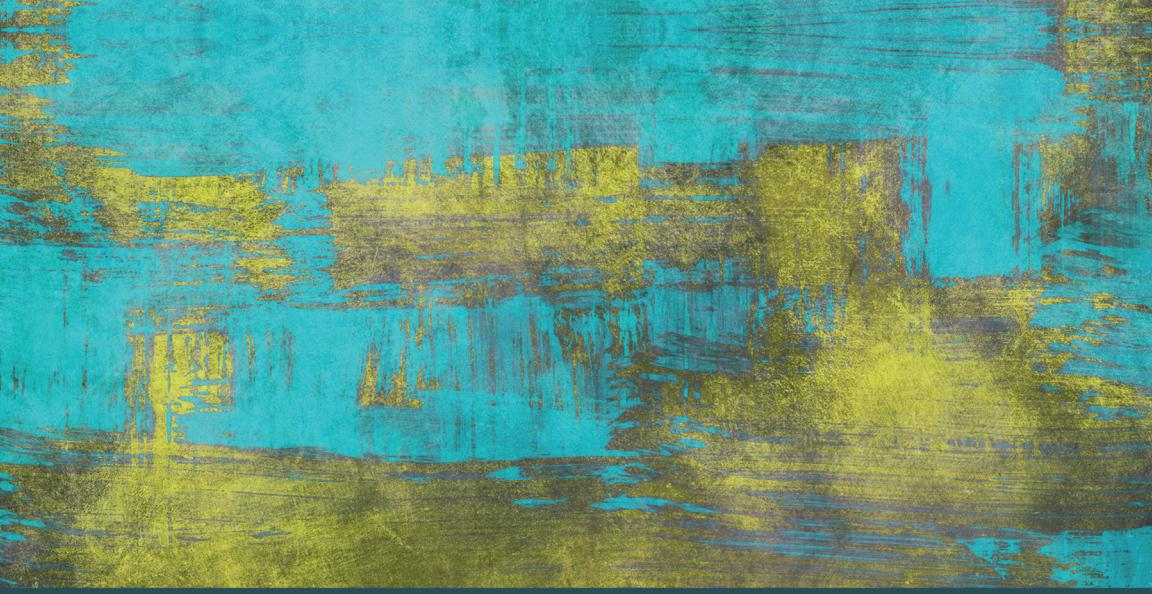

Routledge Studies in Contemporary Philosophy

NATURALISM AND ITS CHALLENGES

Edited by
Gary N. Kemp, Ali Hossein Khani,
Hossein Sheykh Rezaee, and Hassan Amiriara

11 Naturalizing Metaphysics
Epistemological Challenges

Anjan Chakravartty

A groundswell of recent philosophy has targeted the question of how best to think about the relationship between the empirical sciences, on the one hand, and metaphysical theorizing, on the other.[1] More specifically, much of this work has endeavored to articulate conceptions of metaphysical theorizing that are linked to or anchored by, and thereby in some way accredited by, what many regard as the superior epistemic credentials of the modern sciences. This project has been pursued under a number of suggestive headings, such as 'naturalized metaphysics', 'scientific metaphysics', 'inductive metaphysics', and 'the metaphysics of science', and even allowing for some inevitable variation in the precise use of these terms by different philosophers, a shared desideratum seems clear: to give an account of theorizing about aspects of the world on which the sciences themselves are not definitive, in ways that are nonetheless substantially connected to our best science, in hopes that this connection may extend at least some credence from the outputs of scientific inquiry to the relevant metaphysical theorizing.

It is fair, I think, to describe the philosophical impulse underlying this shared desideratum in terms of the aim of 'naturalizing', that is, the aim of demarcating between forms of metaphysical theorizing that satisfy this desideratum and those that do not, and of advocating for the former at the expense of the latter. The hope is that in furnishing such an account, we may see clearly why naturalized metaphysics is privileged in comparison to non-naturalized metaphysics so far as producing knowledge of the world is concerned. In this chapter, after an elaboration of this understanding of the project of naturalized metaphysics, I consider some pressing epistemological challenges to the very idea of naturalizing in this domain. Some of these challenges concern the identification of plausible criteria with which to identify some but not all metaphysical theorizing as naturalized, and others concern the more precise question of *how well* a domain of theorizing must satisfy such criteria in order to meet the requisite standard.

DOI: 10.4324/9781003430568-11

In Section 11.1, I sketch a recent history of philosophical disputes about how to interpret the outputs of scientific inquiry as yielding knowledge, or warranted belief, about the world, and in Section 11.2, I contend that these disputes are helpfully understood as fueled by a number of incompatible presuppositions held by their interlocutors, which I describe in terms of the adoption of different *epistemic stances*. In Sections 11.3 and 11.4, I consider how the possession of a stance shapes one's judgments concerning the epistemic credentials of metaphysical theorizing in different domains of inquiry. These judgments typically turn on assessments of the evidential weight, in any given case, of empirical inquiry and contributions to explanation, which I examine in turn. In conclusion, in Section 11.5, I consider whether the preceding proposal for understanding the epistemic dimensions of naturalism in this sphere is itself subject to a naturalizing impulse, namely, to fortify a philosophical account of the epistemology of naturalized metaphysics with a commitment to the ideal of naturalizing epistemology itself.

