

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>

Anchoring

Prof. Ichthy's cadence amps up toward their lecture's climax as the reconstruction of Broadfoot is projected behind them once more.

"Although Broadfoot is unique—lost—it also contains, and is, multitudes. It is a fossil specimen, and thus can be understood as we understand fossils and the processes that form fossils. It is a mammal, and thus may be compared and contrasted with its mammalian allies, both its rather unimpressive extant relatives and its fellow travellers from the Age of Mammals. But its mammalian alignment doesn't exhaust its phylogenetic properties. It is also a tetrapod, thus partaking in ways of life common across mammals and ourselves. It is also a vertebrate, a bilateral metazoan—an animal—a member of the great ancestrally-spun web of life. But Broadfoot isn't only its ancestry. It is also a predator, likely sometimes prey, it is a secondarily aquatic vertebrate, a venomous organism, an egg-layer, a user of electrolocation. It is a creature built of traits, cells, genes; biochemical, physiological, and immunological processes. And yet more abstractly, it is an anatomical, metabolic—even geometric—object. It is a set of ontogenetic, ecological, and evolutionary processes. Each perspective we take on Broadfoot brings lessons about it. They give us, as it were, a 'place to stand', and each of these platforms gives us information, information which then informs us about other platforms, other 'places to stand'. If I might switch metaphors . . ."

Ichthy pauses, as if expecting some objection from the audience,

" . . . these different perspectives on Broadfoot provide different 'anchors' that, employed with diffidence and care, ensure safe passage on our journey to understanding this ancient, fascinating creature. Many anchors together provide safe harbour for our knowledge. *That*, I think, is the secret to our success."

The chapter has two central goals. First, to introduce our third strategy—“anchoring”—and second, to show how iterative investigation across artificing, grafting, and anchoring can reveal past worlds despite lost denizens.

Chapters 3 and 4 introduced their namesake strategies with sidesteps into non-Ediacaran examples. To introduce anchoring, by contrast, I'll return to an Ediacaran case: Gold et al.'s arguments concerning *Dickinsonia*'s growth. They proposed a synapomorphy aligning that lineage with bilateral animals. I noted that a crucial step in their argument involved conceiving of *Dickinsonia* as a eukaryote that grows via *serial repetition*. Many eukaryotes grow this way: “serial-growing eukaryote” is a diverse historical kind. By conceiving of *Dickinsonia* as such, we can compare and differentiate it across various serially repeating developmental structures in nature. This is the central idea behind *anchoring*: by conceiving of the denizen as a *historical* or an *a historical* kind, we bring it into conversation with a set of other instances and examples—which I'll call *anchors*—allowing us to situate the denizen within the class, thus potentially establishing some of its properties, as well as potentially informing us about the class itself. The set of anchors may just include a single analogue, but, as we'll see, the use of multiple anchors is particularly powerful due to providing more opportunities for iteration. As a warm-up, let's briefly look at how Gold et al. (2015) anchored their analysis of *Dickinsonia*.

Serially repeated bodies are built from repeated, more or less identical segments. Perhaps the most familiar way of conceptualizing serial repetition is to start with segmentation, as seen in annelid worms, centipedes and millipedes, and so on. Each segment (or “annulus”) of an annelid worm, for instance, is composed of the same parts of the various systems (e.g., digestive, excretory, metabolic, etc.) required for the worm to function (only the head and rear are not built of annuli). Thus, to increase length during development, the worm needs only to increase the number of segments. Beyond segmented invertebrates, plants are another familiar serial repeater. They sequentially add repeated parts—leaves—around a stem. In plants, we have two modes of growth: the stem does not grow serially (it gets longer) while the leaves are serial, added around the stem as modular components. Examples of serial repetition of various types are found in every living group that has evolved complex multicellular bodies except for fungi. Gold et al. represent some of these in Figure 5.1.

Gold et al.'s strategy, then, begins by characterizing *Dickinsonia* as a serially-growing organism. This enables them to compare *Dickinsonia* with




FIGURE 5.1 Phylogenetic Distribution of Some Serially Growing Lineages (Fig. 2, Gold et al. 2015). Reprinted with Permission from the *Royal Society B*

various other serial growers, thus navigating among other instances of that developmental phenomenon. These other instances are the *anchors*. Thus, various taxa that share in the historical kind “serially-growing organisms”—trilobites, seaweed, plants, and worms—each act as points of empirical contact for understanding that kind of development generally and for understanding *Dickinsonia* in particular.

For instance, Gold et al. point out differences between the bilaterian mode of growth and growth patterns in plants: “In bilaterians, the development of serially repeated structures often occurs in an anterior to posterior gradient, concurrent with posterior elongation of the primary axis” (Gold et al. 2015, 317). Because *Dickinsonia*’s serial growth more closely resembles bilaterians, this speaks in favour of that alignment. They further point out affinities between *Dickinsonia* and what is found in complex metazoans. This is partially an example of grafting: Gold et al. link serial repeaters in clades to an argument about *Dickinsonia*’s phylogenetic placement. But grafting alone

doesn't exhaust the strategy. Gold et al. are also *anchoring*, that is, characterizing their subject such that we can navigate between multiple systems (past and present) to build an integrated, comparative set of models concerning the relevant lost properties. Note that Figure 5.1 is another example of a mixed perspective, representing both phylogenetic information and historical kind membership ("serial growers"). These different examples of serial growers give Gold et al. (to crib from Prof. Ichthy) a series of "places to stand" when understanding and contextualizing *Dickinsonia's* growth and phylogeny.

So, in this chapter, I'll provide some examples and analyses of anchoring and introduce the role of iteration. I'll then turn to another Ediacaran mystery to tie together the last three chapters into an account of how lost worlds are found: specifically, through a combination of strategic perspectivalism and iteration. I'll close with a brief discussion of the relationship between loss and scientific laws.

As with common causes in chapter 4, there is a clear relationship between anchoring and analogue reasoning, which philosophers have discussed fairly extensively (e.g., Bartha 2022). This discussion has mostly involved identifying various possible schematic justifications of analogy as an argument form (although see Hesse's classic discussion of homology and analogy in biology; Hesse 1966). As I'm here interested in how anchoring works together with other strategies (namely, artificing and grafting), these schematic discussions are less useful. Paul Bartha has recently challenged practice-oriented analyses of analogous reasoning. As he puts it, "To advance the methodological debate, practice-based approaches must either make connections to these general models or explain why the lack of any such connection is unproblematic" (Bartha 2022). On my view, schematic views of analogy should happily fit the examples of anchoring I'll discuss, however (1) their schematic nature makes them inappropriate for bringing out how anchoring works in concert with other strategies, and (2) I'm not interested in the justification of token analogous arguments, but in the overall justification of the strategy in which anchoring plays a part. As such, I suspect I meet Bartha's challenge.

1. Ediacaran Ecology and Its Anchors

Recall Mitchell et al.'s discussion of the neutral-dominated spatial properties of Avalonian ecosystems. A neutral pattern was surprising because modern ecosystems, benthic or otherwise, are generally niche-dominated—that is, what kind of organism you are, or how you make a living, makes a difference

to your spatial location. In light of this, Mitchell et al. raised a worrying idea. In the Avalonian world, did ecology just not work the same way as later times? As they put it: “This stark difference raises the question of whether Ediacaran early animal palaeocommunities had fundamentally different community dynamics to those of the present day” (Mitchell et al. 2019, 2034). The thought is as follows: in the oxygen-poor, sessile landscape of the early Ediacaran, the usual models that ecologists use just don’t gain purchase. In the Avalonian, the laws of ecology are broken.

How should we approach this hypothesis? By anchoring. To do this, Mitchell et al. consider possible anchors and ways of characterizing Avalonian ecologies to make them relevant to one another.

Thus far, there are only a few single spatial studies of fossilized palaeocommunities. Jackson and Blois (2015) review examinations of Quaternary communities (a period spanning the last 2.58 million years to the present). Their analysis focuses on patterns in forests at the scale of thousands of years. Jackson and Blois emphasize the unstable nature of these communities, illustrating “how an ecological community is one small, ephemeral point in a roiling, dynamic unfolding of environmental change, distribution dynamics, and spatially aggregated ecological processes” (Jackson and Blois 2015, 4917). Their results emphasize the role of environmental change in driving niche differentiation. As in modern cases, these (comparatively recent) environments behaved very differently from the Avalonian.

