

THE SCIENCE OF LOST WORLDS

Adrian Mitchell Currie

ISBN 978-1-77385-713-8

THIS BOOK IS AN OPEN ACCESS E-BOOK. It is an electronic version of a book that can be purchased in physical form through any bookseller or on-line retailer, or from our distributors. Please support this open access publication by requesting that your university purchase a print copy of this book, or by purchasing a copy yourself. If you have any questions, please contact us at ucpress@ucalgary.ca

Cover Art: The artwork on the cover of this book is not open access and falls under traditional copyright provisions; it cannot be reproduced in any way without written permission of the artists and their agents. The cover can be displayed as a complete cover image for the purposes of publicizing this work, but the artwork cannot be extracted from the context of the cover of this specific work without breaching the artist's copyright.

COPYRIGHT NOTICE: This open-access work is published under a Creative Commons licence. This means that you are free to copy, distribute, display or perform the work as long as you clearly attribute the work to its authors and publisher, that you do not use this work for any commercial gain in any form, and that you in no way alter, transform, or build on the work outside of its use in normal academic scholarship without our express permission. If you want to reuse or distribute the work, you must inform its new audience of the licence terms of this work. For more information, see details of the Creative Commons licence at: <http://creativecommons.org/licenses/by-nc-nd/4.0/>

UNDER THE CREATIVE COMMONS LICENCE YOU MAY:

- read and store this document free of charge;
- distribute it for personal use free of charge;
- print sections of the work for personal use;
- read or perform parts of the work in a context where no financial transactions take place.

UNDER THE CREATIVE COMMONS LICENCE YOU MAY NOT:

- gain financially from the work in any way;
- sell the work or seek monies in relation to the distribution of the work;
- use the work in any commercial activity of any kind;
- profit a third party indirectly via use or distribution of the work;
- distribute in or through a commercial body (with the exception of academic usage within educational institutions such as schools and universities);
- reproduce, distribute, or store the cover image outside of its function as a cover of this work;
- alter or build on the work outside of normal academic scholarship.



Acknowledgement: We acknowledge the wording around open access used by Australian publisher, **re.press**, and thank them for giving us permission to adapt their wording to our policy <http://www.re-press.org>

Peculiar Worlds

A bespectacled *Allosaurus* takes the next question. They clear their throat and begin speaking in a manner suggestive of someone who, somewhat like Prof. Ichthy, is accustomed to speaking in whole paragraphs.

“Doctor Professor Ichthy, thank you for your interesting lecture and congratulations on your new chair. I must admit, however, that I found much in your lecture deeply troubling. Through the lens of your account of the past, we see a world not of order and stability, but of seemingly endless fragility and chaos. Perhaps better to not say ‘chaos,’ but ‘chaos-ing.’ *Komplexelieferung*, if I may. We can, I’d say, contrast this chaos with the meaning we find in our immediate experience. From our own perspectives, it seems to me, we form the centre of our own local universes, universes containing and shaped by familiar things. Our friends, family, the day-to-day experiences that make up our lives on scales appropriate to our status as finite beings. This stability is the phenomenology of our immediate experience and provides, I would say, the grounding of the very possibility of our lives’ meanings.

Your world vision—‘world-making’ I’d call it—it seems to me, demands that we abstract from that familiar perspective. This confronts us with the sheer fragility of our own *welterschaffend*, of all that is meaningful to us. It seems our beliefs, our relationships, our very moral worldview, are contingent upon not just some tiny happenstance, but some near-infinite chain of tiny happenstances. Where around us, we see a local reality from which we might find glimpses of true meaning in eternity, you would deny us that eternity as an illusion. We light tiny candles of order around us, about which only chaos swirls.

I ask, Prof. Ichthy, how can we make sense of a world of such unerring hostility, how should we find meaning in such a meaningless world?"

Prof. Ichthy pauses, partly making sure the *Allosaurus* has finished, but also wondering how to proceed. They sigh, "No one warned me philosophers would be attending . . ."




ILLUSTRATION 8.1 Ichthy Through the Looking Glass

In the first chapter, I compared Prof. Ichthy with another fictional Ichthyosaur: that depicted in De la Beche's *Awful Changes*. Prof. Ichthy, through their emphasis on Broadfoot's uniqueness, paints a highly contingent picture of the past. I took Professor Ichthyosaurus to represent the steady-state view of Earth's history as defended by Lyell (or as Lyell was interpreted by his contemporaries). As such, we might imagine Professor Ichthyosaurus being less impressed by the contingency of the past. For him, the patterns uncovered in the past reveal regularities. Loss occurs, no doubt, but that loss is in some sense inconsequential. On this kind of view, if you know the broad-scale pattern, then you know most of what matters. The rest, as Ernest Rutherford might say, is stamp-collecting (Johnson 2007).

For Professor Ichthyosaurus, history matters insofar as examining the past might be necessary for discovering the cycles governing the Earth, but once those cycles are established, it matters not. For Prof. Ichthy, by contrast, history does matter. That is, even if we have a good understanding of the Earth's cycles, we still need to attend to the details of history's particular unfolding.

My aim in this chapter is to consider whether the position developed thus far can play any role in the metaphysical dispute between our fictional ichthyosaurs.

I'll start by developing previous work toward an account of when history matters, building on my notion of "peculiarity." I'll then consider under what conditions scientific results or practices might provide insight into metaphysics. This will underwrite arguments that peculiarity is, at least to some extent, supported by the strategies historical scientists adopt in the face of loss. I'll suggest strategic perspectivalism is a method adapted to, and adopted in light of, a peculiar world. Finally, I'll link history's peculiarity to a set of metaphysical views that emphasize the world's disunified, dappled, or processual nature.

My reference to "stamp collecting" may have got you thinking about laws, a topic I covered somewhat dismissively in section 4 of chapter 5. There, I argued that appeals to general laws are sometimes an important aspect of strategic perspectivalism, but their tendency to abstraction typically makes their epistemic utility limited. Suffice to say, in history we see both fragility and robustness, and robust patterns and mechanisms seem to play law-like roles in scientific explanation and investigation, at least sometimes (Sterelny 2016). In a peculiar world, we'll expect law-like regularities to hold fairly

often, but only for a limited scope. Regardless, I'll not directly link my discussion to that of laws here.

As we'll see, I won't be defending a *scale-independent* metaphysics, that is, I'll not maintain that metaphysical views hold regardless of time, place, and level of description. Instead, I'll be interested in what we might call the *palaeobiological scale*. This is an admittedly ambiguous way of putting things, but I'll take it to refer to something like the scales accessible through the fossil record, which includes everything from long-term, millions-of-years-long macroevolutionary trends, through to single-organism properties such as ontogeny, physiology, and so forth. One way of generating the set of scales of interest is to include those captured by the perspectival tools palaeobiologists use. These are certainly broad, but do more or less fall within a certain set of ranges. We might further limit these scales by appealing to some perspectival tools rather than others: perhaps we should exclude the tools that rely on, say, chemical scales insofar as they are used by palaeobiologists to infer larger-scale phenomena. Regardless, I don't think my arguments rely on specifying a scale or set of scales. The point is just that my position isn't scale-independent.

Another way of putting this is to say that the local epistemic activity of the science of lost worlds licences a local, scale-dependent metaphysics, and I'll restrict myself to that.

1. When Does History Matter?

The claim that "history matters" is common, but its meaning is somewhat obscure. Let's consider a few examples and then zoom in on the most relevant to Professors Ichthy and Ichthyosaurus.

In political science, Charles Tilly argues that "explanatory political science can hardly get anywhere without relying on careful historical analysis" (Tilly 2006, 417). Here, history matters for developing decent explanations of contemporary phenomena. John Tosh argues that "to know that the past can illuminate the contours of the present is to be better equipped to make intelligent decisions about difficult public issues" (Tosh 2019, viii). Here, history matters for informing our future-oriented decision-making. For Gerda Lerner, history matters because it informs and challenges political views, describing her work as taking place in the backdrop of culture wars "being fought about the meaning of history, when the history of minority and marginalized is being challenged for its very existence" by those who see history

as the property of ruling elites and as a vehicle for the justification of states and other institutions of power” (Lerner 1998, xvi–xvii). Here, uncovering and telling the stories of once-invisible marginalized people can confront and undermine received histories.

Each of these notions, that specific details of history are crucial for making sense of the past and present, that history matters for tackling the future, and that history matters for challenging myth and emphasizing historically silenced perspectives, is worthy of serious attention. Here, I’m interested in something like that first notion. The idea, I take it, is that for some investigations we cannot simply avail ourselves of a temporal snapshot, or the short-term dynamics, of our investigative target. We cannot screen off the past. We also need to incorporate a detailed, long-term perspective on the target’s trajectory. We need to ask where our subject came from and what processes maintain it. But this remains ambiguous: how long is “long term”? Why might we be forced to venture into deeper time? And what kinds of investigations do we have in mind?

As a first pass, we might understand history mattering for some target when aspects of history must be incorporated into an *explanation* of that target. But this also doesn’t help us much: explanation is notoriously description- and interest-sensitive. If I’m just not interested in those long-term dynamics or the origination of some target, then what is there to force me to consider these elements? Again, comparing Prof. Ichthy with Professor Ichthyosaurus gives some help.