11.1 Metaphysics, Science, and Debates About Scientific Realism

To begin, and as a prelude to grappling with the issue of how, if at all, scientific inquiry and metaphysical inquiry are related to one another, let us pose the more focused question of whether and how it may be possible to do metaphysics in a naturalized sort of way. What should one do, specifically, in order to realize the aspirations of a naturalized metaphysician, and what parameters of assessment might one use to determine whether these aspirations have, in fact, been fulfilled? At a certain level of abstraction, everyone sharing this aspiration will agree that to do naturalized metaphysics is to do metaphysics in way that is appropriately tethered to science. This, however, does not yet begin to tackle the hard work of spelling out concretely what the relation of 'tethering' is, exactly. The devil is in the details that would make the metaphor apt. Perhaps, one might contend, the aim of naturalized metaphysics is to give an *interpretation* of the content of scientific theories and models that goes beyond scientific interpretations. Perhaps the aim is to give a deeper or more fulsome *elaboration* of scientific theories and models. Or perhaps it is to add something *complementary* to them. Such descriptions are suggestive but likewise abstract, and do not take us far.

As a first step toward saying something more concrete, let us be clearer about the epistemic aims of naturalized metaphysics. It is a common assumption that the aim of science is (no doubt among other things) to learn something about the world, even if there is ample disagreement regarding which aspects of our best science fit this description. What it is, precisely, that scientific investigations license in terms of warranted belief, or knowledge, has been widely contested historically by those advocating versions of

instrumentalism on one side, who tend to restrict scientific knowledge to the observable consequences of our best science, and advocates of versions of scientific realism on the other side, who in various ways dismiss that restriction. All parties to these debates nonetheless hold that science is an epistemic project: its goals include something like knowledge and, more specifically, knowledge of the actual world. When it comes to naturalized metaphysics, however, there is no such philosophical consensus.

Leaving aside those who are skeptical about any sort of metaphysics having any sort of value, there are some who take the value of at least some metaphysics to be *non*-epistemic. In the spirit of at least some forms of empiricism and pragmatism, for example, one might think that a reasonable conception of the aim of at least some metaphysics is to describe how the world *might be* or *could be* as opposed to (for all we know) how it is, in fact. On this understanding, we may well be in a position to establish certain epistemically warranted, scientific claims, to which we might then add naturalized metaphysics as a means of theorizing further, to explore a possibility space of finer-grained or more fundamental or general descriptions of whatever is claimed in well-warranted science. For instance, perhaps there are epistemically warranted facts about quantum phenomena as described by our best physics, which may then serve as a basis for further theorizing on the part of naturalized metaphysicians, who aim to furnish possible interpretations of quantum theory, which are viewed most appropriately as attempts to flesh out how the world might be or could in ways that go beyond what is scientifically warranted.

On *this* conception of naturalized metaphysics, given that there is nothing epistemic at stake, the bar for success is low. It may simply give us some satisfaction to explore realms of possibilities. On some occasions, it might even be a useful thing to do with an eye on the future, in the way imagined by W. V. O. Quine and Karl Popper and others, as a potentially fruitful heuristic exercise that is sometimes, eventually, co-opted by the sciences, such as when philosophical theorizing about natural kinds of entities constituting the world was followed in the fullness of time by natural philosophical and then scientific investigations of various categories of entities in the more specialized domains of physics, chemistry, biology, and the human and social sciences. In some cases what may begin as a purely formal exercise, such as John Bell's derivation of his famous inequalities (in 1964) concerning the possible outcomes of certain measurements of quantum phenomena, may be followed one day by empirical assessment, as when Alain Aspect was able to perform experiments to test Bell's theorem (in 1980–82), for which he later won the Noble Prize.

All of this said, if the value of naturalized metaphysics here and now resides only in imagining interesting possibilities—for we generally cannot say in the present when or whether such imaginings may one day prove

fruitful in terms of co-option by empirical science—it is not the case that the value of naturalized metaphysics here and now resides in its epistemic warrant. Even in the best-case scenario, warrant would only kick in later, if and when empirical investigation has entered the picture, which is not something that anyone can guarantee. This sort of merely possible, future-oriented, heuristic value is not, however, what advocates of naturalized metaphysics have in mind. The idea is rather the opposite: that the aim of naturalized metaphysics may be conceived in terms of it being an intrinsically epistemic exercise—that the point of doing it is to tell us something substantive about what the world is like, not merely possibly or one day, but actually and today. The reasonableness of this aim is intended to be a function of the way in which naturalized metaphysics is tethered to those aspects of science that are themselves properly regarded as having epistemic warrant presently.

This focus on the notion of warrant, or justification, suggests a way of proceeding in attempting to spell out more concretely what a plausible account of tethering metaphysics to science might be. Let us first consider the question of how we assess warrant in connection with the sciences, and then think about whether naturalized metaphysics, conceived as an inherently epistemic project, may reasonably be thought of as warranted on the same or a similar sort of basis. Here, though, a challenge arises immediately: there is in fact no consensus regarding how to understand warrant *even with respect to the sciences*. This claim must be carefully unpacked, for it is also the case that scientific communities (and philosophers) do recognize that some theories and models are better than others, epistemically speaking. The challenge arises *after* agreement about which theories and models comprise our best science, in identifying more specifically what their epistemic upshot is, exactly. It is one thing to use a model to make reliable predictions, to build reliable technologies, and so on, and quite another to believe that all aspects of the model are true descriptions of the world. As is well appreciated by scientists and philosophers alike, our best science is rife with idealizations, abstractions, and approximations which for all their utility nonetheless recommend caution regarding their descriptive content.