Mitchell et al.’s strategy is to consider features of the Avalonian Ediacaran that might be responsible for a neutrally-driven ecosystem. This involves characterizing the ecosystem in terms of features that potentially could be associated with neutrality. These features can then be sought in other ecosystems. By navigating between these systems, they can test those ideas similarly to how Gold et al. did with serial repetition. Let’s see this in action.

Mitchell et al. identify four features. First, populations in Avalonian ecosystems are small. Second, they are subject to regular disturbances, meaning that assemblies typically don’t survive longer than three generations. Third, dispersal appears to occur over short ranges. Fourth, the resource constraints of the Avalonian ecosystem are minimal, at least compared to most modern ecosystems. Anchoring often involves building a model through which comparisons and contrasts between empirical cases are woven together. As such, Mitchell et al. use modern and Quaternary ecosystems that have some of those properties as a set of anchors that inform an artifical model.

Mitchell et al. draw on various resources to suggest that the features identified above should lead to neutrally dominated spatial ecologies. For instance, Fisher and Mehta (2014) modelled neutral and niche differentiation as a kind of phase transition, a framework that Mitchell et al. employ as artificing. They conceived of ecological systems as shifting from neutral to niche-dominated patterning in a fashion analogous to ice transitioning to liquid water. Using various simulations, they identify properties that will lead a community between these modes of ecological assemblage. As Mitchell et al. summarize: “Recent simulations show that community dynamics in small populations living in fluctuating environments are dominated by neutral processes, implying a lack of small-spatial scale environmental control on ecological dynamics in such systems” (Mitchell et al. 2019, 2035).

The suggestion is that small populations in insufficiently stable environments are unable to differentiate by niche because the environment keeps shifting, and the population isn’t big enough for competition between organisms to have much effect. A lack of dispersal also limits competition because distant dispersal increases effective community size and thus increases opportunities for competition. Similarly, a lack of resource constraint limits competition: with plenty to go around, there is not much to compete over. Crucially, the ecological models are built in conversation with various contemporary ecosystems that exhibit some of these properties.

In effect, Mitchell et al.’s explanation is quite simple: an unstable environment means niches aren’t around long enough for communities to adapt to them, and competition between organisms can’t get going because various factors—e.g., community size and brief assemblage survival—dampen that.

So, do Avalonian ecosystems play by different rules? In a certain sense, yes, but not in a way that threatens the generality of basic ecological dynamics: “While the dominance of neutral processes within these paleocommunities differs substantively from the majority of the modern marine realm, the underlying dynamics are entirely consistent with models of assembly which include both niche and neutral processes, and are similar to those of modern communities subject to the same conditions” (Mitchell et al. 2019, 2035).

So, the Avalonian ecosystems are weird and, in some ways, unique, but not weird and unique in ways that belie ecological understanding: they, in fact, behave as we’d expect given the particular Ediacaran conditions preserved at Mistaken Point. In other words, they are lost, but we understand how and why.




FIGURE 5.2 Approximation of Mitchell et al.’s Reasoning Process, Showing the Anchors in the Top-Right Corner

Let’s summarize Mitchell et al.’s investigation to highlight anchoring. Mitchell et al. start with a series of fossils at Mistaken Point. They use scanning technologies (laser-line probing) and subsequent data analysis to generate a spatial and taxonomic map. Various taphonomic theories and studies of local context are then brought to bear to argue that the fossil assemblage can be read as something like a “snapshot” of the in situ community. This enables the application of spatial point analysis, which generates the surprising neutral result. We could understand the steps up to this stage as involving the characterization and decontextualizing of the Mistaken Point assemblage, considering only taxic and spatial distribution. This is carried out via various perspectival tools: laser-line probing, systems of data analysis, and so on.

With the neutral result in hand, Mitchell et al. reach for a set of anchors. They contrast Avalonian palaeocommunities with modern and Quaternary communities to identify factors that could explain Avalonian differences. These factors are investigated in various ways: for instance, observationally, as well as with various modelling techniques. Such techniques, together with further details inferred from Mistaken Point (for instance, confirming the short life of communities), enable an explanation of neutral niche differentiation. Figure 5.2 is a kind of messy and approximate sketch of the reasoning process.

Mitchell et al.'s strategy is a classic example of anchoring. In the face of a putatively lost world, they abstract a set of properties and look across many examples to compare and contrast to generate a set of explanatory factors that are then re-applied to the (recontextualized) case at hand.

At base, anchoring involves conceiving of your target as a token of a kind such that it can be brought into empirical and theoretical dialogue with other tokens of that kind. Mitchell et al. use perspectival tools, such as the assemblage itself, laser-line scanning, and spatial point analysis, to isolate the spatial ecological properties of the Mistaken Point assemblage. This isolation enables comparison with other ecologies in those terms. Let's characterize anchoring abstractly:

Anchoring involves characterizing a denizen in terms of some set of historical/ahistorical kinds, identifying other instances of those kinds (anchors), and using these anchors to learn about the relevant kind.

Anchoring can be distinguished from grafting, that is, conceiving of your target as part of a historical individual or line of causal continuity. The difference is that in grafting, we "narrativize" denizens by linking them together as a token causal sequence, while in anchoring, we do not assume causal continuity; anchors are understood as tokens of types, not individuals. For instance, Kolesnikov et al. (2017) examine structures in modern biofilms to suggest new interpretations of fossils of Avalonian microbial mats. This doesn't require taking Avalonian mats and modern biofilms as part of the same lineage. As they put it: "It is not that a specific biological community has survived since the Ediacaran—it is that the biological response of microbial communities that manifested itself quite commonly in certain terminal Ediacaran and early Cambrian environments can still be found (seemingly in much more restricted settings) today" (Kolesnikov et al. 2017, 1).

Kolesnikov et al.'s investigation is anchored both by the fossils and by observations of modern phenomena that display relevantly similar properties. Thus, they take them as historical kinds (microbes responding to particular conditions), not historical individuals.

We can also approximately distinguish anchoring from artifice. Insofar as Kolesnikov et al. don't literally construct their analogues, they anchor rather than artifice. Similarly, Mitchell et al.'s identification of the relationship between, say, small populations and neutral domination doesn't rely on

Table 5.1. Strategies, Definitions

Strategy	Definition	Exemplar Perspectival Tool
Artificing	Denizen conceived of as an historical/ahistorical type, a token of that type is constructed and examined.	Geometric models of ontogeny.
Grafting	Denizen conceived of as part of an historical individual or part of a token historical trajectory.	Cladistic Phylogeny.
Anchoring	Denizen conceived of as an historical/ahistorical type, a set of different instances of that type are examined.	Identification of properties associated with spatial neutrality.

instances being causally continuous or being constructed. They examine different naturally-occurring instances rather than building them. As anchoring involves using multiple instances of a kind to inform us about that kind, it often goes hand in hand with artificing, where we might further explore our understanding of the kind by building a model of it. For example, when Woodruff and Varricchio (2011) artficed a model dinosaur burrow, they were partly influenced by examinations of mammal burrows, and indeed used burrowing mammal bones in the experiment—that is, various anchors. Table 5.1 summarizes the strategies, their definitions, and a paradigmatic perspectival tool for each.

Anchoring makes sense of and unifies several more specific proposals philosophers have made concerning historical reconstruction, so I’ll point to several of these before turning to iteration.

In *Rock Bone and Ruin*, I describe a form of anchoring, which I call “exquisite corpse” modelling: in a nutshell, a single, difficult target is imperfectly modelled with multiple analogues. Perhaps if we squint just right, we could view Mitchell’s investigation in these terms. Avalonian communities are understood as being a combination of various aspects of modern ecosystems: those that are small, that are frequently disrupted, and so on. Joe Wilson (2023) builds on this idea to argue that palaeoclimatology can provide partial analogues of modern and future climate systems (see also Watkins 2024c).