For Prof. Ichthy, the past and present differ significantly. For them, new entities, dynamics, etc., arise over time, and these might underwrite regularities. However, these regularities are fragile. In explaining past denizens, then, we are led into increasing complexity and detail; we must incorporate multiple causal factors, entities, and dynamics. For Professor Ichthyosaur, the differences between past and present are inconsequential: regularities are stable across time, and differences in initial conditions and subsequent paths will nonetheless lead to similar outcomes. As such, much simpler explanations appealing to those laws will do. We’re not led to complexity but generality.¹

In chapter 1’s discussion of loss, I mentioned its inverse: novelty and innovation. Over time, we don’t simply lose denizens, but new denizens arise.

1 This point could likely be cashed out with reference to my distinction between “simple” and “complex” narratives (Currie 2014, 2019f).

Consider Cuthill and Conway Morris' discussion of rangeomorphs and the fall of the Avalonian: "The unique rangeomorph fronds were fractal, surface area specialists of the Ediacaran. At the Cambrian explosion, the ecological and geochemical conditions to which the rangeomorphs were optimized ceased to exist, and their extraordinary body plan was lost from life's repertoire" (Cuthill and Conway Morris 2014, 13125).

We don't simply see the loss of rangeomorphs, but of the "ecological and geochemical conditions" that underwrote their survival. As we saw in chapter 5, those conditions were replaced by those of the later Ediacaran, which included better-oxygenated oceans and processes of bioturbation generated by the evolution of burrowing metazoans. By the Cambrian, the menagerie included critters with specialized adaptations for movement, sensing, and consumption: scavenging, predation, predator avoidance, and mate location. And these underwrite the diurnal vertical migration from the ocean's surface to its depths.

That's an example of history mattering: a complex sequence of events, a trajectory, takes us from one world order to another. To explain how we got from there to here, we must trace that sequence.

Philosophers often make sense of a lawless, or at least fragile, history by appealing to notions of *contingency* (e.g., Desjardins 2011; Millstein 2011; Currie 2012; Powell 2020).² It seems plausible that the loss of the Ediacaran worlds and the arrival of the Mesozoic worlds were not inevitable. It could have been different. But "could," like explanation, is a notoriously tricky idea, highly sensitive to context (Lewis 1976). Do we mean the events were indeterminate? Do we mean they weren't determined by some set of causal factors? Do we just mean we can imagine, or that it is possible, by the laws of physics, that something else could have occurred?

John Beatty divides claims about contingency in biology into two related notions (Beatty 1995, 2006). First, *contingent upon*. This is the idea that the occurrence of some outcome relies upon some earlier event. If Butterfield is right, the evolution of complex metazoans first required the evolution of simpler metazoans, as their structuring of the oxygen environment paved the way to increased complexity. If the biochemical explanation is right, a condition on the evolution of complex metazoans was an increase in global

2 I'm here interested in contingency as a claim about history, particularly macroevolution, as opposed to the contingency of *science* (e.g., Kidd 2011; Soler, Trizio, and Pickering 2016).

oxygen. The views have differing ideas about what the target, the evolution of metazoans, is contingent upon: what enabling conditions were required.

Second, *contingent per se*. This is the idea that some event wasn't guaranteed to occur. A strong version of this idea denies determinism (Beatty 2016). A weaker, and I think more useful, interpretation specifies a set of conditions, and says that those conditions are not sufficient to generate the outcome (Powell and Mariscal 2015). For instance, Alison McConwell and I have argued that palaeobiological arguments denying natural selection's special role in explaining life's shape appeal to contingency *per se* in this sense. When Stephen Jay Gould and others argue from the contingency of, say, the Cambrian explosion to the autonomy of macroevolution from microevolution, they mean that the dynamics and theories associated with microevolution are insufficient—cannot be “scaled up”—to generate the dynamics of macroevolution (McConwell and Currie 2017). As such, macroevolution is contingent *per se* regarding microevolutionary processes. Similarly, if Butterfield is right, appeal to biochemical conditions alone is insufficient to generate the metazoan radiations that occurred after the Avalonian. Referring to those conditions, the radiation is contingent *per se*.

This contingency talk amounts to discussion of the historical conditions required for the emergence of some outcome. The outcome is contingent *upon* a set of conditions, it is contingent *per se* given some (insufficient) conditions. To see why this might matter for our discussion, consider concepts that contrast with contingency, like inevitability and determination. It might be that, say, once oxygen levels rose sufficiently, the evolution of complex metazoans was set: holding fixed some unlikely destabilizing events (say, some kind of extraterrestrial impact), that condition was sufficient for metazoan complexity. Or it could be that no matter what kind of path life took, it would converge on something like complex metazoans (Powell 2020). In these cases, history seems to matter less than cases where the evolution of metazoan complexity requires a set of complex, themselves contingent, preconditions. So, you might think that history matters vis-à-vis some outcome when the conditions of its occurrence are sufficiently contingent. We might say that outcomes that are highly *path dependent*, that is, whose occurrence relies on particular, themselves contingent, sequences of events, are particularly prone to exhibit “historicity” (Ereshefsky 2014).

So, contingency gets us close to what we need. However, as thus far analysed, “contingency” is too focused on outcomes. We don't simply care about contingent events and their fragility, but also about robustness and

stability. The well-mixed oxygenated oceans of Earth have been remarkably stable, partly maintained by diurnal migrations for something like 400 million years. But that stability relies on maintenance from ongoing patterns of metazoan migration, and these only arose due to a complex series of events. This is another sense in which history might matter: by explaining stability in terms of potentially fragile maintaining causes.

In earlier work (Currie 2019, 31), I tried to capture this kind of notion with what I'll call *peculiarity*.³ Here's the definition I'll use here:

Some target is peculiar to the extent that its actual properties and modal profile are sensitive to the properties of the processes that generate or maintain it.⁴ The more processes the target is sensitive to, and how fine-grained the relationship between the target's properties and profile and those generating and maintaining processes, the more peculiar it is.

So, to unpack. We can understand a "modal profile" as what a target can do, what behaviours are open to it, and what dispositions it has. This includes the conditions under which the target's behaviour remains stable and the conditions under which it maintains its identity, that is, remains being the thing that it is. Both a target's modal profile and its actual properties have dependency relations to a set of processes that generate the target (or its properties) and a set that maintains those properties and capacities.

To grab a philosophical mainstay, compare the modal profile of a statue with that of the clay it is made from. The statue has particular generation conditions, say, being shaped by the artist, and particular conditions under which it ceases to exist, say, being crushed and broken.⁵ By contrast, the clay's modal profile differs: it needn't be shaped by the artist and can survive being crushed. Different targets, then, are differently peculiar depending on what their modal profiles are sensitive to. Further, the actual properties (as well as

3 Recently, Robert Northcott has discussed the properties of sciences under conditions of *fragility* in ways that have significant overlaps with the themes discussed in this chapter. For Northcott, "fragility" differs from "peculiarity" due to the latter's modal character, and the former being about the predictability of token relations. I don't use "fragility" in Northcott's specialized sense in this chapter (Northcott 2025). Rob Wilson's work on the "fragile sciences" takes fragility in a similarly broad way to mine (Wilson 2004, 2005).

4 In 2019, I defined *peculiarity* purely in terms of modal properties. I think this was a mistake: we care about the actual distribution of properties as well.

5 Elizabeth Barnes has a useful discussion of this old chestnut (Barnes 2023).

the modal profile) of the statue might be differently sensitive to generating and maintaining processes. Compare statues produced by a mould with those carved by hand. The actual properties of the moulded statues will presumably be more regular in appearance, while the hand-carved ones will be more varied. This is because the properties of the hand-carved statues will depend on features of the artisans, so their peculiarity is partly determined by those features, whereas the others' properties will depend on the mould. Whether we want to say that one set is more peculiar than the others will turn on factors such as the idiosyncrasies of the craftspeople, the moulds, and so on.

Peculiarity differs from other accounts of contingency (although I take it to be part of the same family). An outcome might be *contingent upon* some previous event or process, or *contingent per se* relative to some process or set of processes, but this is only part of whether or not we should consider the outcome highly peculiar. This will depend on the number of maintaining or generating processes involved and how sensitive the outcome is to these. Peculiarity captures a kind of second-order contingency. Outcomes can be highly robust and still be peculiar: it is simply that their robustness is highly dependent on a set of generating or maintaining processes.

Compare the biochemical explanation of metazoan complexity with Butterfield's. We can understand an explanation's "peculiarity" by the extent to which it appeals to its target's peculiarity. In the biochemical explanation, metazoan complexity is sensitive to the Earth being sufficiently oxygenated, a process that they argue occurred more or less as a result of the changes that occurred during the great oxygenation event. By contrast, Butterfield argues that metazoan complexity is a solution to insufficient oceanic oxygen. In his narrative, the event was sensitive to the evolution of (1) various oxygen-structuring critters (rangeomorphs, sponges, etc.), (2) phytoplankton, (3) burrowing, and (4) key morphological innovations in metazoa. Both explanations are peculiar, but we can understand Butterfield's explanation as being *more peculiar* insofar as it appeals to a wider number of denizens at varying scales, which are sensitive in more fine-grained ways to the properties of those denizens.

