empirical metaphysics. It seeks to illuminate the deepest foundational problems without the guidance of experience; and it must do it in a way that comports with the existing, elaborate canon of writing in non-empirical metaphysics. With these constraints, it is a significant intellectual achievement to add to this canon. However, the achievement is just that of conforming with an existing canon. It is not a work of discovery of contingent facts because it lacks the essential element. The analysis proceeds without empirical foundation.

References

Armstrong, D. M. (1983). *What is a Law of Nature?* Cambridge: Cambridge University Press.

Bennett, Karen (2009) "Composition, Colocation, and Metaontology," Ch.2 in Chalmers et al. (2009).

Bliss, Ricki and Kelly Trogdon (2021), "Metaphysical Grounding," *The Stanford Encyclopedia of Philosophy* (Summer 2024 Edition), Edward N. Zalta & Uri Nodelman, eds., https://plato.stanford.edu/archives/sum2024/entries/grounding/

Bridgman, Percy W. (1927) *The Logic of Modern Physics*. New York: The Macmillan Company.

Bridgman, Percy W. (1959) "P. W. Bridgman's 'The Logic of Modern Physics' after Thirty Years," *Daedalus* , **88**, pp. 518-526.

Brown, James R. (1991) *The Laboratory of the Mind: Thought Experiments in the Natural Science*. London: Routledge.

Bryant, Amanda (2020) "Keep the chickens cooped: the epistemic inadequacy of free range metaphysics," *Synthese*, **197**, pp. 1867–1887.

Callender, Craig (2011) "Introduction" in C. Callender, ed., *The Oxford Handbook of the Philosophy of Time*. Oxford: Oxford University Press.

Chalmers, David; Manley, David; and Wasserman, Ryan, eds. (2009) *Metametaphysics: New Essays on the Foundations of Ontology*. Oxford: Oxford University Press.

Cohen, Robert S.; Horne, Michael; and Stachel, John; (1997), *Experimental Metaphysics. Quantum Mechanical Studies for Abner Shimony. Volume One*. Dordrecht: Kluwer.

Craig, Edward (2005) "Metaphysics," pp. 656-59 in E. Craig, ed., *The Shorter Routledge Encyclopedia of Philosophy*. London: Routledge.

Emery, Nina (2023) *Naturalism Beyond the Limits of Science: How Scientific Methodology Can and Should Shape Philosophical Theorizing*. New York: Oxford University Press.

Fine, Kit (2011) "What is Metaphysics?" Ch. 1 in Tuomas E. Tahko, ed., *Contemporary Aristotelian Metaphysics*, Cambridge: Cambridge University Press.

Frank, Philipp (1938-1955) "Foundations of Physics" pp. 424-504 in O. Neurath, C. Morris, R. Carnap, eds., *International Encyclopedia of Unified Science*. Voume 1, Part 2. Chicago: University of Chicago Press.

Jackson, Frank (2021), "Conceptual analysis in metaphysics," Ch. 31 in Bliss, Ricki and Miller, J.T.M., eds, *The Routledge Handbook of Metametaphysics*. London: Routledge.

Kincaid, Harold (2013) "Introduction: Pursuing a Naturalist Metaphysics," Ch.1 in In Ross, Don; Ladyman, James; and Kincaid; Harold, eds., *Scientific Metaphysics*. Oxford: Oxford University Press.

Kölbel, Max (2023) "Varieties of conceptual analysis," *Analytic Philosophy*, **64**, pp.20–38.

Le Poidevin, Robin (2009) "General Introduction: What is Metaphysics?" Ch. 1 in R. Le Poindevin et al., ed., *The Routledge Companion to Metaphysics*. London: Routledge.

Leslie, John (2009) "Cosmology" pp. 195-200 in J. Kim, E. Sosa and G. S. Rosenkrantz, *A Companion to Metaphysics.* West Sussex: Wiley-Blackwell.

Lewis, David (1986). *On the Plurality of Worlds.* Oxford: Blackwells.

Loux, Michael J. (2006) *Metaphysics: A Contemporary Introduction.* 3rd ed. New York: Routledge.

Lowe, E. Jonathan (2002). *Survey of Metaphysics.* Oxford: Oxford University Press.

Mellor, David H. (1995) *The Facts of Causation*. London: Routledge.

Mellor, David H. (1998) *Real Time II*. London: Routledge.

Manley, David (2009) "Introduction: A Guided Tour of Metametaphysics," Ch. 1 in Chalmers et al. (2009).

Morris, Charles W. (1938-1955) "Scientific Empiricism," pp. 63-75 in O. Neurath, C. Morris, R. Carnap, eds., *International Encyclopedia of Unified Science*. Voume 1, Part 1. Chicago: University of Chicago Press.

Nagel, Ernest (1961) *The Structure of Science: Problems in the Logic of Scientific Explanation.* New York: Harcourt, Brace & World.

Newton, Isaac (1761) *Four Letters From Isaac Newton to Doctor Bentley*. Pall Mall: R. and J. Dodsley.

Ney, Alyssa (2014) *Metaphysics*: *An Introduction*. London: Routledge.

Norton, John D. (2003) "Causation as Folk Science," *Philosophers' Imprint*. Vol. 3, No. 4 reprinted in pp. 11-44, H. Price and R. Corry, *Causation, Physics and the Constitution of Reality*. Oxford: Oxford University Press.

Norton, John D. (2004) "Why Thought Experiments Do Not Transcend Empiricism," pp. 44-66 in Christopher Hitchcock (ed.) *Contemporary Debates in the Philosophy of Science.* Blackwell, 2004.