Meghan Page’s demonstration that sometimes, as she puts it, “the past is the key to the present” (Page 2021; Dresow 2023) also involves anchoring. She

describes scenarios where some regularity has been confirmed under present conditions, but we are nonetheless interested in the regularity's fragility. That is, is the regularity likely to continue to hold under different conditions? What is the range of the regularity's applicability? As she points out, by testing the regularity against conditions in historical records, we are able to further understand the circumstances under which it fails and under which it holds. By the past-present interactionist view I'm pushing here, we should expect the circumstances described by her account to be common. Mitchell et al. demonstrate that our ecological models of spatial distribution (perhaps a little enriched) can extend into the Avalonian world. They thus teach us something about the robustness of ecological models and the relevant regularities, much in the way that Page describes.

Finally, Aja Watkins examines how various scaling strategies—basically transforming data from imperfect analogues and data generated from the target—are required to gain evidential relevance in many historical investigations (Watkins 2023). She is particularly interested in how this occurs in palaeoclimatological contexts, where differences in scale between events occurring millions of years ago and those currently occurring (or soon occurring) threaten to undermine using *then* to understand *now*. In the terminology of this book, the ways data from analogues and targets might be scaled is a specific kind of characterization that enables comparison, particularly in anchoring.

What makes these accounts examples of anchoring is that all involve characterizing a denizen such that it falls under a kind that includes a bunch of other targets. This differs from grafting insofar as grafting involves unifying the target via a token causal trajectory—making them part of the same narrative, if you want—while anchoring involves taking targets as more or less independent instances of the same kind. It differs from artficing insofar as natural analogues are not constructed. The difference between artficing and anchoring is messy, depending on how much “construction” and modeling is required to connect the analogue and target. But I think the distinction is useful because the practices they highlight—building an instance of the relevant properties as opposed to navigating between various naturally occurring exemplars—differ sufficiently.

These strategies of characterization—strategic perspectivalism—enable iterative, mutually-informing investigations across perspectives. Under the right conditions, this can reveal past worlds despite loss. It is time to see how.

2. Epistemic Iteration




ILLUSTRATION 5.1 So Many Perspectives, So Many Anchors

“With multiple platforms, or multiple anchors,” Prof. Ichthy pauses as if struggling under the weight of juggling so many metaphors, “we can jump between perspectives, gaining a little insight here, discarding a once-promising idea there, all the while generating a richer picture of our target. Starting with a rather simple, likely wrong-headed conception of Broadfoot, we navigate between various ways of testing and probing that conception, and these are refined, made richer, and made more plausible. This jumping between perspectives, I say, is a self-correcting process, taking us from bad ideas to good ideas. Thus, iteration between perspectives is how we’ve discovered so much about Broadfoot despite their uniqueness, despite their loss.”

What is advantageous about having multiple anchors? For Mitchell et al., what is critical about accessing a set of contemporary benthic ecosystems, the Quaternary system, the fossils themselves, and various simulations? One immediate answer—and certainly not a false one—appeals to induction. Multiple anchors mean more evidence, and thus a richer data set to draw on in support of claims. But I don’t think a simple increase in enumerative inductive power captures the epistemology of anchoring specifically, or strategic perspectivalism in general. For this, we need to understand iteration.

Throughout the last three chapters, we have repeatedly seen a basic structure to applications of strategic perspectivalism, first laid out in chapter 3. First, decontextualize: characterize a set of properties using perspectival tools. As Helen Longino has discussed in the context of the study of behaviour, decontextualization is often required in scientific inquiry, since isolating properties renders them measurable and—crucially for our purposes here—comparable across contexts (Longino 2021).

Second, use these isolated properties to apply various investigative strategies. We might artifice, building something analogous to those properties; graft, linking them to a historical individual; or anchor, navigating between different targets that share those properties or contrast with them.

Third, recontextualize: return to the past world and attempt to make sense of the results.

This strategy is fruitfully understood via what philosophers have called *epistemic iteration*. Drawing inspiration from practices in mathematics, where iteration involves an approximate proof acting as a scaffold or springboard for a less approximate proof, Hasok Chang (2004) argues that similar processes occur in the development of measurements, in his case, temperature. He sees epistemic iteration as “a process in which we throw very imperfect

ingredients together and manufacture something just a bit less imperfect” (Chang 2004, 226). As Kevin Elliot points out, such ideas are fairly common across the history and philosophy of science (Elliot 2012; see also work on scaffolding: Chapman and Wylie 2018; Currie 2018a; Walsh 2019; and the papers collected in Caporael, Griesemer, and Wimsatt 2014), and Daniel Swaim has suggested that the iterativity of speculative narratives has important epistemic functions in the historical sciences (Swaim 2024).

I want to emphasize how iterative practices interact with strategic perspectivalism to generate knowledge of lost worlds. I don’t think what I have to say here clashes with these previous discussions, but I do think that the highly localized and particular ways in which iteration plays out in the examples that interest me differ from what occurs for more standardized products, such as measurement concepts and practices, that have energized folks like Chang. This picks up a major theme of my earlier work: that historical science is successful in virtue of *methodological omnivory* (Currie 2015, 2018a). As opposed to following a particular method (e.g., Cleland 2002; Kleinhans et al. 2005), historical scientists throw whatever methods they can against the wall, hoping to see what will stick. Methodological omnivory isn’t simply the idea of using multiple threads of evidence, but actively and opportunistically generating often highly imperfect evidence in highly tailored, local ways. I think the combination of strategic perspectivalism and iteration, as I’ve described it, partly underwrites why methodological omnivory works.

Having multiple anchors is good not only because they provide more data for stronger inductions. Multiple anchors are useful because they provide multiple points for progressive iteration. Consider again Mitchell et al.’s work. The reconstruction and analysis of Ediacaran ecology relies on strategic perspectivalism—characterizing the fossils as recording *in situ* benthic spatial ecology—which enables the application of various anchors (the fossils themselves, simulations, other ecosystems, and so on, as approximated in Figure 5.2). But the relationship between Mitchell et al.’s anchors is iterative in the sense that knowledge generated from one investigation can motivate and be transported into another context, and thus be transformed and transported back to the original context. From spatial point analysis, we learn that the ecosystem is neutrally dominated. Based on a comparison with other ecosystems, we learn that this is unusual. This raises a question: what explains the exceptionalism of Avalonian ecosystems? This then leads Mitchell et al. to navigate between the Avalonian and current ecosystems, thus identifying various properties that appear to correlate with neutral dominance, such




FIGURE 5.3 Possible Lines of Iteration between Three Perspectives; Light Arrows Feed Forward Changes, Dark Arrows Feed Backward

as small populations. These hypotheses then lead to various tests—appeals to simulations, for instance—before finally being brought together in a re-contextualized explanation of the Avalonian Ediacaran. Iteration goes beyond enumerative induction because what is learned from one perspective can, under the right conditions, be transported to another perspective and back again. This is due, I take it, to the lack of perspectival isolation as discussed in the last chapter. Under the right conditions, as we’ll see, this brings epistemic benefits.

Iteration in this context, then, occurs when (1) some information is generated via some perspective (that is, a characterization and analysis with some perspectival tools), which (2) is then applied to another perspective, which results in (3) that information being transformed and/or (4) that perspective being transformed. This then leads to (5) that information, transformed, being returned to the original perspective and/or (6) the information being brought to a third perspective. Figure 5.3 captures the basic idea.

With iterativity on the table, the task now is to show how, in combination with strategic perspectivalism, loss and putative loss can be overcome. That is, how historical scientists discover lost worlds.

Recall a problem raised in our discussion of de-extinction: if we attempt to clone an extinct animal, how would we be sure that it, in fact, behaved and looked the way that its past world cousins did? We only know mammoths

from their carcasses, their elephant relatives, and their genomes. So, if we had a cloned “mammoth,” how would we be sure that elements of its phenotype (especially its behaviour) were relevantly similar to past mammoths and not due to our cloning process or other sources of noise?