This notion of "peculiarity" is a good candidate for understanding when history matters. In circumstances of high peculiarity, explanations are likely to have to appeal to a wider range of particular historical factors, whereas in low peculiarity contexts, significantly less detailed explanations, appealing to broader generalizations, will likely be satisfactory. If Professor Ichthyosaurus is right, the scales and entities of interest to Earth scientists and

palaeontologists are *highly unpeculiar*. If Prof. Ichthy is right, the scales and entities of interest to Earth scientists and palaeontologists are *highly peculiar*.

What is the relationship between peculiarity and loss? Peculiarity is a partly modal relational property concerning persistence and origination, while loss concerns the actual existence of a denizen (characterized as a kind or an individual). If the past is highly peculiar, we should expect it to contain highly complex, local dynamics due to the variety of generating and maintaining processes. In such a world, loss might be rampant, due to the fragility of those conditions. Not necessarily, however: the world might have a lot of contingent origination, but much persistence. That world would be peculiar but not contain much loss. Conversely, a world full of loss is likely highly peculiar. So, if we think the world peculiar, we should potentially expect loss, but it is not guaranteed.

I understand peculiarity as a graded and comparative notion. As with loss, claims about peculiarity rely on a particular characterization of the target at hand. For instance, understood as the historical kind *macroevolutionary radiation*, the Cambrian explosion might be relatively unpeculiar: arising due to the normal dynamics of that kind of event. Understood in terms of the specific adaptations and lineages that arose during the Cambrian, the radiation might be highly peculiar, taking the shape it does in virtue of a set of particular local features. But, again, as with loss, this doesn't make peculiarity a mere feature we project onto the world. With a fixed characterization, I take it, there is a fact of the matter about peculiarity. Peculiarity is also comparative: at least for all of the denizens that palaeobiologists and their ilk are concerned with, everything will be peculiar at least to some extent. But some are more peculiar: they have modal profiles and actual properties that are more sensitive to various factors than others.

So, when I say "peculiar world" going forward, I mean *highly peculiar*.

There could be other reasons for history to matter or not. For instance, we might take the *shit happens* view of history: events in the past do not follow any regular patterns, it's just one thing after the other. Or we might think that history unfolds toward inevitable, perhaps predetermined outcomes, as seen in "progressivist" views of evolution. Fortunately, for my purposes, these can be safely dismissed. We clearly have some grip on regularities in history, and the thought that history progresses "toward" anything requires teleological commitments that I'm very happy not to incur. A quite different approach to denying history's importance appeals to epistemic pessimism: we just don't have sufficient evidence to tell a story about peculiarity at geological scales.

On Secord's reading of Lyell, this is just what we get: Lyell "claimed that *any* kind of global narrative would prove impossible to reconstruct, because too much of the record had been lost [*italics in original*]" (Lyell 1990, 151). Secord's Lyell refuses to commit to a progressivist or steady-state model of life because he thinks the fossil record is too incomplete to support either. No doubt there are cases where pessimism is the right attitude to take about historical knowledge, but no doubt as well, there are cases where optimism is a better bet (Currie 2018a, chapter 12).

So, history matters to the extent that some target is peculiar. Professor Ichthyosaurus takes much of interest to be relatively unpeculiar; Prof. Ichthy takes much of interest to be comparatively peculiar. What makes this a *metaphysical* disagreement? "Metaphysics" is an infamously ambiguous term, but we can make some progress with contrasts. Metaphysical claims are not epistemic but ontic: that is, they don't concern what we know, but what the world is like (or could be like). Metaphysical claims concern the general rather than the particular: that is, they don't make claims about some instances, but systematically. Metaphysical claims concern fundamentality, not superficial appearances: they tell us about the structures underlying the world as it appears to us. This doesn't necessarily make big-picture metaphysical views and empirical science independent, so long as one commits to there being, as John Dupre has put it, "a dynamic interdependence between such general views and the specificities of scientific opinion" (Dupre 2023, 2).

Thus, in claiming that the past at the scales of interest to palaeontologists and Earth scientists is peculiar, I'm saying that the past itself is structured in a peculiar manner. Our fictional ichthyosaurs have a metaphysical disagreement: the two have differing views on the structure of the past at palaeobiological scales. How might we make some headway in distinguishing them? To answer that, we need to consider approaches to the metaphysics of science.

2. Metaphysics and Science

Metaphysics and the philosophy of science have a vexed history, far too dull to rehearse here. Needless to say, how features of science might inform metaphysical questions is an open and much-disputed question (Chakravartty 2010; Hawley 2016). Given my arguments in this book thus far, I'm interested in exploring how the strategies palaeobiologists employ in the face of loss—strategic perspectivalism—might provide inroads into the metaphysical nature of the deep past. Specifically, I'm interested in whether they can

underwrite an argument for the peculiarity of the palaeobiological past. To get there, I think it is helpful to situate my approach, which shifts from practice to metaphysics, with others in the metaphysics of science. As such, I'll draw a set of distinctions regarding the relationship between metaphysics and science, zooming in on what I'm interested in. In the next section, I'll make that argument, but also zoom out once more and suggest that a peculiar history can be defended from multiple methodological standpoints.

We might call the most common approach in analytic metaphysics *minimal naturalism*. At base, our metaphysical views should be consistent with the results of our best science. If your view of the metaphysics of time, for instance, requires an absolute present, then this must be reconciled with the lack of absolute simultaneity in relativity theory. This approach is often coupled with the idea that the methods of metaphysics are in some important respect continuous with the methods of science. A popular approach, for instance, argues that both appeal to various theoretical virtues in forming inferences to the best explanation (Swoyer 1999; Paul 2012).

Philosophers of science have denied both the legitimacy of appeals to mere consistency with science and the idea that science and metaphysics are methodologically continuous (Ladyman and Ross 2007; Novick 2017). What distinguishes a *metaphysician of science* from a bog-standard metaphysical naturalist is, as Kerry McKenzie puts it, that the metaphysician of science “does not spin her view out of air but rather develops it in close consultation with the best science of her time” (McKenzie 2025). That is, the views of a metaphysician of science are in some sense inferred from interpretations of scientific results or scientific practices. But which results and which practices?

McKenzie argues that the relationship between science and metaphysics is best understood in terms of falsification. That is, metaphysical ideas are something like theses about the (general, fundamental) nature of the world, and these theses might clash with scientific discoveries. I take it that on this approach, some minimalist commitment to scientific realism is required: we take scientific theories to provide at least an approximately correct view of the world. Thus, a sustained interpretation of such theories is required. These theories, the idea runs, disallow some metaphysical views while accommodating others. It may be that accounts of time that posit an absolute present simply do not align with the best theories of physics, yet those physical theories may still leave open—underdetermine—other mutually exclusive metaphysical options (Ladyman 2012). As such, McKenzie argues, the metaphysician of science should be less in the business of arguing for or against various

theoretical virtues—that is, constructing inferences to the best explanation from features of theory to the world—and instead aim at the falsification (or near-falsification) of metaphysical hypotheses. They should seek inconsistencies between our best scientific theories and the world.

But scientific theories are not the only potential sources for informing a metaphysics of science. Some philosophers defend approaches that emphasize scientific “practices,” which might include theorizing but are more inclusive (Bausman, Baxter, and Lean 2023; Kaiser and Suarez 2025). As Marie Kaiser and Javier Suarez put it recently:

Metaphysicians of scientific practice typically consider the results or products of scientific investigation (e.g., scientific theories or explanations), as well as the scientific activities that lead to these results, that is, the different epistemic practices (e.g., measurement practices, reasoning practices, explanatory practices, individuation practices) that characterize a scientific field. Hence, this method builds metaphysical claims on empirical information from scientific practice (results, products) and empirical information about scientific practice (activities, practices). (Kaiser and Suarez 2025, 309–310)

Notice the distinction drawn between scientific *products* and scientific *processes*. The former are outputs of scientific work: theories, evidential arguments, datasets, and so on. This includes Butterfield’s narrative trajectory from the “boring billion” to the metazoan-filled, well-oxygenated Mesozoic world is a product, Hoekzema et al.’s model of *Dickinsonia* development via deltoidal addition, and Mitchell et al.’s claim that Avalonian ecosystems are spatially neutral. The latter involves what scientists do to generate those products. As we’ve seen, historical scientists construct and deploy various models and experimental apparatuses (they artifice), use phylogenies and other narrative devices to link past denizens to each other and the present (they graft), and link together various instances of the phenomena they’re interested in (they anchor). More generally, they iteratively apply strategic perspectivalism. These are the practices I’ve focused on in this book, but of course, there are many more. Consider, for instance, the discovery, extraction, preparation, storage, archiving, and analysis of fossils.

I think that focusing on scientific products leads to much continuity between the methods of the metaphysics of science discussed by McKenzie on

the one hand and Kaiser and Suarez on the other. Both require a thorough understanding of the product and asking about its likely truth. At base, we are inferring from how scientists represent the world to how the world is. Butterfield's explanation is highly peculiar, so at least a little inductive weight is given to the idea that the world he's in the business of explaining is a peculiar one.