Those familiar with the philosophy of science will recognize in this a central debate in the field concerning what sorts of beliefs are warranted by our best science: a debate between varieties of scientific realists and antirealists. Roughly speaking, realism holds that our best theories yield truths (or something close by) regarding both observable and strictly unobservable aspects of a mind-independent world, and antirealism is any denial of this, often taking the form of a more restrictive conception of the scope of warranted belief, to scientific claims pertaining only to what is observable. I cannot do justice here to the vast terrain over which this

debate unfolds, but for present purposes, it will suffice simply to note two things. First, the debate as a whole, and the many ancillary debates it has grown to encompass—concerning the nature of truth and reference, conflicting ontologies of the world championed by different scientific realists, and conflicting conceptions of knowledge championed by different antirealists—are unresolved, as one might expect in any sufficiently complex philosophical arena. Second, disagreement here is apparent not only between philosophers but also between scientists, which is especially evident in historical case studies and in cutting-edge work in all eras, where explicit or unspoken disagreement about how best to assess the epistemic upshot of our best science is hardly uncommon.

In the following section, I provide a framework for understanding the nature of these disputes in terms of different (and likely irresolvable, or so I will later suggest) presuppositions regarding how to assess warrant in contexts of scientific knowledge. This will set the stage for a more detailed consideration of how these disagreements shape the beliefs that different agents take to be warranted and, concomitantly, an understanding of how they conceive the epistemic credentials of metaphysics in relation to the sciences.[2]

11.2 Relations Between Scientific Ontology and Epistemic Stances

Forms of scientific realism and antirealism have diversified significantly over the past few decades, which has witnessed a transition from speaking of the possible truth or approximate truth of scientific theories *simpliciter* to more refined conceptions of what if anything described *in* our best theories and models are plausible subjects of realism—certain entities, structures, properties, and so on, whether observable or unobservable. Let me use the term 'scientific ontology' as a catch-all for the different sorts of more specific claims that realists and antirealists alike may view as warranted. Given that ontology generally is concerned with questions about what exists and the natures of those things, we may view scientific ontology as concerned with the ontology of the world as revealed by our best science. Thus conceived, there are two aspects of scientific ontology that seem crucial to appreciating how different agents come to hold different beliefs about the epistemic upshot of any given domain of science. The first has to do with the *epistemic risk* one takes in believing any given proposition, and the second has to do with the nature of *risk assessment*. Let us consider these notions in turn.

As I intend the term, 'epistemic risk' is something we attribute to propositions, corresponding to our levels of confidence in their truth or falsity. If one views the claim that there is dark matter as especially epistemically risky, what this means is that one does not feel confident in assessing the

truth value of that claim given the evidence at one's disposal; in that case, it is entirely appropriate simply to suspend judgment and remain agnostic about the status of dark matter. On the other hand, if one views this claim as not being especially risky or not risky at all, this means that one has confidence in the judgment that it is true or false, based on one's evidence. To use the now-familiar terminology of 'degrees', assignments of high or low degrees of belief correspond to high levels of confidence in the truth or falsity of a proposition, respectively, and degrees toward the middle plausibly correspond to suspension. Now, what is most interesting presently is not merely that one may associate different levels of epistemic risk with different propositions, but that different epistemic agents often disagree about how much risk is properly associated with *one and the same* proposition, and this, as we will see, has profound implications for scientific ontology.

Consider the following sequence of epistemic judgments. Suppose that our best immunology has produced some putative information concerning how vaccinated people fare in comparison to unvaccinated people after infection by a virus, in terms of some strictly observable symptoms and test results. No doubt, to believe that these correlations obtain is to accept a measure of risk and, upon consideration of the evidence available, one may not find oneself convinced. Imagine, though, that one *is* convinced. Believing in these observable correlations does not, presumably, involve quite as much risk as believing in the underlying entities and processes in terms of which immunology theorizes about them—involving the spike proteins of viruses and the antibodies in human blood that are produced in response to the presence of a virus or vaccine—and on this point, considered judgments may once again diverge. One might, for example, regard this talk of underlying, unobservable entities and causal processes as merely a useful instrument for helping to produce well-warranted observable correlations. Or, alternatively, one might take the available evidence to furnish sufficiently strong support for belief in these ostensibly detectable (though nonetheless unobservable) things.

It is clear why extending belief in the reality of certain observable correlations to yet further beliefs about underlying, strictly unobservable entities and mechanisms involves an increase in epistemic risk. After all, the reality of observable correlations does not entail the reality of any particular underlying things, even if one takes the latter to provide good explanations of the former. Believing more in this way and adding further risk goes hand in hand. It is now perhaps also clear that this process of believing more and taking further risk can be iterated. The unobservable proteins, properties, and interactions described in molecular biochemistry in relation to vaccine efficacy are ostensibly amenable to putative detection and study via a number of scientific techniques of investigation, including microscopy. But one might go further: what about the existence and nature of

the molecules (amino acids) ostensibly making up these proteins, and the atoms that make up these molecules? In contemplating these yet further underlying realities we eventually enter realms of physics where, if we were to continue to iterate the process of asking deeper and deeper questions, we would find ourselves theorizing about things that are not even putatively amenable to detection, but which are posited to explain observations and detections of other things.