Norton, John D. (2007) "Do the Causal Principles of Modern Physics Contradict Causal Anti-Fundamentalism?" pp. 222-34 in *Thinking about Causes: From Greek Philosophy to Modern Physics.* eds. P. K. Machamer and G. Wolters, Pittsburgh: University of Pittsburgh Press.

Norton, John D. (2009) "Is There an Independent Principle of Causality in Physics?" *British Journal for the Philosophy of Science*, **60**, pp. 475-86.

Norton, John D. (2015) "The Burning Fuse Model of Unbecoming in Time," *Studies in History and Philosophy of Modern Physics*. **52**, pp. 103-105.

Norton, John D. (2016) "Curie's Truism," *Philosophy of Science*. **83**, pp. 1014-1026.

Norton, John D. (2021) *The Material Theory of Induction*. BSPS*Open*/University of Calgary Press.

Norton, John D. (2022) "How to Make Possibility Safe for Empiricists," pp. 129-159 in Yemima Ben-Menahem, ed., *Rethinking the Concept of Laws of Nature: Natural order in the Light of Contemporary Science.* Springer, 2022.

Norton, John D. (2024) *The Large-Scale Structure of Inductive Inference.* BSPS*Open*/University of Calgary Press.

Norton, John D. (2024a) "The Metaphysics of Causation: An Empiricist Critique," pp. 58-94 in in Yafeng Shen, ed., *Alternative Approaches to Causation*. Oxford: Oxford University Press.

Norton, John D. (2025) "A Simple Minkowskian Time-Travel Spacetime," *American Journal of Physics*, **93**, pp. 241–249.

Rose, David, ed., (2017) *Experimental Metaphysics*. London: Bloomsbury Academic.

Shimony, Abner (1978) “Metaphysical Problems in the Foundations of Quantum Mechanics,” *International Philosophical Quarterly,* **18**, pp. 3-17. Reprinted as pp. 517-528 in R. Boyd, P. Gasper, and J.D. Trout (eds.), *The Philosophy of Science*. Cambridge, MA: MIT Press, 1999.

Shimony, Abner (1989) “Search for a Worldview which can Accommodate our Knowledge of Microphysics,” in *Philosophical Consequences of Quantum Theory: Reflections on Bell's Theorem*, eds. J. T. Cushing and E. McMullin. University of Notre Dame Press, 1989; reprinted as Ch.3 in Abner Shimony, *Search for a Naturalistic World View. Volume I. Scientific Method and Epistemology.* Cambridge: Cambridge University Press, 1993.

Stalnaker, Robert C. (1976) “Possible Worlds,” *Noûs*, **10**, pp. 65-75.

Stalnaker, Robert C. (1988) “Critical Notice on the Plurality of Worlds. By David Lewis.” *Mind*, **XCVII**, pp. 117–128.

van Inwagen, Peter; Sullivan, Meghan; and Bernstein, Sara, “Metaphysics,” *The Stanford Encyclopedia of Philosophy* (Winter 2025 Edition), Edward N. Zalta & Uri Nodelman (eds.), https://plato.stanford.edu/archives/win2025/entries/metaphysics/

Williamson, Timothy (2013) *Modal Logic as Metaphysics*. Oxford University Press.

Williamson, Timothy (2016) “Modal Science,” *Canadian Journal of Philosophy*, **46**, pp. 453-492,

Wilson, Jessica M. (2014) “No Work for a Theory of Grounding,” *Inquiry*, **57**, pp. 535-579.

Woodward, James F. (2003) *Making Things Happen. A Theory of Causal Explanation.* Oxford: Oxford University Press.

Zenneck, Jonathan (1901) “Gravitation” in *Encyklopädie der mathematischen Wissenschaften mit Einschluss ihrer Anwendungen*, Vol. 5, Part l (1903-21), pp. 25-67.

Afterword

Afterword

This volume is a project in history and philosophy of science. Its philosophical claims are not recovered by generalizing from a few artfully selected case studies in history of science. Rather its claims are supported by independent argumentation and then are illustrated, wherever possible, by case studies in the history of science. The goal is to display that the claims advanced here are not fanciful inventions, unrealized in actual science, but are instantiated in science.

A major goal of this volume is to correct a trend in the field of history and philosophy of science itself. Too much work in history and philosophy of science pursues what I regard as an unfounded skepticism. The leading idea is that scientists are naïve and credulous in believing what is in their textbooks. The mistake, we are to suppose, is revealed by a closer examination of the messiness of the history and by attention to the skeptical philosophical theses that are supposed to follow from the history. Science, in this view, is indissolubly conditioned by its historical context and scientists overestimate the reach of their methods.

Here I dissent. It is undeniable, I believe, that science has enjoyed immense success, unmatched by any other human enterprise. Skepticism is easy. All that is needed is to attribute a naïve and oversimplified methodology to the sciences and then to point out how it fails. The hard task is to figure out how the science really operates and how that operation has delivered the successes science enjoys. How did science achieve something quite singular in its successes: they transcend the historical circumstances of their creation and are portable across time and culture.

That success is not distributed uniformly across all the sciences. Some sciences are still in a rudimentary, unstable state and should be met with caution. If we focus on these sciences, we do find failures. The danger is to generalize from them to all sciences; to propose that all sciences fail in the same way; and, worse, to deny the very possibility of the sort of success scientists claim.

The failure of these generalizations is demonstrated by the profound successes of many sciences. These include the sciences I know best, physics and chemistry. I do *not* generalize their success to all sciences. That would make a mistake comparable to the mistaken generalization I see in present work in history and philosophy of science. Instead, my point is that science has available to it methods that can yield this level of success and that they have been employed successfully in at least some sciences.

History and philosophy of science can offer significant insights into science. Some will be negative in displaying failures. Others, such as offered in this volume, will justly celebrate the successes of science and will do so with a fuller understanding of the messiness and complications of real science.