There are at least three challenges to metaphysical views that proceed from analyses of scientific products. First, why think that our "best" science aligns with metaphysical truths? "Best" need not accord with "true," and we shouldn't accept the best of a bad lot. Second, what scientific theories in fact say or mean is a matter of philosophical interpretation—and it isn't obvious that we can identify a "correct" interpretation (Kinzel 2015; Currie 2025). Second, scientific claims are not on their face general, structural claims about the fundamental nature of reality, so we must generalize and systematize them to generate metaphysical claims, even if we've nailed the interpretation. Third, plausibly scientific theories underdetermine metaphysical theories (McKenzie 2025). We can't, then, take scientific theories to simply mirror metaphysics.

There is nonetheless space for this approach to underwrite peculiarism about the past. Research into the deep past has revealed much that is peculiar: often shifting from simple to complex narratives, suggesting highly sensitive and interdependent dynamics (Currie 2014; Dresow 2021). We've seen plenty of examples throughout the cases in the first half of this book.

However, in the context of my approach in this monograph, I am not well placed to make an argument from scientific products. Astute readers will have noticed that I haven't really committed to any particular scientific claim. I'm not in the business of arguing that, say, Avalonian ecosystems were neutral, that pizza disks are not a true taxon, that *Oryctodromeus* burrowed, or even that dinosaurs felt pain. My arguments have not turned on the truth or otherwise of token propositions about the deep past, but concern the strategies scientists adopt when confronted with loss or its possibility. As such, I'm much better placed to speak to practices than products.

How would understanding scientific practices be a route into metaphysics? The approach must be less direct than a product-driven one. After all, at least sometimes, scientific products are supposed to be about the world, and the inductive metaphysician takes those results and says, "Well, they're true, and thus they partly reveal the metaphysical structure of the world." Practices aren't like this. They do not represent the world in the same way.

However, scientific practices are not independent of the world. Far from it: they are rather part of, and situated within, the world. And they are directed toward discovering features of that world. As such, practices might serve as an indirect way into metaphysics. As Ken Waters has put it: “The best strategy underlying scientific metaphysics, I suggest, is based on the idea that the world provides constraints on scientific inquiry, and that philosophers can inform metaphysics by investigating these constraints” (Waters 2017, 83).

This relies on the idea that not just any scientific approach will succeed for some investigative target. Sometimes, at least the world will resist our attempts to make empirical or theoretical inroads. In the best case, we hope, the practices that scientists adopt are *adapted to* the system they are interested in understanding. As I’ll expand below, and following Will Bausman (2023), something analogous to biological adaptation might be applicable. That is, much as we might infer from the platypus having webbed feet that it evolved in aquatic environments, perhaps we can infer from scientists adopting particular practices that there is a world with particular properties they seek to understand.

Ian Hacking’s argument for entity realism is potentially an example of this kind of approach (Hacking 1983; Miller 2016). Hacking argued that some unobservable entities posited in physics shouldn’t be understood as theoretical postulates—entities inferred from the results of experimental interventions—but instead understood as entities that are *used* to perform experimental interventions. The classic example involves incorporating electrons into an experimental apparatus built to detect properties of other non-observable particles. For Hacking, the existence of those entities, or at least entities with the properties made use of in the experiments, is in some sense presupposed by the experimental practice. Thus, the practice of these experimental scientists is revelatory of the world they examine, in virtue of how the world is incorporated into and used in those practices. We’ll consider how exactly that inference might look in the next section, but here Hacking is useful as an example of a potential inference from practice to metaphysics.

Much work along this vein has focused on scientific concepts. Waters, for instance, argues that the flexible concepts of “gene” developed in experimental biology are made sense of in light of a biological world that lacks a general structure. I suspect that much concept-focused work in the philosophy of biology could be understood in something like these terms. If scientific concepts are developed in dialogue with the systems they are intended to make sense of, that is, if scientific concepts are shaped by what we could call “ontic

interaction,” then the concepts scientists employ will reflect that interaction to some extent. For instance, biologists make use of many differing concepts of “species,” some appealing to phylogeny, others to interbreeding, yet others to genetic relatedness, and so on. We might think that this conceptual pluralism is a reaction to the complexity of the world biologists encounter. Something similar could be said of work focusing on classification in science: the plurality of classifications maps onto a complex world (Dupre 1993; Ereshefsky and Reydon 2015; Kendig 2015).

However, in this book, I have not had much to say about concepts in historical science. Although a very loose notion of kinds—token-historical, type-historical, and ahistorical—is useful for understanding characterization, my analysis has not been focused on concepts and their deployment. Rather, my focus is on the strategies these scientists adopt in tackling loss. So, we want to focus on strategies or methods. We might argue that methods, too, are shaped by ontic interaction.

Kirsten Walsh and I have developed an account of what we call *ontic-driven explanations of scientific method* (Currie and Walsh 2018). At least sometimes, when we try to make sense of why scientists adopt the methods they do, we point to particular features of the system they are trying to understand. We explain the notion as such: “An explanation of scientific method is ontic-driven when the explanatory target and the relevant contrasts diverge in virtue of properties of the systems investigated” (Currie and Walsh 2018, 122).

Explanations are often contrastive; we want to know why one thing happened rather than another. In the cases Walsh and I are interested in, we want to know why some group of scientists adopted one method instead of another. The explanation is *ontic-driven* when it appeals to features of the system the scientists are trying to understand: they adopted that method *because* they’re trying to understand a system with particular properties that the method is attuned to.

So, ontic-driven explanations target scientific methods. We might ask, say, why scientific community *A* adopts lab-based experimental methods while community *B* uses observational methods. An answer might appeal to many not mutually exclusive factors. These include various social and economic factors: *A* might have the means to develop, house, and maintain expensive experimental apparatuses, while *B* might not. They could also be due to differing epistemic traditions: *A* might value intervention, while *B* values examining their target “in the wild.” It could be due to the aims and purposes of the science: *A* might be primarily interested in developing technological

applications, while *B* might be interested in uncovering natural patterns. At least sometimes, an important additional factor is *the kinds of systems A and B are interested in*. *B* might be interested in complex, large-scale systems. Community *B* might be targeting, say, meteorology, planetary dynamics, or ecosystems. Meanwhile, community *A* might be interested in more localized mechanistic systems, perhaps developmental systems, or chemical reactions. As such, we might agree that *A*'s experimental approach is due to the kinds of investigations that the system affords, while *B*'s more observational work is adapted to what their target will allow. When we say *A* and *B* differ in their method because of differences in the systems they target, we are providing an ontic-driven explanation. Ontic-driven explanations can also be brought to bear on the success of various methods. A scientific community might have more success with one method rather than another because the former is better suited to their explanatory target.

When we provide an ontic-driven explanation of some scientific method, it seems to me, we also provide *prima facie* reason to accept the ontology cited in the explanation. If we don't think what we're saying about the world is true, then presumably our explanation of why those scientists use those methods fails. This is highly defeasible. We might just be wrong about the method. It might be that the explanation is wrong, that the method is adopted primarily for some non-ontic (or a different ontic) reason. It could be due more to social or epistemic factors, say, or perhaps the ontology isn't playing the explanatory role we thought. It might also be that the method is inappropriate: that scientists are making some kind of mistake in adopting it. Katherine Valde, for instance, has recently argued that some metaphysical commitments in cancer research can lead to egregious methodological choices, instead suggesting that agnosticism about underlying metaphysics would lead to better methodologies (Valde 2024). But a sufficiently well-supported, ontic-driven explanation, I take it, provides some route into the ontology the method is adapted to.

At this point, the shape of the argument I'm interested in exploring (perhaps even making!) is emerging. Some of the strategies historical scientists adopt, particularly in light of loss or putative loss, are adapted to a peculiar world. And this gives us some reason to think the world is, in fact, peculiar.

3. The Peculiarity of History

I've thus far provided an account of "history mattering" in terms of peculiarity, and zeroed in on a set of pathways from scientific practice to metaphysics: namely, that an ontic-driven explanation of method underwrites a metaphysics committed to that ontology. It's now time to see what headway we can make in applying that approach to peculiarity. Here's a sketch of how that argument might go:

1. Historical scientists are *strategic perspectivalists*.
2. Ontic-driven explanation: strategic perspectivalism is best understood as a response to the peculiarity of the entities and scales historical scientists are interested in.
3. Bridging principle: ontic-driven explanations underwrite metaphysical claims.
4. Therefore, the scales and entities of interest to historical scientists are peculiar.

The first premise was defended through chapters 3, 4, and 5, at least pertaining to the science of lost worlds. In response to loss or potential loss, historical scientists adopt perspectives that facilitate artificing, grafting, and anchoring, perspectives that are enabled by various tools that often partially overlap, thus allowing for iterative progression. My task here, then, is to consider premises 2 and 3. Is there an ontic-driven explanation to be had of strategic perspectivalism and, indeed, is there some principle that might take us from that explanation to the past in fact being peculiar? I suspect so.