Let us take one last step in this imagined sequence of epistemic judgments and concomitantly increasing degrees of epistemic risk, into (arguably) realms of scientific ontology that are not even strictly part of what we call science, but are rather home to metaphysical theorizing intended to provide deeper or more fulsome interpretations or elaborations of recognizably scientific claims. This, recall, is the ambition of naturalized metaphysics. Throughout what I have described as a process of iteration, at each stage asking and attempting to answer questions about scientific ontology, with what we call science shading into what we call metaphysics over what may reasonably be described as a vague boundary, we have moved a significant distance along a spectrum of epistemic risk.

Having illustrated the notion of a spectrum, however, it remains to be clarified why agents disagree about where to draw lines between domains in which levels of risk are compatible with warranted belief, and domains in which they are not, resulting in suspension of belief instead. The most commonly invoked criteria of assessment are implicit above. First, there is what I (2017) have elsewhere called *empirical vulnerability*, which is a measure of the susceptibility of a claim to confirmation on the basis of our (actual) empirical evidence. This includes much of what is commonly discussed under the heading of empirical 'testing', but the idea of vulnerability is preferable, I think, in that it avoids the baggage of some previous debates. It is not uncommon, for instance, to find discussions of testing mired in disputes about whether there are such things as 'crucial experiments' (ones that definitively settle the choice between two or more rival theories or hypotheses), or disputes about how 'direct' the link between empirical evidence and a proposition must be in order for the production of the former to count as a probative test of the latter. As I intend it, empirical vulnerability is gradable, admitting of degrees, and immune to the distraction of debates over imagined thresholds of genuine testability.

A strict focus on empirical testing might also be thought to be insufficiently discerning in the sense that satisfying other criteria may also be regarded as 'tests' of the credibility of an ontological claim. A case in point is the second criterion I will flag here, namely, *explanatory power*, which is a measure of how well a theory or hypothesis explains something of scientific interest. A discussion of what, precisely, determines the quality of an explanation is beyond my ambitions here, but a number of markers

of good explanation are widely acknowledged: internal consistency; coherence with background knowledge; the possible unification of phenomena previously conceived as distinct; etc. In some cases, where one hopes to explain the results of experiments or other empirical interventions, assessments of empirical vulnerability and explanatory power may be interwoven. In other cases, though, these assessments come apart, as when, for example, something is posited merely or primarily on the basis that its existence would have remarkable explanatory power (say, in unifying an account of some domain of inquiry), and the posited entity is not itself susceptible to detection, or is susceptible only to highly indirect detection at best, in which case empirical vulnerability is low.

Just as in the case of empirical vulnerability, explanatory power is something that admits of degrees, and, as indicated earlier, there is often disagreement about what conclusions are properly drawn, in any given domain of scientific ontology, concerning whether a particular theory or hypothesis or claim is warranted on the basis of criteria such as these. Answers to questions regarding how empirically vulnerable an ontological claim is, or how much explanatory power it affords, or what sorts of balance of empirical vulnerability and explanatory power are sufficient to lower epistemic risk to the extent that belief is warranted, often vary dramatically between different agents. It should not be surprising that there are often no easy answers to such questions, or answers that are likely to generate universal assent among those who consider them. The resulting variety of agents' perceptions of epistemic risk are precisely what is at stake in debates about scientific ontology, which brings me now to the idea of stances.

An *epistemic stance*, as I will use the term here, is a collection of attitudes, values, aims, and policies concerning the assessment of evidence, which underwrites one's judgments about how far one should go, along any given spectrum of epistemic risk, in making ontological commitments. It thereby determines where an agent draws lines between domains in which belief seems warranted, and ones in which it seems more appropriate to be agnostic. As the list of elements comprising stances just mentioned suggests, they are not propositional—that is, they do not amount to factual claims, such as 'dark energy and dark matter combined make up 95% of the universe', or 'an asteroid impact on the Earth 65 million years ago was a major cause of the extinction of the dinosaurs'. Instead, the commitments making up a stance combine to form an orientation that shapes epistemic judgments. Stances are thus not believed as such but are rather adopted, and they play substantial roles in determining the doxastic attitudes agents form on the basis of considerations of our best science.

While I will not attempt to elaborate further here the notion of a stance, it may be helpful in passing, at least, to mention a few exemplifications merely as proofs of concept.[3] James (1897/1956) famously held that one

may make different tradeoffs in belief, in consideration of evidence, in hopes of striking a balance between believing too much and thereby opening the door to too many falsehoods, and believing too little and thereby missing out on too many truths. It is easy to find different tradeoffs arguably fitting something like this Jamesian mold in relation to scientific investigation. van Fraassen (2002), for instance, holds that empiricism is best conceived not as a doctrine—say, a factual claim or claims about legitimate sources of knowledge—but rather as a stance, which may incorporate anything from a relative disinterest in to an ample suspicion of attempts to produce knowledge by answering demands for explanation in terms of unobservable entities. In contrast, Psillos (2018) adopts what he calls a 'realist stance', on which many such attempts are, in fact, epistemically fruitful. In contrast to both empiricist-oriented stances and the more metaphysically oriented stances often associated with scientific realism, Fine (1986/1996) adopts a deflationary stance, incorporating something of disdainful attitude toward these sorts of epistemological disputes—an attitude not uncharacteristic of pragmatist approaches to scientific knowledge.