This is a general problem when dealing with loss. The past could always have been different in surprising ways, which undermines uniformitarian assumptions. Historical scientists always labour under the possibility of loss: it could just be that dinosaur bones worked differently from rabbit bones, that the microbial mats of the Ediacaran affected decay differently from today’s bacterial films under experimental conditions, and so on. We need to be very careful about how we articulate this challenge. Simply pointing out the possibility of loss and demanding a general answer, if we set the bar too high, more or less amounts to demanding an answer to the problem of induction. That’s above my paygrade, seems to be a problem with any inductive process, let alone those involving lost worlds, and frankly, isn’t something I find productive or enormously interesting (perhaps it’s a little like a philosophical version of asking whether dinosaurs feel pain: if you’re not careful, it doesn’t lead to a very productive discussion). What is significant for our purposes is how iteration plays a critical role in identifying loss when it occurs and partially overcoming the challenges posed by loss and potential loss.

The strategies I’ve outlined all have potential failure conditions. In articulating, one might build something that isolates a property the target denizen lacks; in grafting, you might over-emphasize the continuity and similarity between denizen and putative members of the same historical trajectory (or misidentify a historical individual); in anchoring, the apparent similarities between instances of types might be just that, only apparent. Each strategy more or less involves taking a perspective—using tools to isolate some set of properties—which can go wrong in multiple ways. As with any inductive inference, attributions of loss can involve various kinds of errors. We might commit a false negative: attribute loss when there has not been any; we might commit a false positive: think there isn’t loss when it has occurred.

I argue that iteration mitigates against these errors. How? I won’t claim that iteration and strategic perspectivalism guarantee success—we’re working with induction after all—but rather that, under the right conditions, iteration constitutes a self-correcting process that gives us reason to think the investigation at hand is either on the right track or, if we’re barking up the wrong tree, this will be discovered.

Let’s start by going abstract.

Take some characterization of a denizen, call it p , and some characterization of an anchor, call it p^* . That is, there is a perspectival tool such that we can conceive of p and p^* as being of the same kind. Thus, we've taken a perspective such that p and p^* are putatively of the same kind. How do we determine whether, in fact, $p = p^*$? That is, that p and p^* are, in fact, the same kind in the relevant way, and thus that loss hasn't occurred (at least pertaining to p)?

To answer such a worry, we might conduct some investigation of p^* , maybe model it, measure it, etc. And then see whether we get a similar result by performing an analogous investigation of p (or vice versa). Notice that whether we get a positive result or not, we still learn something important about p , at least in relation to p^* . If we get a positive result, then presumably we should be more confident that $p = p^*$. We might then turn to another anchor, measuring or modelling that. However, if we get a negative result, we will seek an explanation for why the results differed. This might involve adopting a different perspective on p or p^* , for instance, or shifting to another putative anchor. Maybe we would try to use the discrepancy between p and p^* to hypothesize a more general model that explains the differences, or we might hunt for further instances that also show that discrepancy. Failures, then, also act as potential scaffolds for further perspectives.

What we have, I hope, is a kind of "self-correcting" procedure. The bet is that if the isolation has gone wrong, then recontextualizing will backfire; if you've contextualized incorrectly, this will be revealed through iterative tests. There are no guarantees or perfect solutions here; chains of unlikely coincidences can happen, but we have good reason here to think that errors are at the very least minimized.

As we move forward, navigating between different perspectives, we identify false positives and negatives in loss-attribution, gain new knowledge about p , while also gaining new knowledge about p^* . As we saw in the last chapter's discussion of living fossils, our investigation of the past and our investigation of the present are coupled. Thus, I think, iteration both helps us discover and understand the properties of past worlds even if they are lost, in part by making sure that we've characterized them adequately.

You might complain about the abstractness of this argument. You shouldn't. Consider the relationship between Mitchell et al.'s strategy and all of those p s and p^* s. Mitchell et al.'s investigations more or less follow the abstract pattern. With the relevant changes, you could describe most of the cases we've examined thus far in these terms (I'll leave this as an exercise for zealous readers). The power of iterativity becomes clearer, I think, when

we connect it explicitly with arguments I've made over the last few chapters about perspectives themselves and how they relate. I'll particularly emphasize the role of recontextualization.

Massimi urges us to imagine perspectives "as opening up a 'window on reality' that extends well beyond the boundaries of the representation itself and where the depth, angle, and scale of the representation leave enough room to make inferences about the space and what's in it" (Massimi 2022, 40). In Charlotte Kenchington's palaeoartistic reconstruction of the Ediacaran Avalonian (Figure 2.3), the vanishing point, created by a horizontal line and differently darkened or lightened rangeomorphs, provides a sense of space that reaches well beyond what is literally depicted in the image. Massimi wants to draw an analogy between that effect and scientific perspectives. My interpretation of this idea is fairly simple: a perspective on a target involves characterizing it as some kind, and this has further implications for that kind and other things that are included within it.

Recall the disagreement between Gold et al. (2015) and Hoekzema et al. (2017) from chapter 4. There, we saw that although both employed different perspectives (Gold et al. characterized rangeomorphs as exhibiting terminal-end growth within a phylogenetic bracket, while Hoekzema et al. characterized them in terms of a developmental sequence abstracted from a set of fossils), these were not independent from one another, as Hoekzema et al.'s approach involved empirically testing an assumption underwriting Gold et al.'s work. Further, both perspectives overlapped insofar as they took *Dickinsonia* to grow via serial repetition. Perhaps unlike perspectives in visual art, scientific perspectives can sometimes be superimposed: *Dickinsonia* is simultaneously both a phylogenetic animal and grows via serial repetition.

Such overlaps, links, and extensions from and between perspectives enable a widening scope of rich knowledge. They extend the range of alternatives available, and thus the investigative options and opportunities for developing knowledge. It turns out, so long as you're clever with perspective, problems of loss are often surmountable. Just don't let N equal zero.

Do we have good reason to think these iterative processes are in fact self-correcting, and that those corrections will lead to knowledge of lost worlds? I think—under the right conditions—we do, and to see this, I'll turn to discussion of iterativity in the measurement literature that follows from Hasok Chang's influential work (Chang 2004).

Here are two related challenges to developing a measurement of some quantity or phenomenon. First, generally, how do we validate a measurement

and know when one is “correct”? If you wish to test your measurement, that is, check whether it gives the correct result, presumably, you need to already know what the correct result is, that is, already have a measurement. Many of the things we wish to measure are complex, only partially known to us, and often highly ambiguous prior to developing a measurement. So, what do we check our devices against, if not their own results (Chang 2004 calls this “the problem of nomic measurement”)?

Second, how do you develop a measure for something when, at the outset of the investigation, you don’t quite know what that thing is? Chang highlights a difference between iterative proofs in mathematics, where “the true function we are trying to approximate is already known or at least knowable” (Chang 2004, 216), and empirical measurement cases, where the motivating phenomenon is often highly unsystematic and unstable. Temperature, as we experience it, for instance, is highly variable, context-sensitive, and messy, witnessed by our various attempts to incorporate “wind chill,” “dry heat,” and other bids to make sense of the difference between temperature as measured by sensible Celsius (and I guess Fahrenheit if you must) and how we experience variation in phenomenal temperature.

Chang’s answer to both challenges appeals to the iterative testing and tweaking of measurement across various, related apparatuses, thus generating “a process in which the successive stages of knowledge, each building on the previous one, are created in order to enhance the achievement of certain epistemic goals. For each stage, the later stage is based on the earlier stage, but cannot be deduced from it in any straightforward sense” (Chang 2004, 226; see also Tal 2016 and 2017). What is self-corrected in Chang’s cases, then? Correction involves (1) the convergence of iterations toward a fixed point—similar in a sense to consilience of evidence—and (2) “agreement between the concrete image of the abstract concept and the actual operations that we adopt for an empirical engagement with the concept (including its measurement)” (Chang 2004, 216). That is, we have some idea of what we’re trying to measure, and through mutual adjustment, this and our measurements come into agreement. Our experiences of temperature provide a starting point and reference point for the development of increasingly sophisticated apparatuses for measuring temperature and theorizing about it. This then leads us to identify useful fixed points, such as the boiling point of water, to build further apparatuses and measurements for.