Let's remind ourselves of strategic perspectivalism's essential features. "Strategic perspectivalism" is not the same kind of view as "philosophical perspectivalism"—it isn't about the nature of knowledge itself—rather, it is an account of the strategies scientists adopt in tackling loss. This strategy involves (1) characterization: using a perspectival tool to pick out a set of denizens and their properties (and potentially a set of standards and expectations associated with that tool and characterization). This characterization forms the basis of some combination of (2) artificing, grafting, or anchoring. That is, the characterization affords various investigative routes into understanding those denizens. Finally, (3) recontextualization: results from investigations

are brought into dialogue with features of the past world targeted (and other perspectives). I've further emphasized the *bespoke* nature of the investigations: in artificing, for instance, the details of the model are often specified to highly particular features of the past world or its traces. I've also emphasized the *pluripotent* nature of the investigations: tools, results, and interpretations often play multiple epistemic roles, informing us about the target, aiding in the construction of models of regularities, and iterating between various perspectival tools.

To what kind of world would this method be best applicable? A highly peculiar one. To see this, consider what an investigation of *unpeculiar* targets might be like. An unpeculiar target is one that is insensitive to generative and maintaining processes; that is, it maintains its properties and persists across a wide range of background conditions, processes, and possible disruptions. For such a target, knowledge of the dynamics of its properties themselves, and the relatively few conditions of its persistence, would be sufficient for explanation and prediction. An investigation into such a system would no doubt involve characterization, but relatively few perspectives would be required to make sense of it. Considering investigations of the past, in an unpeculiar world, a narrow uniformitarian method would work highly reliably. One would be able to investigate current processes and simply project them backward, scaling up to the deep past's horizons.

But a strict uniformitarian approach doesn't work. As Meghan Page (2021) has emphasized, what appear to be stable regularities often only stretch so far into the past. Some examples include the following: assuming niche differentiation in spatial ecologies when interpreting Avalonian rangeomorph-dominated systems would yield the wrong results; the size limits on today's obligate osmotrophs do not necessarily apply to rangeomorphs; and the relationship between the biosphere and atmospheric carbon on Earth is potentially unique to our planet and has been stable only for the last few hundred million years.

In a highly peculiar world, the existence and persistence of denizens will be sensitive to the properties of various *other* denizens. And those *other* denizens will themselves have various sensitivities. The world would both be highly interdependent and highly fragile. The well-mixed oxygenated oceans of the Mesozoic and Cenozoic have been stable for over 400 million years. But if Butterfield is right, that stability is sensitive to processes of bioturbation, the biological pump, large populations of krill and their allies, and so on. Rangeomorphs appeared to play a crucial role in Ediacaran ecologies for

something like 30 million years. But their dominance was fragile: their sessility made them unfit in ecologies sporting more dynamic metazoans. As I discussed in previous work (Currie 2019a), from the dawn of the Jurassic until the “Mid-Cretaceous Terrestrial Revolution” around 100 million years ago, the world’s macroflora and fauna could be characterized in terms of sauro-pod dinosaurs consuming gymnosperms. With the advent of angiosperms, we’ve seen millions of years of flowers and pollinating insects. These patterns are highly stable (often due to the interdependence between their constituent systems) but, nonetheless, are fragile due to reliance on various stabilizing processes. In other words, they are peculiar.

Peculiar worlds, due to their multiple fragilities and interdependencies, must be understood from multiple scales across multiple entities and processes. Many denizens have crucial similarities and continuities between them, but those similarities are typically only partial. More generally occurring regularities will only be capturable by highly abstract models. In trying to understand a peculiar world, scientific tools must be bespoke to capture the specific details and thus be epistemically relevant. Investigations of peculiar worlds will be open-ended as they navigate the various partial similarities that history has generated, and varying tools will be required to characterize the varying denizens and their interdependency. In other words, a peculiar world is best tackled via strategic perspectivalism. Alisa Bokulich (2021a) has argued that Earth scientists navigate among various concrete and conceptual models to investigate phenomena that reach across scales. Models are used to make sense of both hierarchical relations (identifying “the functionally relevant scales given the application of interest, as well as the constraints coming from the higher levels of the hierarchy” (Bokulich 2021a, 14179) and threshold effects “identifying the relevant dynamic thresholds in the system where there are qualitatively discontinuous changes in the system’s behavior”) (Bokulich 2021a, 14179). That is, different tools are used to examine both diachronic dependencies, such as that between the evolution of multicellularity and a well-oxygenated ocean (if Butterfield is right) or fossils and their living progenitors; and synchronic dependencies, such as between an organism’s traits and its capacities. In the parlance of this book, the variety and complexity of generating and maintaining processes lead to scale-relative complexity, which requires multiple perspectives to make sense of. The world’s peculiarity—or at least its complexity—demands strategic perspectival investigation.

Consider once more Butterfield’s arguments concerning the emergence of complex metazoans. Both his narrative and its supporting evidence require

switching between a variety of different scales and denizens. We must consider: models and data concerning global biochemical interchange, the optimality or otherwise of various morphological adaptations (often in their biochemical context), the structure of benthic ecologies, niche construction as applied to oxygen environments, and the evolvability of metazoan innovations. We bounce between the global, the ecological, the physiological, and the evolutionary. This switching between and combining of perspectives is, I suggest, necessitated by the past's peculiarity. If the past were *less* peculiar, then we wouldn't need to pick out so many differing sensitivities to explain phenomena.

So, we have at least a putative, ontic-driven explanation of strategic perspectivalism. Historical scientists are strategic perspectivalists due, in a significant way, to the peculiarity of the past worlds they aim to understand. So, how might we link this with the claim that the past world is, in fact, peculiar: what bridging principles between practice and metaphysics are suitable here?

Boaz Miller (2016) has dissected Hacking's argument for entity realism, providing a set of possible interpretations that I'll canvass before drawing on William Bausman's (2024) idea that practice might be adapted to the world via a process analogous to natural selection.

Inference to the best explanation is the most discussed bridge between science and metaphysics. Following Peter Lipton's turn of phrase, an inference to the best explanation takes us from an explanation's loveliness to its likeliness (Lipton 2000). Beyond an explanation's empirical adequacy and its capturing the phenomenon we're interested in, we might also think that certain other properties—elegance, simplicity, and so on—could lead us to think it more likely to be true. Like some philosophers of science, I'm sceptical that inference to the best explanation, as traditional metaphysicians appeal to it, is continuous with the epistemology of science (Novick 2017). Regardless, I think there's a disanalogy between what I'm doing and what philosophers who attempt to argue from scientific results are doing. I'm arguing from a philosophical explanation of scientific practice to a claim about the world: I don't take myself to be making anything like a scientific argument. Rather, I'm making a philosophical argument based on a philosophical explanation of scientific practice. This doesn't commit me to saying that philosophical and scientific arguments aren't in some sense continuous. However, I don't think that appeal to inference to the best explanation (IBE), insofar as it occurs in science, gives philosophical appeals to IBE any special license.

Shifting from IBE specifically, what would we need to establish to connect our ontic-driven explanation to a peculiar ontology? Ideally, such an argument would establish the truth of a conditional along the lines of the following:

If the past were not peculiar, historical scientists would not be strategic perspectivalists.

However, a strict conditional relationship is too much to hope for, as we'll see. But there is reason to commit to something weaker.

When I considered what the science of a non-peculiar past might be like, I suggested that it would be considerably more uniformitarian. Scientists would seek to discover the basic dynamics by which the past unfolds, and these dynamics would not require bespoke models to capture them (see Walsh and Currie's discussion of Newton's astronomical work, 2018). At least sometimes, then, a lack of peculiarity underwrites a different methodological strategy for scientists.

Thus, we might argue that the best explanation of historical scientists adopting strategic perspectives is that the world is peculiar. To apply this approach, we'd both need to give reason to think it is the best explanation, and reason to think it is, as it were, sufficiently best. Bausman's account is, I think, promising here: "Successful scientific practices are adapted to features of the world for purposes as organisms are adapted to features of their environment for functions" (Bausman 2023, 28). If the sciences follow a process analogous to adaptation while investigating their target systems, then we can make sense of the above conditional: over time, ontic interaction has shaped the methods and practices of historical scientists, thereby fitting them to their target system.

While ontic-driven explanations do not exclude other possible reasons for scientists to adopt a method—explanations can cite multiple factors—there might be explanations that exclude or minimize the importance of peculiarity, thus undermining the argument. In the last section's discussion of ontic-driven explanation, I named a set of non-ontic factors, such as social, economic, or ideological factors, that could explain strategic perspectivalism. Here I'll discuss some ontic alternatives.

A set of potentially pressing alternatives is also ontic-driven. It might not be peculiarity but erasure that underwrites strategic perspectivalism. Historical scientists often face "epistemic scarcity" (Currie 2021a). They

tackle a gappy record, with often hard-to-manage specimens and data, and often lack the background theory required to “fill” that record’s gaps. In response, the explanation goes, they adopt strategic perspectival methods. As peculiarity underwrites loss, this explanation amounts to claiming that most loss is only apparent and that a lack of uniformity is due to the degradation of traces. In chapter 1, we saw Lyell argue something like this: the apparently “progressive” pattern in the fossil record, taking us from fish-dominated to reptile-dominated to mammal-dominated macrofauna, is due to differential preservation (and is moreover a highly selective reading of the record). We lack sufficient evidence to conclude it is a biological signal.