There is more to be said about the varieties of epistemic stances one might adopt, and the consequences of their adoption for conceptions of scientific ontology. For present purposes, however, rather than trawling through a compendium of stances that appear to be operative in the philosophy of science, and with the issue of naturalizing metaphysics firmly in view, let us proceed instead with a more detailed consideration of the two most important criteria of assessment of the outputs of scientific inquiry introduced earlier: empirical vulnerability and explanatory power. Since different stances conflict largely with respect to how they view the epistemic import of these criteria in relation to scientific theories and models and associated metaphysical theorizing, some clarity on the epistemological challenges inherent in their assessment will help to illuminate the prospects of naturalizing metaphysics. In the following two sections (Sections 11.3 and 11.4), then, let us dive deeper into empirical vulnerability and explanatory power, beginning with the former.

11.3 A Posteriori Versus A Priori Inquiry I: Empirical Vulnerability

Defining or clarifying the extension of 'what is natural' is not, on reflection, an easy thing to do in terms of some imagined, intrinsic properties of 'naturalness'. It is thus no surprise that definitions and clarifications often appeal, optimistically, to an extrinsic property instead: a contrast with the putatively *super*natural. On this rendering, and since the subject matters of the sciences are expressly not supernatural, it is understandable why naturalism is often understood in terms of some or other sort of deference to our best science on questions of ontology. Let me simply state rather

than discuss some of the ways in which this is problematic. For one thing, even if the modern sciences do not engage with the supernatural, historical antecedents often identified as scientific, such as ancient science, medieval science, and natural philosophy, were not devoid of supernatural considerations. And even restricting ourselves to the modern period, it is unclear that the natural and the supernatural are exhaustive categories. Arguably, some phenomena belong to neither or extend beyond both, such as the subject matters of the social sciences, social and political philosophy, ethics, and aesthetics. If this is so, then deference to the (natural) sciences on all questions of ontology seems overly restrictive, even for those who reject the supernatural.

Most important here, though, is the fact that *even if* the natural sciences *were* the ultimate arbiters of what there is and what those things are like, for reasons we have already noted, 'being science' does not carry much if any weight all by itself. It carries weight only insofar as some parts of science are well confirmed. Given the ubiquity of idealizations and abstractions in scientific theories and models, the continual growth and change in scientific descriptions over time, and the fact that no one thinks that this ongoing development has come to an end, confirmation in this sphere is generally held to be *selective*. In this context, empirical vulnerability is an especially weighty consideration, and this is especially transparent in debates about scientific realism, where warranted belief is often associated most strongly with specific parts of theories or models that are judged to be well confirmed in virtue of specifically empirical investigations: observation, detection, intervention, manipulation, experimentation, etc. All of this said, the question of which aspects of theories and models are thus confirmed is contentious. As I will now illustrate, different epistemic stances foster different judgments regarding the epistemic credentials of different parts of science, resulting from different assessments of their empirical vulnerability.

Imagine a seventeenth-century physicist, discussing Newton's newly proposed account of gravitation with a colleague. Both think the theory passes impressive empirical tests, but one holds that this confirms, to some impressive degree, the existence of gravitational forces, while the other holds that the evidence does not take one quite that far, and she is agnostic about the existence of forces *per se*—after all, what are they?; Newton does not say. Nonetheless, she takes the evidence to confirm the existence of empirically detectable regularities predicted by Newton's equations. As we know, natural philosophers at the time were divided on questions about the nature of forces. Or consider a more recent case. The existence of the neutrino, a subatomic particle, was posited by Enrico Fermi and Wolfgang Pauli in the 1930s. In detections of certain processes of atomic decay, certain quantities of mass-energy and angular momentum seemed to be

missing, but a hitherto undetected particle would account for the conservation of these quantities. In other words, a neutrino would fill the gap. But the first experiments to detect neutrinos themselves, conducted by Frederick Reines and Clyde Cowan, did not occur until the 1950s. One might tell a similar story about the positing of the Higgs boson in 1964, again, not as a flight of fancy but for empirical reasons. However, the detection of the Higgs particle, announced by CERN (the European Council for Nuclear Research), did not come until 2012, almost 50 years later.

Let me extract two points from these examples. The first is that empirical vulnerability is not a bivalent property. That is, it is not generally the case that in light of empirical evidence, a scientific description is either vulnerable or invulnerable. Rather, empirical vulnerability is something that admits of degrees. This is why the sorts of cases mentioned earlier, of contentions regarding the existence of forces or neutrinos or bosons, are so interesting. Over time, their empirical vulnerability may be assessed as growing, resulting in a lowering of epistemic risk. That said, an assessment of the consequences of this for epistemic warrant is inevitably challenging. Even if epistemic risk is judged to decrease as a function of growing empirical vulnerability, there is room for disagreement regarding *how* vulnerable any particular contention may be at any given time—about how 'tight' the connection must be between specific ontological posits and specific empirical tests—to warrant belief. Add to this the ubiquity of idealizations and abstractions, and the fact that there is no rule book for distributing whatever confirmation flows from empirical evidence among aspects of theories and models that may be regarded as better established or more speculative to begin with, and the challenges to assessing warrant mount further.