What might license the results of such iterative processes? As Valde and Scarffe have recently put it, “Rather than resting on a non-social or objective

foundation, their processes get their legitimacy from reflexive interactions that are, at least in part, governed by social norms” (2025).¹ Insofar as such processes are “objective,” it is in some form of coherentist, socially-mediated way. But this needn’t involve vicious circularity, nor does it involve losing empirical contact. To see why, let’s use the term *agreement* to do a little schematization of iterative accounts of measurements.

We can distinguish between the *target*, the *operational constraints*, and the *criteria* of agreement. Chang is interested in measuring temperature; his target phenomenon is (say) differences in heat and cold. The operational constraints come from the apparatus used and the derivable data—the perspectival tools, if you will, and the theories of temperature developed. Chang’s criteria of agreement are coherentist in nature, but coherence understood in a broad sense: it concerns not merely the relationship between propositions, but how the components of the system contribute to the success of that system’s aims. So, what it is to *agree* for Chang is for the target and operational constraints to converge coherently. Other philosophers interested in measurement, looking at differing forms of measurement, have different accounts of agreement.

Alisa Bokulich, for instance, appeals to iterative coherence tests across radiometric data, arguing that these increase accuracy. The points of these tests are not “to arrive at a specific age for the event but rather to pinpoint potential sources of error in the initial methods and assess the magnitude of their effects” (Bokulich 2020, 433). Here, an iterative process in the development of a measurement leads to its increased accuracy in measuring age. So, the target is geological time, the operational constraints involve the apparatus and theory of carbon decay and other radiometric phenomena, and the criterion is increased accuracy. Here, increased accuracy concerns the ability to identify and control for error.

A third example is Sharon Crasnow’s (2021) discussion of political-scientific measures of democracy. These measures, at base, take a bunch of measurable proxies as an indication of how ‘democratic’ a country is. She develops a form of “objectivity” that holds when a measure “is subject to theoretical, empirical and pragmatic constraints” (Crasnow 2021, 1209). She particularly emphasizes how the normative aims of the measure interact with empirical results in an iterative manner: “While conceptualizing democracies can be

1 Valde and Scarffe (2025) draw a convincing analogy between the socially mediated coherence of iteration in science and in law.

a starting point for developing measures, concepts may alter in response to changes in regimes around the world, the results of empirical research on the causes and effects of democracy, debates about the values democracies exemplify, the effects on policies of different ways of conceiving of and measuring democracy, and on the goals of research” (Crasnow 2021, 1224).

So, here agreement occurs across multiple elements: the target (democracy and its proxies), the operational constraints, the means of generating data about those proxies, the practical effects of those decisions on policy (depending on how we carve things up), and a set of criteria involving values and goals. Crasnow argues that agreement here (she uses the language of objectivity) occurs when these come together to succeed across these various criteria.

I think the message we should take from the work on iteration in measurement is clear: in the good case, we can argue that an iterative procedure is self-correcting when we can point to a set of constraints emerging from the target, operational constraints, and criteria, which plausibly lead to the relevant kind of agreement.² Let’s now turn from measurement to consider lost worlds.

Attempts to measure temperature had to grapple with the unsystematic and ambiguous initial phenomena of temperature as we experience it. Historical scientists tackling lost worlds face a similar challenge: is the past world in fact lost and, if so, what properties in particular are lost? When delivering their initial neutral result, Mitchell et al. were unsure whether this could be understood via our usual ecological models, or whether new models needed to be developed specific to the Ediacaran. That is, is the Avalonian Ediacaran lost insofar as it is an unusual edge-case of usual ecological dynamics that we don’t see today, or is it lost due to being a completely different kind of ecosystem, demanding a new set of dynamics? To resolve this ambiguity, a complex investigation was required, which involved iterating between various anchors, models, and what we know of the Avalonian systems themselves. Agreement, here, was between these various resources and, I think profoundly, *toward* the past world.

So, the agreement we want is toward a past world, and thus knowledge of what is lost within it. So, what resources do we get from strategic

2 If you find this kind of coherentism distasteful, a view like Khalifa and Goldberg’s (2022) might be more attractive to you. Although in section 4, I’ll provide some reason to doubt that the epistemic phenomena that interest them are common or central in the science of lost worlds.

perspectivalism? I'll suggest that the role of recontextualization, and particularly that grounded in knowledge derived from traces, puts historical scientists in a relatively strong position regarding whether their processes are self-correcting toward how the past world in fact was.

In the cases I'm interested in, we're not developing a measurement (although, as mentioned in chapter 1, there is a lot for philosophers to think about concerning measuring the deep past!) but aiming to empirically discover a past world. Here, the agreement is not between an "abstract object" ("temperature" in Chang's work), various phenomena related to that object, and various apparatuses and procedures of measurement. Instead, it is between things like fossils and their interpretations, models of general dynamics, currently existing analogues, and so forth. Here, recontextualizing plays a critical role. It is a grounding in traces that plays a large role in ensuring self-correction throughout the procedure.

Derek and Michelle Turner have recently highlighted the role of "contextualized research" to explain the difference between "ancient-alien" theories and legitimate archaeological theories of past human structures and cultures (Turner and Turner 2021). Ancient-alien theories typically hold that some achievements of our ancestors couldn't have been accomplished by such "primitive" folks unaided, instead positing that some extraterrestrial intelligence was to blame. They often involve highly fine-grained analyses of particular artifacts (for instance, von Däniken's interpretation of a Mayan fresco as depicting an astronaut), combined with a form of analogous reasoning that bears a superficial similarity to anchoring. For instance, ancient alien theorists will point to multiple examples across the globe and history to underwrite their claims. However, this similarity between anchoring and conspiratorial analogous reasoning is indeed only superficial, and this is due to the role of recontextualization. Turner and Turner's discussion is worth quoting at length:

When archaeologists are investigating a particular question about a particular site, they do so in a rich informational context that includes established investigative practices and lots of relevant background knowledge, as well as multiple lines of evidence. This is a way of leveraging previous epistemic successes. Understanding the imagery on a Mayan ruler's tomb does not involve just looking at it, as von Däniken did with Lord Pakal's sarcophagus lid before declaring it to be an

image of an ancient astronaut. Instead it requires understanding context such as the identity of the dead king and his place in Mayan history as revealed by past research on Mayan texts. Understanding how Mayans adorned their bodies might be relevant to understanding the object near Pakal's nose. Also highly relevant are Mayan cosmology, symbolism and artistic conventions as understood from texts, inscriptions, thousands of objects of art, and architecture at many other sites. One might want to consider the social and political nature of the city where Pakal ruled, as understood from texts and excavations, as well as the layout and social meanings of the pyramid in which the king was buried. Studies of the bones found inside the tomb and the tools and materials used to make the sarcophagus might also be relevant . . . When archaeologists seek to interpret something like Lord Pakal's sarcophagus lid, they rely heavily on a supporting structure of previous research, even if very little of that previous work actually focused on the artifact in question [references removed]. (Turner and Turner 2021, 20)

We've seen similar recontextualization in the palaeobiological examples that have energized this book. Analyses of *Dickinsonia*'s development are recontextualized into later Ediacaran environments as well as phylogenetic hypotheses; trackways putatively demonstrating limping dinosaurs are contextualized with dinosaur physiology and possible environmental variation; Mitchell et al.'s neutral ecological result is made sense of by considering further features of Avalonian assemblages: their small size, short duration, and so on. This shows the crucial role that records—traces—play in overcoming challenges from loss. Downstream information that hasn't been erased provides a crucial lynchpin tying analyses built from artficed experiments and models, and various anchors, to the actual past world we're interested in. There is continual empirical contact—the world “pushes back”—and for historical scientists the agreement is to a past world revealed through traces.