On this alternative ontic-driven explanation, historical scientists are strategically perspectival because of a lack of traces. This amounts, perhaps, to saying that information preservation is peculiar, but the sources of that information—the non-trace denizens in the past—are not, or at least that we should remain agnostic on that score.

I have an answer, I think, to this alternative. Many of the troubling cases of loss or putative loss we’ve discussed have been in relatively trace-rich contexts. Take, for instance, Mitchell et al.’s examination of rangeomorph ecosystems. The fossil beds of Mistaken Point are remarkable: there is good reason to think that volcanoclastic processes have preserved the spatial positions of the rangeomorphs *in situ*. However, loss and potential loss nonetheless present major epistemic challenges in this trace-rich context, and Mitchell et al. adopt a paradigmatically strategic perspectival strategy in light of this. David Sepkoski (2019) and I (2018a) have argued that the capacity of historical scientists to “re-read” the fossil record, distinguishing between biological and geological signals and thus partly overcoming a lack of traces, is remarkable. No doubt there are many cases where we should be pessimistic about our knowledge of the past because of epistemic scarcity, but this pessimism should only go so far. To be sure, strategic perspectivalism is in part a response to a lack of traces, but if that were the only driver, we would see historical science become significantly more uniformitarian as records became less gappy or better understood. And I don’t see evidence of this.

A further set of alternative ontic-driven explanations might appeal to non-peculiar ways the world might be that strategic perspectivalism is a good response to. One alternative is *complexity*.⁶ The world might be highly

6 Thanks to Alisa Bokulich for pushing me on this alternative.

complex and heterogeneous, that is, made up of many different components with a wide variety of arrangements and relations, without being peculiar. These complex properties and relations could have few dependencies between generating and maintaining processes; they would be complex, but nonetheless not peculiar. Plausibly, under those conditions, successful scientists would also adopt strategic perspectivalism. I consider this a live and alternative hypothesis, although I suspect that such a world would demand explanation in a way that a peculiar world does not. As I'll argue in the next section, a peculiar past *makes sense of* a complex and heterogeneous present. At the very least, a complex-but-unpeculiar world would require some explanation.

So, although these alternatives undermine a strict conditional relationship between peculiarity and strategic perspectivalism, I do think there are reasonable grounds to take at least some of the practices of historical science to be adapted to a peculiar world.

Miller discusses two alternative readings of Hacking's argument that are worth mentioning. He calls these the "indispensable" and "transcendental" arguments. Both routes more or less argue that a particular ontological posit is required for a scientific practice. In Hacking's case, if there were no entities like electrons, then apparatuses like electron microscopes wouldn't make sense: without committing to the existence of electrons, the practice is rendered incoherent. Similarly, the notion of *electron* (or at least some entities with electron-like-properties) is indispensable to the use of electron microscopes. To apply these approaches to the case at hand, we'd need to establish that the coherence of strategic perspectivalism turns on there being a peculiar past. I think there is potentially *something* to this kind of argument: historical scientists tackle the past with the expectation and understanding of peculiarity. And understanding peculiarity makes sense of why they use the strategies that they do. But I don't think anything as strong as indispensability can be established. As we've seen, there are plenty of other contenders.

Overall, linking peculiarity and strategic perspectivalism is best done by arguing that the method is adapted to a peculiar past. I imagine that scientists who begin investigation taking the world as non-peculiar—and thus adopting a uniformitarian approach—would quickly find their investigation either stalling or gradually becoming increasingly perspectival. Although, as we've seen, historical science is often very concerned with understanding the present and linking it to the past in various ways, attempts to uncover non-bespoke general regularities have typically failed. There are at least two ways one might try to strengthen this argument. First, one might take it as a

serious historical claim and go looking for cases across the history of palaeontology to see whether such a pattern presents itself (I suspect claims that historical explanations increase in complexity over time could be co-opted here: see Currie 2014; Dresow 2021). Second, one might take it as a thought experiment aiming to pump and explore our intuitions about the connection between method and the kind of system we're targeting.

Bausman distinguishes between inferring from practice, the metaphysics of a particular domain, and inferring across domains. Waters' argument that there is no general structure, for instance, first involves establishing that the practices of geneticists are not adapted to a generally structured world, and then requires universalizing that claim. I am interested in the former but not the latter inference. That is, I am interested in the kinds of scales investigated by palaeobiologists and Earth scientists (already a pretty rich domain!). There might be ways of expanding this—maybe peculiarity in one domain entails peculiarity in others—but I think such arguments require care. Fragility and sensitivity in one domain need not guarantee the same at “lower levels” of description (maybe those atoms behave highly regularly), but also not at “higher” levels (order can arise out of chaos; see Strevens 2009).

Of course, this doesn't mean that I am not making various kinds of inductions. I've hardly conducted a survey of practices across the Earth sciences to show that strategic perspectivalism is the norm. Although I'm being metaphysically fairly cautious, I'm being rather reckless in my claims about methodology.

None of my arguments would force a really determined detractor to agree with me on pain of irrationality. But I don't think metaphysical arguments need to be “knock-down” in that sense. Nor need we rely on a single argument for a conclusion. Nancy Cartwright has decried the use of “long, skinny” deductive arguments on the basis of their fragility—lose one premise, and the whole shebang is up in logical smoke (Cartwright 2015). Instead, she recommends arguments that are perhaps less logically impeccable, but are significantly more robust. I see the argument here as following the latter approach. (1) That historical scientists have discovered peculiarity in the past over and over again, (2) that an ontic driven explanation of their method appeals to peculiarity, and (3) that their method is plausibly adapted to a peculiar world, are all imperfect arguments pointing toward the same metaphysical conclusion: at the scales that interest palaeobiologists and Earth scientists at the very least, the past is peculiar. That is, history matters.

4. From Peculiar to Dappled

There is a weighty pause whilst Prof. Ichthy awkwardly shuffles on their lower flippers, mentally corralling their own paragraphs. Scratching their lower jaw, they respond to the *Allosaurus*.

“I don’t think the world is pure chaos: history is highly structured in various ways. Broadfoot, we have good reason to believe, lived in shallow fresh water. We think they used their echolocation to detect prey, the best bet being small crustaceans. They existed, then, in an environment that they were adapted to, and were linked with other organisms through various trophic interactions. I don’t know if there was anything we might call a ‘harmony’ in that ecosystem, but there was at least some kind of dynamic stability. And, I would say, that stability itself had a history: we can make sense of why the mammals arose to such prominence and perhaps make sense of why they fell.

“History, then, is neither chaotic nor ordered, but a bit of both. And I like to think that understanding ourselves in light of that history—how the order and, yes, the chaos arises from complex historical processes—is a way of making sense of our lives. I guess that’s likely as philosophical as I can get, I’m afraid.”

Prof. Ichthy and the *Allosaurus* both half-shrug as if to signal a conversational cease-fire.

Everyone agrees that the representational power of contemporary science provides us with, at best, a patchwork. Insofar as science provides an empirical grip on reality, it does so impartially, imperfectly, requiring various fudges and kludges to scale the gaps between perspectives. Science is disunified in that sense. Descriptive claims concerning science’s disunity have been traditionally drawn on to undermine unificationist and reductionist conceptions of science. No doubt, seeking underwriting mechanisms for various processes—understanding the “macro” through the “micro”—is often a successful scientific strategy (Machamer, Darden, and Craver 2000), but it doesn’t often lead to more general unification, and nor does it often lead to applications that can be applied “off the shelf,” as it were, to domains beyond those that they were developed to make sense of.

A common way of making sense of disunity appeals to the limits of human knowledge. A philosophical perspectivist, for instance, might infer from the claim that all knowledge is from a perspective—that is, requiring both a characterization and a standpoint—to the idea that all knowledge is

partial. Further, we might point to the limits of representation generally to make sense of science's disunity (Mitchell 2002). Our theoretical understanding is necessarily partial, so no wonder the picture of the world provided by science is a patchwork.

One needn't lean on representationalist accounts of scientific knowledge nor on theoretical understanding to provide an epistemic explanation of science's disunity. We might also appeal to the varying aims of science and how these aims shape scientific research into differing forms. Alan Love's notion of a *problem agenda* (Love 2008), for instance, conceptualizes scientific research as driven by particular questions or "problems" (say, understanding rangeomorph feeding). Tackling these questions sets an agenda, which requires marshalling multidisciplinary evidence and structuring that evidence into various roles (traces that inform us about morphology, geometric models, studies of osmotrophy, etc.). As differing questions—differing aims—require differing complexes of disciplines and evidence, we should expect science to take the form of a complex web of disunities and integration (see also Brigandt 2013).