This leads to a second point one may extract from case studies in purportedly scientific ontology: adopting a simpleminded division according to which science is viewed as empirical, or *a posteriori*, and metaphysics is viewed in contrast as non-empirical, or *a priori*, is a misleading way to proceed in thinking about the possibility of naturalizing metaphysics. Judgments about the extent to which empirical vulnerability lowers epistemic risk in any given domain of scientific ontology generally make recourse to both *a posteriori* considerations, involving empirical evidence, *and a priori* considerations, involving attitudes toward epistemic risk and policies for belief formation inherent in the epistemic stances we bring to these assessments (as discussed in Section 11.2). Instead of an overly simple picture according to which one form of inquiry, metaphysics, is conceived as being tethered in some way to another, wholly distinct form of inquiry, science, what we have instead is a spectrum of possible ontological commitments ranging from the highly empirical vulnerable to the increasingly empirically invulnerable. And where any particular ontological posit falls along this spectrum, and how great the epistemic risk is *there*, are matters of

variable judgment. Where science ends and where metaphysics begins, if such a question is intelligible at all, is not a matter of simple definition.

11.4 A Posteriori Versus A Priori Inquiry II: Explanatory Power

Earlier I indicated that in addition to empirical vulnerability, explanatory power is also a major factor in assessing the proper scope of scientific ontology. A significant disanalogy between assessments of the two, however, is that while everyone agrees that the former is an evidentially weighty consideration (even if there is disagreement about the extent of this weight in particular cases), there is huge disagreement about whether the latter has much if any evidential weight at all. On the other hand, assessments of these two factors are analogous in that, in both cases, epistemic stances play a substantial role in determining the extent to which agents hold them to be relevant to lowering epistemic risk, and thus warranting belief. Let us now consider how this works in the case of explanatory power, and the epistemic challenges that arise in assessing it.

With prospects for naturalizing metaphysics in mind, it is helpful to reflect on the question of where exactly explanatory considerations are most contentious. Very abstractly speaking, many would agree that good explanations are at least epistemically suggestive: an explanans that gives what seems like a convincing account of some explanandum of interest is generally taken seriously as a promising candidate for belief. This is the kernel of what is often called 'inference to the best explanation': good-making features of explanation are suggestive of their truth; an explanans that is superior to rival explanantia with respect to such features is thus more likely to be true, and may well be warranted as a result. Unqualified, however, this is immediately contentious, because it places no restriction on the domains in which this sort of inferential practice is reasonable—it gives no indication of the circumstances in which explanations are, in fact, convincing. One who is ready to grant that inference to the best explanation regarding mechanisms of enzymatic catalysis in a cell is epistemically probative is surely not committed, on that basis alone, to the veridicality of such inferences concerning mechanisms of instantiation of universals by concrete particulars. These cases involve very different sorts of evidence and perhaps even different sorts of reasoning, as further examples may now help to illuminate.

It is the strength of specific explanations in specific contexts that is relevant to epistemic warrant, not the abstract idea of explanation. Consider, for instance, what one might call framework assumptions, or presuppositions, of particular bodies of scientific practice associated with a theory or theories and concomitant sets of models in application to target systems in the world. In order for such practice to get off the ground, we must render

phenomena of interest into systems of ontological categories and properties in terms of which they can be described and investigated. This regimentation of concepts and principles is a necessary condition, as part of a background one must have in place before it is even possible to engage in systematic inquiry. In this way, Euclidean geometry is presuppositional in the context of Newtonian physics, furnishing a structure in terms of which one may describe spatial relations. Later, in the context of Einstein's development of his theory of Special Relativity, an analogous status pertains to the principle that the speed of light is the same in all directions. Framework assumptions are not generally empirically testable in any direct sense, via causal isolation or detection or manipulation, but serve as a platform for theory building, modeling, and experimentation, and in this way form part of an overarching explanation for the empirical success of a theory.

In addition to *ab initio* concepts and principles that are at best indirectly empirically vulnerable, explanatory roles are also played by aspects of scientific ontology that are ostensibly revealed by inferences *from* our best science—inferences that are made *post hoc*, as it were. Some of this is implicit in the earlier discussion of empirical vulnerability: taking a claim about an ontological posit to be highly empirically vulnerable and thus believing (or disbelieving) in the entity posited may be described in explanatory terms, that is, in terms of finding convincing the explanation of the empirical data one would possess if the entity did (or did not) exist. We often say that the fact that something exists or has certain properties would explain or partially explain the data we produce in detections or experiments, for instance. However, there are other cases of inferences to aspects of scientific ontology in which explanatory power is largely detached from or only very loosely connected to empirical vulnerability.

Consider, for example, debates about the ontological nature of biological species. Is a species something best described in terms of the long-standing philosophical tradition of theorizing about natural kinds, that is, as a category of things comprising members that share certain qualitative properties? Or is it better described as an individual, which comes into existence during a speciation event and departs from existence with extinction? If a species is an individual, an organism belonging to it is not properly described as a member *per se*, but rather as a part, which implies a different mereological relationship and perhaps further consequences for our conception of evolutionary biology. Now, these are just two ways of thinking about species—there are many more—and I will not attempt to do justice to the many debates surrounding them here. The point for present purposes is simply to illustrate the idea that in some cases, theorizing about scientific ontology takes the form of inferences made on the basis of our best science in which empirical vulnerability is largely irrelevant. What matters instead, in a debate like this, are assessments of the quality

of the explanations such theorizing provides in relation to the scientific phenomena at issue.