This emphasis on iteration brings together, and I think makes sense of, a bunch of philosophical claims about the nature of historical inference. Both Chapman and Wylie (2018) and Thomas Bonnín (2019) have used Toulmin schemas to understand palaeoscientific arguments. These schemas are extremely helpful, eschewing hypothetico-deductive or probabilistic

reconstructions of arguments in favour of various claims with various relations such as “warrants,” “rebuttals,” etc. I take their work to be a description of the kind of iterativity I’ve tried to highlight here. It also makes sense of the claims Kim Sterelny and I have made about the epistemic benefits of narratives (Currie and Sterelny 2017; Currie 2017a). We argue that one advantage narratives have over more abstract theoretical explanations of historical episodes is that the narrative, stretched over time as it is, has more opportunities for empirical testing. This can be understood as a claim about grafting: taking an episode as part of a token trajectory enables multiple locations for iterative testing.

Another of Chapman and Wylie’s claims about archaeology is that it lacks epistemic foundations. As they say of archaeological data: “Neither these data nor the evidential claims based on them constitute a self-warranting empirical foundation, and yet they can powerfully challenge and constrain the reconstructive and explanatory claims we project onto the cultured past” (Chapman and Wylie 2018, 6).

On this view, the reason that archaeological interpretation is not (always) beholden to archaeological presuppositions is because such presuppositions do not play unquestionable, foundational roles. Rather—to put the point in the terms I’ve developed here—they consist of sets of perspectival tools that, when judiciously applied, generate various perspectives and provide the basis of iterating studies (we’ll return to this thought in section 4).

Thus, iteration and strategic perspectivalism allow historical scientists to uncover lost worlds.

3. Bringing It All Together: How the Metazoan World Arose

“Finally, I want to ask: what sealed Broadfoot’s fate?” The next transparency shows a series of black lines stretched across a horizontal axis, each ending with the outlined representation of some lineage. Most lines associated with mammals, including Broadfoot, reach about halfway across, while many graced with reptiles stretch across to the right of the image, where a red line indicates ‘present day’.

“What pushed these animals aside and made space for their replacement? That is, what heralded the so-called ‘Age of Reptiles’?” Prof. Ichthy indicates everyone around them with a sweep of a flipper. “This is a difficult question. Understanding a

lost world is one thing; understanding how it became lost is another entirely. Happily, understanding both lost things and how they became lost are not independent projects. That is, examining Broadfoot and their relatives, and trying to understand how their world fell, and ours arose, are complementary scientific endeavours.

“The iterativity of knowledge just keeps paying dividends, it seems.”

Let’s briefly characterize the conceptual resources I’ve been collecting thus far (take a deep breath!).

I’ve introduced the notion of a *past world*, a subsection of actuality indexed to some set of *denizens*, that is, the entities, processes, properties, relations, and dynamics that populate that world. *Loss* occurs when denizens of some past world do not exist in the present, *erasure* when the traces of those denizens are not retained. A world’s being lost is, then, a derivative notion: a world is lost when some of its significant denizens are lost. *Perspectives* (in this context) are ways of characterizing denizens that are enabled by *perspectival tools* such as fossils, models, comparisons, experiments, observations, and so on. In trying to understand putatively lost worlds, historical scientists make creative and opportunistic use of perspectives (“strategic” perspectivalism), developing iterative investigations between and within them. These investigations involve some combination of artificing (building simulacra of denizens), grafting (unifying denizens with the present and other times via causal trajectories), and anchoring (characterizing denizens as tokens of a type and comparing them to other examples). Because perspectives reach beyond the target that we’ve taken the perspective on, under the right circumstances (where there are sufficient perspectival resources!), historical scientists can solve challenges from loss by adopting self-correcting processes of comparison and testing across perspectives, thus uncovering lost worlds.

Phew, it’s been a busy few chapters.

Most of the case studies I’ve drawn on use some combination of the strategic package I’ve developed. However, my descriptions aimed to bring out a particular aspect (artificing, grafting, anchoring). This was somewhat artificial, downplaying how strategic perspectivalism involves weaving variations of these strategies together. So, in this section, I will provide an example that highlights this pluralistic aspect, thus demonstrating the epistemic prowess of the bag of tricks I’ve been describing. To drive the argument home, then,

let's turn to another set of questions arising from the Ediacaran: why did the complex metazoans of the later Ediacaran and Mesozoic turn up when they did? Why did the Avalonian world fall? We'll look at Nick Butterfield's controversial (see Cole et al. 2020) answer. As we'll see, it involves a pretty remarkable example of strategic perspectivalism that prominently involves grafting, artificing, and anchoring together.

Let's begin with some very big-picture and very deep-past context.

At some level of description, life as we biochemically know it arose in the "Great Oxygenation Event" around 2.4 billion years ago (Ligrone 2019; Lyons et al. 2021). Before that, the Earth's atmosphere was dominated by methane, and from around 3.5 billion years ago, the biosphere consisted of single-celled critters sporting anaerobic metabolisms. There was a steady supply of both oxygen and hydrogen, as sunlight split water vapour in the atmosphere into its constituent elements, but it wasn't until cyanobacteria arose that oxygen began fixing in the atmosphere. The major metabolic innovation of cyanobacteria is photosynthesis, which uses sunlight to oxidize water, thus producing excess oxygen. As cyanobacteria steadily did their work, an oxygen pump was created, and the atmosphere became the nitrogen- and oxygen-dominated beast we know and love today (and the anaerobic organisms, for whom excess oxygen is poison, sought refuge in the deep oceans and, eventually, our guts).

Fossil evidence shows that eukaryotes arose around 800 million years after the Great Oxygenation Event. Complex multicellular eukaryotes—animals and so on—took another billion years (Mukherjee et al. 2018). What gives? Why would it take a boring billion years (one thousand million) for complex metazoans to arrive on the scene? A common answer appeals to the biochemical requirements of eukaryotes, especially their being oxygen-greedy. Most eukaryotes³ must absorb oxygen to function, so oxygen-availability is a constraint on eukaryotic complexity and size. Oceanic oxygen levels fluctuated heterogeneously throughout the Phanerozoic, not reaching high levels until the dawn of the Cambrian, and only stabilizing around modern levels in the late Jurassic (Lu et al. 2018). That this correlates with the Cambrian explosion—complex metazoans finally getting their act together—has not gone unnoticed. We have, then, what seems like a neat explanation for what has been called the "boring billion," the eternity between eukaryotes arriving on the scene and metazoans arising. Perhaps our ancient ancestors were awaiting the right oxygenic conditions.

3 For examples of non-aerobic eukaryotes, see Fenchel (2011).

Butterfield argues that this oxygen-driven explanation gets things backwards: it was the eukaryotes themselves that generated the oxygen-rich conditions that underwrote their evolution. Let's see why he thinks this, and along the way, see him graft and anchor (with a dash of artificing).

Butterfield makes a few arguments against the oxygen-driven explanation. I'll mention two. First, he points to monophyly. All evidence suggests that animals form a "well-behaved" clade, that is, they share a common ancestor. That is odd: if the only thing holding a bunch of microbes back from complex multicellularity is oxygen availability, then wouldn't separate lineages take advantage of those conditions once they turned up? So, Butterfield argues we should expect polyphyly under such conditions. If so, then the monophyly of complex animals speaks against an oxygen-driven trajectory. Second, he highlights "the fundamentally non-uniformitarian nature of the pre-metazoan biological pump (precluding any simple extrapolation of atmospheric composition from marine redox signatures)" (Butterfield 2009, 3). It's worth pausing to understand this, not only because (as we'll see) it is crucial to Butterfield's positive account, but because it involves an appeal to loss.

We can understand the biological pump as the process by which carbon is transported from the sea's surface to the ocean floor (de La Rocha and Passow 2006). Some proportion of this occurs by the passive sinking of corpses and excreta (sometimes called "marine snow"). But another proportion is "active." Basically, large numbers of organisms travel from the ocean surface at night to the lower oceans by day. A great migration occurs wherein herds of smaller aquatic life, krill and the like, head up to feed at night on phyto- and zooplankton, followed by a vast ecological web of fish, cephalopods, and cetaceans (oh my!). As the sun rises, krill lead a caravan of trophic hangers-on back to the lightless depths. It seems basically accepted that this process is driven by predation. Krill (and similar organisms) feed during the night because darkness provides cover, making them trickier prey, and return to the safety of the ocean's "twilight zone" during the day (Lambert 1993). As Pinti et al. have put it, we can understand the migration as "the product of a co-adaptive 'game' where many animals seek to optimize their migration patterns relative to the migration patterns of their respective prey, predators, and conspecifics" (Pinti et al. 2023, 997).