Finally, in chapters 6 and 7, I defended a form of *artifactualism* about the imagination of historical scientists and the modal knowledge they provide. An artifactualist explains the justification of modal knowledge by appeal to the dispositions of the artifacts scientists construct. We might generalize this idea, claiming that scientific knowledge is *fundamentally constrained by its medium*, that is, the physical thing that bears or instantiates the dispositional properties that constitute scientific knowledge. Different media instantiate different modal properties, and this necessitates a diversity of media. As these media exemplify slightly different modal spaces (see Currie 2020), knowledge is always partial due to its medium-boundedness. I have some sympathy for this idea, particularly if coupled with a know-how account of understanding. If artifactualism about scientific knowledge generally could be defended (and doing so is well beyond the scope of this book! Maybe the next one . . .), it seems to follow that this knowledge is of necessity a patchwork.

In the metaphysical spirit of this chapter, we might go beyond these epistemic explanations of science's disunity to construct an ontic-driven explanation. It isn't only constraints on knowledge, nor the nature of scientific knowledge, that underwrite disunity. Additionally, science is disunified because *the world is disunified*. John Dupre (1993), Nancy Cartwright (1999), and (as we've seen) Ken Waters (2017) have argued along these lines. Cartwright, focused on economics and physics, argues that the theories of

those sciences, despite the best efforts of the scientists working on them, fail in their goals of generality. And this failure is explained by the world being “dappled”: there is one material world, which is to say whatever there is, is spatio-temporally connected; however, the various relations between those things cannot be systematically described.

The sciences are each tied, both in application and confirmation, to the same material world; their language is the shared language of space-time events. But beyond that there is no system, no fixed relations among them . . . Their boundaries are flexible: they can be expanded or contracted; they can even come to cover some of the same territory. But they undoubtedly have boundaries. There is no universal cover of law. (Cartwright 1999, 6)

Dupre focuses primarily on how biologists go about carving up the world: the categories they employ are often not intertranslatable, cross-cut one another, and serve different purposes. For Dupre, there is no basis for metaphysically privileging one way of carving over another. When we find order, “There is no reason to suppose that whatever uniformities apply to a particular ordered domain will apply to every such pocket of order” (Dupre 1993, 4). As such, we should commit to what he calls *promiscuous realism*: the world is constituted by a wide variety of complex, overlapping kinds, all of which deserve to be taken metaphysically seriously.

Certainly, Cartwright and, to an extent, Dupre, provide product-first metaphysical arguments, and as such face some of the same challenges I described above for that route from science to metaphysics. Waters, as we’ve seen, argues along similar lines to mine: from the practices of scientists. What unifies these views, however, is their *synchronic* nature. They paint the world as a complex patchwork, but have little to say about *why* the world might be such a patchwork. Let’s quickly point to some synchronic arguments for disunity, before turning to a diachronic account.

Taking the metaphysical view of the causal closure of the physical as our target, we can discern several synchronic arguments for the world’s disunity. Causal closure is the idea that the causal powers of whatever the smallest bits in the world are—at least in principle, the things studied by current or future fundamental physics—exhaust the causal powers of everything else. Arguments for metaphysical disunity often involve denying causal closure

in some way (Dupre 2023). One argument is *from context*: the behaviour of many low-level entities turns not just on their internal properties, but often on highly context-sensitive factors within their environment. These contextual factors are crucial for understanding how messy the world is (Dupre 1993; Havstad 2018; Baxter 2023). Gene expression doesn't just rely on that particular gene, but on the influence from other genes, on the biochemical environment they are in, and on interactions from various other factors in the developmental sequence (Griffiths and Gray 1994). Another argument is from explanatory autonomy: if some explanations at a higher level capture explanatorily ineliminable information that is not discernible from the lower level (Jackson and Pettit 1992; Sterelny 1996; Grantham 1999)—and we read explanation in ontic terms—then either causation doesn't exhaust ontology, or there are “higher-level” causes. Further, there are arguments from reductive failure: there is a long list of entities that were once taken as putative places where the world might bottom out, or that might be the underlying determinant of some system, which, on investigation, have turned out to be significantly messier. The relationship between Mendelian and molecular genetics is perhaps the classic instance of these phenomena (Kitcher 1984).

These synchronic arguments can be supplemented with a diachronic argument, which we might call *the argument from peculiarity*. If we think the past is peculiar, then we should expect the present (and future!) to be dappled and disunified. Let's see why.

A peculiar world is one in which new kinds arise as new conditions for existence emerge. The post-Avalonian evolution of metazoans blessed with movement and mouths represented not only a set of new individual taxa but also enabled new ecological dynamics: predation and scavenging, for instance. These are often highly stable, but also highly peculiar. This should lead to a differential structuring of different parts of the world at different times. Globally, the Earth has housed fairly stable relationships between atmosphere and biosphere for 400 million years. But there are more local variations within that global story, and earlier there were different global stabilities. Peculiarity, then, provides diachronic grounding for a dappled world: the world is disunified *because* the past is peculiar.⁷ I'll illustrate this with a simple ecological case.

⁷ In section 4 of chapter 5, I mentioned Lange's distinction between temporary and eternal but time-dependent laws. One might insist that the right set of eternal but time-dependent

Aotearoa (New Zealand), prior to the arrival of our species and their accompanying troop of placental mammals, boasted a fauna populated by birds and insects. Islands like Aotearoa often have weird animals, and one way of making sense of these weirdos is by appeal to familiar niches being occupied by surprising taxa due to founder effects. As such, Aotearoa is full of birds filling niches typically associated with mammals. One of the most iconic is the flightless ratite moa, which ranged in size from around a meter to the enormous 3.6-meter-tall *Dinoris*. They were highly populous (Gemmell, Schwartz, and Robertson 2004) and took on many ecological roles, from sheep-like grazing to deer-like browsing (Wood et al. 2020). So, various moa species acted something like deer or sheep, fulfilling grazing and browsing roles. Meanwhile, their equally iconic relatives, the kiwis, are birds attempting to act like nocturnal burrowing insectivores, such as badgers. It is sometimes assumed that the introduction of mammals into Aotearoa involved “ecological replacement,” that is, mammals took up the various roles that birds once filled. At an extremely coarse-grained level of analysis, I suppose something like this is right: deer and moa are both browsers after all. This perspective runs into both philosophical and empirical trouble, however.

We might have philosophical worries about this notion of “niche.” As Dennis Walsh puts it: “A (nonmetaphorical) niche is a space, or a recess, into which something (say, a statue) might fit. Like a real niche, an evolutionary niche is a set of properties of an environment, to which an organismal form may fit” (Walsh 2015, 169). However, Walsh points out, this seems to assume that niches somehow exist independently from their occupants, as if Aotearoa was “waiting” for kiwis to evolve. Richard Lewontin ties this independent and concrete conception of niches to adaptationist thinking: “To make the metaphor of adaptation work, environments or ecological niches must exist before the organisms that fill them. Otherwise, environments couldn’t cause organisms to fill those niches. The history of life is then the history of coming into being of new forms that fit more closely into these pre-existing niches” (Lewontin 2001, 63).

I think Walsh and Lewontin’s philosophical skepticism has some merits, although I doubt it applies across all niche-talk among scientists. In the spirit of strategic perspectivalism, we might think that conceiving of niches in coarse-grained ways might gain epistemic purchase in some contexts.

laws, coupled with sufficiently complex initial conditions, would render a non-dappled, law-like world.

Regardless, there are empirical problems with applying ecological replacement to Aotearoa, and this has everything to do with the peculiarity of trophic interactions.

The transition from a bird-rich to a mammal-rich fauna drastically changed some of Aotearoa's trophic dynamics. Here's one example, taken from Atkinson and Greenwood (1989). In this study, they test replacement by seeing if it can make sense of the changes that occurred in Aotearoa when mammals were introduced. "Discussions of the effects of introduced browsing mammals on New Zealand plants and vegetation have sometimes included the suggestion that they are not fundamentally different from those formerly exerted by moas. This suggestion is treated here as a null hypothesis and tested by some comparisons between moas and mammals" (Atkinson and Greenwood 1989, 67).

A fair few of Aotearoa's plants exhibit a form of cryptic mimicry where healthy parts appear to be unhealthy or dead. For instance, the leafless pōhuehue (*Muehlenbeckia ephedoides*) has lower shoots that take on the appearance of dead twigs. This is presumably an adaptation, a surprising kind of mimicry that makes the plant seem less palatable to grazing animals. This makes sense for a bird-rich biota, as birds primarily hunt by eyesight (with kiwis being a notable, mammal-like exception). However, looking dead is very little help if you're trying to avoid being eaten by mammals like sheep, who rely not on their eyes but primarily on their noses. So, the leafless pōhuehue is adapted to the bird world, and sadly maladapted to the mammal world. Atkinson and Greenwood provide many examples of this phenomenon: although mammals and birds both act as grazers and browsers, they do so in quite different ways, and these differences matter.