Mirroring the conclusions above regarding empirical vulnerability (at the end of Section 11.3), let me extract two morals from the preceding examples in connection with explanatory power. The first is that explanatory power is not bivalent; rather, it is something that admits of degrees. In cases where it is invoked in assessments of epistemic risk, ranging over various kinds of framework assumptions as well as inferences from scientific theories and models to aspects of ontology, the power of explanatory considerations in gauging warrant is something that varies with the perceived quality of explanations in specific contexts. Furthermore, there is room for disagreement in these assessments about whether warrant is enhanced sufficiently for belief or disbelief, or whether we should suspend belief instead, which leads to the second moral. Assessing the confirmational import (or lack thereof) of explanatory power is not a matter of separating cases into two types—ones in which explanations are empirical or *a posteriori*, thus 'belonging' to science, and ones in which they are non-empirical or *a priori*, thus 'belonging' to metaphysics. Judgments of explanatory power in cases of scientific ontology turn on both *a posteriori* and *a priori* considerations, the latter involving attitudes toward explanation that vary between those with different epistemic stances.

It is evident once again that conceiving of (some part of) metaphysics as a form of inquiry that may be tethered in some manner to science, conceived as something distinct, is wrongheaded. What we have instead is a spectrum of possible ontological commitments ranging from those associated with high degrees of explanatory power to those having little or none. Where precisely any given aspect of scientific ontology falls along this spectrum and what impact this has on assessments of epistemic risk are susceptible to variable judgments driven by different epistemic stances. Once again, these challenges of determining warrant problematize any naïve hope of neatly distinguishing what is 'genuinely scientific' from what is not scientific but 'metaphysical'.

11.5 An Open Question About Naturalizing Scientific Epistemology

Let us take stock. My central focus in this chapter has been the aim of naturalizing metaphysics, which has attracted significant attention recently from philosophers interested in relationships between the empirical sciences and metaphysical theorizing. This aim incorporates the hope of doing metaphysics in a way that enhances its epistemic credentials by connecting it in a suitable manner to scientific inquiry, which naturalists commonly regard as a privileged source of knowledge of the world. On reflection, though, it emerges that the question of what sort of knowledge the *sciences*

yield is itself subject to dispute, which complicates the notion that we may naturalize metaphysics by simply anchoring it, or some portion of it, to our best science. I have suggested that different takes on what beliefs scientific theories and models license can be traced to different epistemic stances adopted by agents, which shape their assessments of the warrant of specific claims about scientific ontology in virtue of their empirical vulnerability and explanatory power. Challenges inherent in these assessments, and in conflicts between the different stances that inform them, are illustrative of the significant complexity of attempting to naturalize metaphysics.

It is not my aim to dissolve this complexity. Indeed, I do not regard it as a problem in need of resolution, but rather as a constitutive feature of scientific ontology conceived as an epistemic project. That said, and in conclusion, let me address a potential objection to the picture I have sketched above. Ironically, this objection concerns the question of whether the account itself is properly subject to yet another sort of naturalizing impulse, namely, to the ideal of naturalizing *epistemology*.

I have elaborated a view according to which answers to questions about what beliefs are warranted by science, along spectra of ontological depth characterized by decreasing empirical vulnerability and explanatory power, with highly empirical science at one end and speculative metaphysics at the other, are facilitated by different epistemic stances. These stances reflect differences in the sorts of things agents value, epistemically, including certain kinds of information, explanation, evidence, argument, and intuitive judgments about how these things support ontological claims and to what extent. As a result, agents holding different stances draw lines between science and metaphysics in different ways, and even those sympathetic to the idea that some metaphysics may produce warranted beliefs may disagree about which parts of it (if any), in virtue of their relation to science, fit the bill. This is, no doubt, a philosophical account of naturalizing, based on a philosophical analysis of reasons given by parties to debates about scientific ontology for why they believe what they do and not otherwise. But now, one may wonder: in just the way that naturalists may think it necessary to rescue scientific ontology from the excesses of armchair metaphysics, should they not also be concerned about the excesses of armchair epistemology?

Granted, as I have described them, epistemic stances are comprised of things—attitudes, values, aims, etc.—that may be amenable to scientific study, perhaps in the social sciences or empirical psychology, or perhaps even in the brain sciences, given the claim that stances function cognitively in certain ways with respect to belief. One might thus contend, in a naturalist spirit, that the very account given above, of the ways in which we manage epistemological challenges inherent in scientific ontology, depends for its credibility on support from the relevant domain or domains of empirical

science, to vouch for the existence and function of stances (cf. Bryant 2021: 12). This is an intriguing suggestion, especially for a thoroughgoing naturalist, but as I will now contend, it does not serve to undermine the view of naturalizing metaphysics and scientific ontology presented here. At best, it points to a tantalizing question regarding the prospects for naturalized epistemology in this sphere, and whether possible answers to this question are likely to generate skepticism about epistemic stances.

Let us first acknowledge that, of course, it may be possible to study belief-generation processes (ostensibly) involving stances in scientific ways. In noting this, I do not mean to downplay the rigor of systematic, philosophical pursuits as recognized in the inclusion of philosophy and other humanities disciplines as *scientific* as per the German conception of *Wissenschaft*. I simply mean to acknowledge that this does not preclude, all by itself, the further investigation of aspects of human reason, including how we assess evidence and form beliefs, by the aforementioned natural and social sciences. The question here is not so much whether they *can* investigate, but rather what we might reasonably expect them to reveal. It is useful here to distinguish two quite different questions, one concerning what reasons we may have to believe that there are such things as epistemic stances in the first place, and a second concerning the variety of things we might learn about them via different modes of inquiry. This will allow us to appreciate, I submit, that whatever prospects there may be for naturalized epistemology in this arena, they pertain to the latter hope of learning more about the nature of stances and how they function, not the former question as to whether they exist and function in our mental lives at all.