What is the relationship between all this and lost worlds? These days—and by that I mean more or less since the Cambrian—the relationship between oceanic and atmospheric oxygen has been basically stable, meaning we can infer from information about deep-water geochemistry to atmospheric

oxygen and carbon (Mackensen and Schmiedl 2019). Butterfield has argued against “the assumption that ancient oceanic structure was comparable to today’s” (Butterfield 2009, 3) on the basis that, well, it couldn’t have been: without complex metazoans, you don’t get the biological pump, without the biological pump, the biochemical cycling of oxygen and carbon is different, and it is that cycling that ensures continuity between atmospheric and oceanic carbon.

Butterfield’s argument here is a nice example of Page’s notion of the past being the key to the present. Nowadays, there is a stable relationship between oceanic and atmospheric carbon and oxygen (well, sort of stable—Butterfield 2009 points to a bunch of exceptions), which is handy because it lets us infer from proxies concerning oceanic carbon to the composition of atmospheric carbon over time (which is important for establishing fluctuations in the relationship between global temperature and atmospheric carbon). But how stable is that regularity, and what does it depend on? By examining a time before the biological pump, Butterfield is able to examine that relationship.

This is a classic example of anchoring: a regularity’s robustness and limits are tested by linking them to various instances across time that are relevant to them. We can then iteratively shift between instances and models of that regularity. This allows us to potentially identify cases of continuity and of loss: if Butterfield is right, then pre-Cambrian biochemistry played by different rules than those obeyed from the Cambrian onwards. And he thinks he has a reason for the posited discontinuity, specifically, the organisms of the Ediacaran themselves influenced patterns of oceanic oxidation: “Multicellularity and large size introduce an entirely novel range of fluid-dynamic properties and ecophysiological opportunities for organisms, not least an unparalleled capacity to circulate and swim through the water” (Butterfield 2009, 3). In short, metazoans didn’t need to wait until oxygen levels were right. They got there themselves through some nifty ecological engineering.

Early (non-complex) multicellular eukaryotes were single-celled organisms knitted together without any particular architectural or structural flourishes (compared to, say, the elegance of the rangeomorphs), and—as we saw in chapter 3—the dynamics of obligate osmotrophy makes getting bigger tricky. As energy dispersal relies on passive diffusion, and as surface-area-to-volume ratios get less efficient as size increases, the traditional osmotrophic system breaks down as you get larger. However, Butterfield points out that the common ancestor of eukaryotes were single-celled organisms with flagella. Flagella are the basic means of bacterial locomotion: probably most

well-known as the “tails” of sperm cells. Butterfield argues that the effect of these tiny wriggling organelles “fundamentally changes the nature of transmembrane gas exchange” (Butterfield 2018, 3). They do this more or less by creating a current that enhances oxygen diffusion, pulsing it through their bodies. By banding together, they amp up this effect: “Simply by arranging themselves into colonies, flagellated cells gain access to fundamentally greater levels of gas exchange than would be possible in isolation. Far from being frustrated by low levels of oxygen, then, multicellularity presents a unique opportunity to escape the ‘tyranny of diffusion’” (Butterfield 2018, 4).

“Simple” metazoans, such as coral, create significant vortices that structure oxygen supply, especially flow, in their local environments. This increases the availability of oxygen and the efficiency of oxygen absorption. This underwrites Butterfield’s arguments against an oxygen-driven explanation of the timing and nature of complex metazoan evolution. Oxygen isn’t a constraint on getting bigger: getting bigger or banding together in communities is a solution to a lack of oxygen.

So, why the long wait for complexity? What explains the boring billion? Butterfield argues that the answer lies not in an external constraint like oxygen availability, but in the sheer difficulty of evolving the phenotypic and developmental components themselves, what we might call the internal infrastructure constitutive of complex metazoan life: “The challenge appears to lie in the design and assembly of the machine itself” (Butterfield 2018, 5). Developmentally speaking, simple multicellularity only requires iteration. But tissue-grade multicellularity, where various cells come together as a tissue to perform a function, is significantly more complex. Butterfield’s explanation needn’t commit to the idea that complex metazoan traits have low fitness—that they are selected against—rather, it could also be that their evolvability turns on highly-specific and unlikely developmental trajectories (that is, they are highly “peculiar” in the parlance we’ll introduce in chapter 8).

So, Butterfield’s negative argument relies on adopting a perspective enabled by a set of tools: examination of currently living analogues of simple metazoan life; metabolic and fluid-dynamic analyses of flagella and other cellular organelles; and various artifactual models of those dynamics, all contextualized to a lost world that existed before complex metazoans and the biological carbon and oxygen pumps. How then does Butterfield account for the emergence of those “modern” systems? Here, he takes on a further set of perspectives, and for our purposes, is quite explicit about this: “Animals are large aerobic heterotrophic organisms, so it makes sense to view them




FIGURE 5.4 Modern Oceanic System (Detail from Butterfield 2018, Fig. 3). Reprinted with Permission from *Geobiology*

as oxygen sinks. But as swimmers, pumpers and tethered sources of drag, they also contribute importantly to environmental mixing. And as collectors, transformers and translocators of organic carbon, they fundamentally alter rates and styles of biological oxygen demand” (Butterfield 2018, 7).

Let’s unpack this. First, animals consume oxygen (are aerobic) and do not generate energy from the sun, but get it by consuming others, using their oxygen supply (they are heterotrophic). Thus, they are oxygen sinks. Second, animals shift water around with all their swimming and suchlike, thus acting as environmental mixers: they shift the biochemical nature of the ocean, in particular, leading to well-oxygenated surface waters. Third, they change the carbon game by absorbing and excreting large amounts of carbon. For instance, krill interact with oxygen systems differently from phytoplankton: they don’t produce their own, for one thing. Thus, by consuming large amounts of phytoplankton, krill convert one kind of carbon-oxygen system into another. Each of these perspectives enables the isolation of particular properties, their comparison with relevant instances, and their subsequent recontextualization. For instance, add these aspects together with the diurnal migration from the surface to the depths, and we have the remarkable biological pump represented in Figure 5.4.

Figure 5.4 represents the oceans as we know them: the daily migration (“diurnal vertical migration”) transports carbon and oxygen to the sea floor in a great pump, while swimming biomasses mix the oceans, generating fairly homogenous oxygenation, although with higher concentrations in the surface waters. The representation is highly schematic, focusing on biological interactions and feedback while ignoring physical processes such as tides.

But how did we get here? Butterfield might have convinced us that the oceans of the deep, deep past are lost—that our own oceans work by fundamentally different dynamics—but what were those previous oceans like and how were they transformed into our own? Here, we find grafting coming to the fore: there is a story to tell connecting the unfamiliar past to the familiar, and by linking these through a series of trajectories, they become evidentially relevant not because they are the same, but in part because they are different.

Let’s begin in the pre-Metazoan oceans, seen in Figure 5.5. In the pre-Cryogenian (earlier than 720 million years ago), primary production in the oceans was dominated by cyanobacteria. Reliant on sunlight, these spread throughout the surface waters and shallower seas. The evolution of sponges and other filter feeders in shallow benthic environments cleared cyanobacteria from those contexts (shifting them to the surfaces of deeper oceans). This sets the stage for the evolution of algae (Brocks 2018) and phytoplankton: larger eukaryotic photosynthesizers. By the Avalonian period, rangeomorphs, sponges, and other filter-feeding metazoans create an increasingly well-mixed and concentrated oxygenic and carbonic environment for shallower benthic communities.

Butterfield’s narrative continues. By the Cambrian, the late Ediacaran evolution of moving metazoans (worms and our friend *Dickinsonia*) had begun processes of “soft sediment bioturbation”: basically, the seafloor was transformed as various critters burrow through it, creating richer carbon sequestering and oxygenic exchange (Teal et al. 2008). Further, the evolution of early muscled and skeleton-based metazoans enabled mobile filter feeding and eventually the evolution of predation, and thus the modest beginning of something like the modern biological pump, as in those parts of the ocean deep enough to escape light, eukaryotes begin their daily plunge and nightly surfacing, as we see in Figure 5.6.