It is unlikely that moas bit and stripped bark from stems, in the way that deer and goats bark *Pseudopanax* spp. and some podocarps. It is also hard to imagine them ploughing the soil and destroying roots in the way that pigs do even though they may have scratched and torn at the ground with their feet. Moas were presumably not capable of climbing trees, so there would be no parallel with the arboreal feeding of brush-tailed possums. (Atkinson and Greenwood 1989, 89)

A picture emerges: Aotearoa's flora and fauna co-evolved; thus, truths about bird-rich Aotearoa do not carry over into mammal-rich Aotearoa. A peculiar

process—co-evolution—leads to a disunified world. The leafless pōhuehue is adapted towards a bird-filled ecology in which visual mimicry is a good strategy. Much less so in a mammalian ecology. So, the regularity *visual mimicry is a good way of avoiding grazing and browsing animals* held true in Aotearoa prior to the arrival of humans and other mammals, but no longer held true afterward. 1500 years ago, it was true in Aotearoa but false in, say, North America.

One might complain that more coarse-grained conditional regularities nonetheless held true, say, “If grazers rely primarily on vision, then playing dead is a good strategy.” However, this conditional itself captures the results of peculiar historical processes. Peculiarity doesn’t necessarily undermine fairly coarse-grained conditional properties or relations. But note that many of the properties above, e.g., “vision,” “browsing,” and “playing dead,” are themselves the outcomes of fairly peculiar processes, and how they play out—what acts as vision, how one browses, what it takes to play dead—are often highly sensitive to evolutionary histories. Due to the co-evolution of biota, such regularities only gain purchase under specific local circumstances. That is, the peculiarity of the past leads to a dappled picture.

I want to close with two discussions. The first concerns the limits of experimentation, and the second concerns the relationship between what I’ve established here and a more general metaphysical view: processualism.

Much has been made of the importance of experimentation to science (rightly, I think: Weber 2004; Currie and Levy 2019); however, philosophers committed to a dappled or disunified world sometimes point to the limitations of experimental systems, Cartwright most paradigmatically: “Physics in its various branches works in pockets, primarily inside walls: the walls of a laboratory or the casing of a common battery or deep in a large thermos, walls within which the conditions can be arranged just so, to fit the well confirmed and well-established models of the theory, that is, the models that have proved to be dependable and can be relied on to stay that way” (Cartwright 1999, 2).

Similarly, in biology, Jim Griesemer has emphasized the relationship between experimental systems and mathematical models, explaining how the design of experimental systems involves the application of scientific agency toward that system behaving in ways related to mathematical models (Griesemer 1988; Griesemer and Wade 1988). The expected experimental behaviour is not necessarily analogous to behaviour in the natural target: “Expectations about effects stemming from knowledge of causes in the

laboratory system are applied to natural systems through the use of the same data model. To justify using the same data model, the experimental design structure of a laboratory model must be recognized in the natural system” (Griesemer and Wade 1988, 84).

Experiments give the results that they do because, in part, they have been designed in a particular way to “capture” the regularities at hand—they provide “walls,” as Cartwright puts it—so, if they are to apply to natural systems, the differences between the structure of the lab and the world must be understood.

This take on experiments can be understood using peculiarity: established experimental regularities are highly peculiar to human-manufactured conditions. As such, understanding the limits and otherwise of an experimental system is extremely difficult from *within the system*. This is why experimental approaches need to run in parallel with non-experimental, more “natural” observations. It is also why history is so important in peculiar worlds, one reason why history matters. It is, in some sense, “easy” to establish a regularity experimentally (not *practically* easy!): we set up the conditions and control much of the subsequent behaviour with much tweaking. No wonder, after all this labour, that we get something relatively well-behaved. But in a peculiar world, those experimentally established regularities will only extrapolate inelegantly and to a highly restricted range of targets. Both the peculiarity of the world and the peculiarity of experiments themselves restrict the epistemic power of experiment.

Let’s finish with processualism.

I’ve argued that the world at scales investigated by the Earth and historical sciences is peculiar; that is, the world’s actual and modal features are highly sensitive to features of a large number of generating and maintaining processes. I’ve further argued that this provides a diachronic argument for a disunified present. This is closely aligned with an old philosophical view that’s been reinvigorated in the last ten years or so: *processualism* (Nicholson and Dupre 2018; Dupre 2025). In this final discussion, I’ll argue that committing to “peculiarism” aligns well with processualism—they are fellow travellers, if you will—but nonetheless one may adopt my view without thereby being a processualist. And indeed, the scale-relativity that underwrites peculiarism as I’ve induced it from strategic perspectivalism might motivate refusing to swallow processualism whole hog.

I understand *processualism* as a metaphysical view about the fundamental constituents of reality, according to which the world is built of processes.

What, then, is a *process*? The concept is often explained in contrast to *things*, where “things” are taken to be intrinsically stable objects with essences. Take the early-modern view *corpuscularianism*, according to which the world is constituted by units (“corpuscles”) that have particular essential properties, such as size and shape. Accordingly, observable behaviours and effects are due to the interactions of corpuscles. This is a thing-like view. By contrast, processualists deny there are stable entities such as corpuscles (or atoms, or whatever); rather, there are, well, dynamic, shifting, not-things. One way of capturing the notion of processes, and one that well-distinguishes processes from things, is to define them as “entities for the identity of which change is essential” (Meincke 2023, 234). Things are what they are in virtue of stable properties, and processes are what they are because of how they change.

There is a lot for someone fascinated by lost worlds to like about processualism. In particular, the processualist assumption of change and flux highlights the importance of stability as worthy of explanation. One needn’t be a processualist to explain surprising stability, but it is nonetheless a useful perspective correcting a bias toward change as explananda. Further, analyzing the science of lost worlds highlights the importance of the historical individuals and long-term trajectories that are both so crucial for grafting strategies, and a processualist metaphysics naturally cashes out individuals not as things but as temporarily stable homeostatic processes or trajectories. Further, historical kinds are what they are in virtue of the processes by which they form, not their essential properties: processualism thereby highlights these.

However, despite the appeal to process, if you agree with the metaphysics of this chapter, you’re not thereby committed to processualism, at least not by the definition of peculiarity. This is because admitting that an individual is sensitive to generating and maintaining conditions does not thereby commit you to thinking an individual is thus *metaphysically subservient* to those processes. Causal dependence is not in itself sufficient for metaphysical dependence. It seems we are logically and metaphysically free to describe, say, various taxa in terms of stability—as things—or in terms of various processes, while happily admitting their peculiarity.

Further—and this is the step that pulls me back from the processualist brink—once we have the idea of a “thing ontology” and a “process ontology” on the table, it is natural to think that whether we should use “thing-language” or “process-language” turns a fair bit on the scale at which we work. When Cuthill and Conway Morris model the development of rangeomorphs, it seems natural to describe their target as temporarily stable processes, whose

temporary stability is made sense of in light of developmental (and other) processes. When Mitchell et al. model rangeomorph ecology, it seems natural to think of rangeomorphs as a population of things. After all, in that context, the rangeomorphs are treated as objects with set properties, that is, a bunch of individuals in a population.

To establish processualism, we need an argument to say that the levels of description, temporal duration, or scales by which things are best described as processes should take metaphysical priority, and I'm unsure why this should be the case. It might be that at the widest possible scale, we find processes, not things. This is an empirical conjecture, and I'm not sure how we might establish it, but let's go along with it for the argument's sake. Why would processes dominating the widest scale lead us to say processes are metaphysically prior? We'd need an argument to say that the widest possible scale is the most important, a claim I find symmetrical to, and problematic for the same reasons as, claiming that the "lowest" level of description, the base of the mereological heap, should have ontological priority. Another approach might be that the monism and unity of processualism itself is reason to prefer it. Even if you're someone who sees these as theoretical virtues, I'm not sure why we should think a monistic or unified world is more likely due to those virtues. And further, such an argument seems to favour a thing-ontology as much as a process-ontology.

Further, recall the commitment to strategic perspectivalism that underwrote my argument for peculiarity. I suggested that strategic perspectivalism is best explained as an adaptation to a peculiar world. And strategic perspectivalism shifts between thing-ontology and process-ontology (and probably a bunch of others as well!), so adopting a monist ontology in its name seems, well, against the spirit.

This is not to say that processualists and those who commit to peculiarity do not significantly agree; indeed, processualists should expect a lot of peculiarity. Recently, John Dupre has argued that processualism is another route to providing a diachronic argument for disunity: "Process ontology provides the diachronic dimension that unifies the disordered world" (Dupre 2023, 13). In this instance, I think it is fair to think of the argument of this chapter as an *application* of processualism to the scales of interest to palaeontologists and Earth scientists. But a commitment to peculiarity doesn't require one to commit to processualism. Thus, I'll keep to agnosticism about things beyond the scales of the Earth and historical sciences, and stick to a claim about peculiarity rather than processualism.

So, I take myself to have provided an argument to favour Prof. Ichthy's view of the world over Professor Ichthyosaurus'. I've argued we should understand their differences in terms of how peculiar they take the past to be. Insofar as the science of lost worlds has uncovered much loss, and loss is anticipated by a peculiar world, we've reason to think the world peculiar. Insofar as historical investigation uncovers peculiarity, we've reason to think the world is peculiar. Insofar as strategic perspectivalism is adapted to a peculiar world, we've reason to think the world is peculiar. It doesn't follow from this that all things in the past are highly peculiar, nor that they are peculiar to the same extent. But, if we want to commit to a metaphysical view about the structure of the past at the scales of interest to palaeobiology and related sciences, we have a host of reasons to pick peculiarity.