Take first the latter hope. Perhaps it is within the capacities of neuroscience or empirical psychology (or will be one day) to reveal facts about stances—perhaps there are, say, correlations to be found between cognitive or psychological features of individuals and the stances they adopt. Perhaps the investigations of a present or future social science would reveal different sorts of correlations between the possession of certain stances and the social, institutional, or cultural environments or commitments of epistemic agents, all of which may be complementary to or further explicative of a more strictly philosophical, phenomenological study of how different stances strike agents as appealing or unappealing in ways that accord with their values, epistemic and otherwise. But note that all of this imaginary science is *predicated* on the aim of investigating the nature of epistemic stances; it is not geared toward establishing whether there are such things. Consider an analogy: each of the natural and social sciences investigate causal interactions and processes in terms of their own subject matters and modes of analysis. But none of them is equipped to answer the prior question of whether there *is* such a thing as causation and, if so, what it is, exactly. These questions are not within their ken. They belong to philosophy.

What, then, of the prior question of what reasons we have to believe that there are such things as epistemic stances, which play the roles I have attributed to them in determining how we think about the relationship between science and metaphysics? Here I am reminded of Garfinkel's (1981) energizing resistance to a certain kind of reductionism. Garfinkel's interest was in social theory; he argued for the salutary claim that it is often a mistake simply to assume that all questions can be answered in explanatorily satisfactory ways by focusing on some more basic, micro-foundational science. Once we are clear on what question is being asked, we may then find that a compelling answer must be given at a certain level of analysis, not necessarily a micro-foundational one. 'Explanation seeks its own level' (Garfinkel, 1981: 59 ff.). I suspect that it is a failure to appreciate this that drives the mistaken assumption that only evidence from the natural or social sciences could establish the reality of stances, and produce an understanding of how they shape assessments of scientific ontology and, thereby, how we may understand the quest to naturalize metaphysics.

There are compelling philosophical reasons to theorize about the nature and role of epistemic stances. After all the arguments about scientific ontology are laid on the table and considered by experts on all sides, there are irresolvable differences. Those more inclined toward sensibilities associated with stricter forms of empiricism or instrumentalism view anything beyond detection by means of the unaided senses as metaphysical and unsusceptible to belief. Others extend belief to what seem empirically well-confirmed or explanatorily-crucial ontological posits but no further, drawing a line between scientific ontology and the merely metaphysical in a different place. Others go further, believing in aspects of ontology—the nature of objects, properties, modality, laws, and so on—that are not the subject matters of science *per se*, but that are connected to science proper, they hold, in such a way as to merit belief. When we trace, as only philosophers are equipped to do, the foundational commitments that eventuate in such differences, we ultimately come to the epistemic stances of agents. It is only by exploring *this* terrain that we may understand what it could mean to naturalize metaphysics, and all of the epistemological challenges this entails.

Notes

1 For just a few book-length examples, see Morganti (2013), Ross et al. (2013), Schrenk (2016), and Chakravartty (2017).
2 A more detailed and comprehensive exploration of this account can be found in Chakravartty (2017).
3 For further details, including more extensive examples and references, see Chakravartty (2018).

References

Bryant, A. (2021) 'A Thousand Flowers on the Road to Epistemic Anarchy: Comments on Chakravartty's *Scientific Ontology*'. *Dialogue* 60(1), 1–13.
Chakravartty, A. (2017) *Scientific Ontology: Integrating Naturalized Metaphysics and Voluntarist Epistemology*. New York: Oxford University Press.
Chakravartty, A. (2018) 'Realism, Antirealism, Epistemic Stances, and Voluntarism'. In: J. Saatsi (ed.), *The Routledge Handbook of Scientific Realism*. Abingdon: Routledge, pp. 225–236.
Fine, A. (1986/1996) *The Shaky Game: Einstein, Realism and The Quantum Theory*. 2nd edn. Chicago, IL: University of Chicago Press.
Garfinkel, A. (1981) *Forms of Explanation: Rethinking the Questions in Social Theory*. New Haven, CT: Yale University Press.
James, W. (1897/1956) 'The Will to Believe'. In: *The Will to Believe, and Other Essays in Popular Philosophy*. New York: Dover, pp. 1–31.
Morganti, M. (2013) *Combining Science and Metaphysics: Contemporary Physics, Conceptual Revision and Common Sense*. London: Palgrave Macmillan.
Psillos, S. (2018) 'The Realist Turn in the Philosophy of Science'. In: J. Saatsi (ed.), *The Routledge Handbook of Scientific Realism*. Abingdon: Routledge, pp. 20–34.
Ross, D., Ladyman, J. and Kincaid, H. (eds.). (2013) *Scientific Metaphysics*. Oxford: Oxford University Press.
Schrenk, M. (2016) *Metaphysics of Science: A Systematic and Historical Introduction*. Abingdon: Routledge.
van Fraassen, B. C. (2002) *The Empirical Stance*. New Haven, CT: Yale University Press.