With the emergence of large fish in the late Devonian, we get the architecture of the modern oceans. In the last chapter, we saw how Ediacaran fauna might not be failed experiments in evolution after all, but rather include stem-metazoans capable of helping us make sense of the evolution




FIGURE 5.5 Pre-Metazoan Oceans (Detail from Butterfield 2018, Fig. 3). Reprinted with Permission from *Geobiology*




FIGURE 5.6 The Cambrian Ocean (Detail from Butterfield 2018, Fig. 3). Reprinted with Permission from *Geobiology*

of, say, bilaterans in the case of *Dickinsonia*. Here, we see another example of grafting: the fauna of the Ediacaran, not only sponges but presumably rangeomorphs and their ilk, began processes of filtering and structuring their aquatic environments, and concentrating, mixing, and distributing oxygen and carbon, in ways that in part created various conditions required for the eventual evolution of the oceans we know today.

By adopting multiple perspectives, Butterfield iteratively builds a rich picture of a lost world, but not only that, he also provides a picture of how the world became lost and its role in creating the conditions leading to more familiar times. Crucially, the model is continually re-contextualized with features of the past worlds we're interested in. This continual return to past contexts makes it more likely that mistakes, when they arise, will be discovered. The more alignment we find between accounts like Butterfield's and context concerning the past world, the more confident we should be that the lost world has been found.

It's worth noting that I'm not here committed to Butterfield's model being correct. For instance, Austin Booth and Ford Doolittle have provided empirical and conceptual reasons to question the uniqueness of both eukaryogenesis and the special powers of eukaryotes (Booth and Doolittle 2015). Although their ideas don't conflict directly with Butterfield's, they could undermine part of his picture. My claim is that the process of iteration engaged in should make us optimistic that if Butterfield is wrong, we'll have the tools to find this out.

Thus, we can see how iterative strategic perspectivalism generates a self-correcting investigative process that can uncover past worlds in spite of loss. Before turning to more philosophical discussion, I want to finish with a brief discussion of scientific laws.

4. Coda: Scientific Laws

In chapter 8, I'll discuss the "peculiarity" of the past—approximately, its contingency—but our discussion of anchoring raises questions about laws. If various anchors can be understood as tokens of historical or ahistorical types governed by the same set of rules, then shouldn't we think of those rules as putative laws of nature? In this coda, I'll briefly discuss the *epistemic* features of laws as they play out in strategic perspectivalism, leaving metaphysical discussion for that later chapter.

A fairly common refrain among historical scientists is that their inferences about the deep past are underwritten by invariant scientific laws, particularly those revealed by chemistry and physics. Might these provide foundationalist anchors, playing necessary or in some sense privileged roles in the science of lost worlds? I don't think so, and will try to articulate why here. Note that the central point of this discussion doesn't turn on whether laws generally hold, but rather where and how they are applicable.

Philosophers often understand scientific or natural "laws" as *non-accidental, universal* regularities (Carroll 2020). That is, a type-level process (rangeomorphs grow via a three-branch structure, not *this* rangeomorph grows via a three-branch structure) that ranges over all relevant systems (fractal growth leads to self-similarity, not self-similarity is found in some fractally-growing things), that possess those properties for some appropriate reason (geometric relationships, say, ensure the relationship between fractal growth and self-similarity, rather than it just happening to be the case that surviving fractally-growing things exhibit self-similarity). Such regularities, if only we could find them, were taken to play an important role in science because—at base—they enable prediction (if we find a rangeomorph, it will grow fractally) and they enable explanation (rangeomorphs have self-similarity because they grow fractally).

One philosophically exciting idea concerning loss connects it with laws of nature. If a law can be "lost" or a law can "arise," then doesn't this imply that the laws of nature could be (in fact have been) different? If a law of nature is some kind of metaphysical necessity—true in all metaphysically possible worlds—then laws cannot be lost and cannot be found. It might be that various perimeters within the laws differ across worlds or even times, or vary across "initial conditions." Marc Lange distinguishes between *temporary laws* and *eternal but time-dependent laws* (Lange 2008). The former are laws that only hold for some period, while others are laws that come "online" once the right conditions arise. The practice-oriented discussion here cannot distinguish between these, but I consider this a feature, not a bug.

We don't need full-blown laws to predict or explain (Mitchell 1997), so we shouldn't get too concerned about their availability or otherwise in science. However, I want to consider a fairly common refrain from scientists seeking to understand putatively lost worlds: that appeals to scientific laws will ground their discoveries and justify their investigations. Here's an illustrative example from the introduction to Maynard-Smith and Szathmary's *Major Transitions*:

There are obvious difficulties in discussing unique events that happened a long time ago . . . [however] we have agreed theories both of chemistry and of the mechanism of evolutionary change. We can therefore insist that our explanations be plausible both chemically, and in terms of natural selection. This places a severe constraint on possible theories . . . Further, theories are often testable by looking at existing organisms. (Maynard-Smith and Szathmary 1997, 3)

Maynard-Smith and Szathmary are proposing a general strategy for understanding a particular set of important changes in the history of life, “major transitions.” These changes involve shifts in the transmission of ancestral information or shifts in biological individuality. A classic example is the evolution of multicellularity (see chapters collected in Calcott and Sterelny 2011). The investigative strategy proposed could be understood as a kind of combination of anchoring and artificing. Maynard-Smith and Szathmary’s proposal is to understand major transitions in terms of the abstract mechanics of population genetics and similar modelling approaches. Roughly, build an agent-based model that explains under what evolutionary conditions “de-Darwinization” occurs, for instance, an evolutionary process that transitions from a bunch of “selfish,” competing single-celled organisms to an organism consisting of cooperating parts. The models must be consistent with evolutionary and biochemical laws and tested on extant model organisms.

Such a model would then be applied across those major transitions. This is an ambitious project, no doubt, but what matters here is what Maynard-Smith and Szathmary appeal to in constraining their project: it must fit within the theories of chemistry and evolution. That is, they must appeal to fairly generally applicable laws. One way of understanding the putative role of laws here is through what Khalifa and Goldberg call *epistemic outsourcing*: “The phenomenon whereby results taken as basic or given in one scientific field are justified owing to the intellectual efforts, evidence, and methods in another scientific field” (Khalifa and Goldberg 2022, 383). Epistemic outsourcing is *social* and *foundational*. The former insofar as it relies on (something like) the testimony of the outsourced scientists, and the latter insofar as it is taken as “basic,” that is, unquestioned, by the outsourcing scientists. If laws from physical (and other) sciences are epistemically outsourced by historical scientists, then we should see them playing foundational, taken-for-granted roles

in their reasoning. They might be applied, for instance, in an “off-the-shelf” manner to their particular purposes.

No doubt, well-verified robust regularities from physics and chemistry (but also from evolutionary biology, ecology, physiology, etc.) play critical roles in uncovering lost worlds. Mitchell et al., for instance, use fairly straightforward models of spatial niche-differentiation; for another example, the geometric growth models we’ve seen (of *Dickinsonia* and rangeomorphs) follow extremely general mathematical rules; and for another example, we’ve just seen how crucial various basic biochemical dynamics are for Butterfield’s account of the arrival of complex metazoans and the biological pump. But to think that, say, Butterfield’s investigation is a primarily biochemical one is to miss the epistemic trees for the forest. As he puts it: “All biological exchange ultimately depends on chemical diffusion, but it is the associated fluid-dynamic context that determines physical and ecological properties” (Butterfield 2017, 10).

Models of chemical diffusion are crucial, but these must be put in the context of an aquatic medium, which itself—as Butterfield argues—is deeply affected by physiological and anatomical features of the organisms living within those environments. So, no doubt, what we might call “laws,” or at least models of them, are important perspectival tools, but I don’t see them as playing a privileged or necessary role. That is, there might be some level of epistemic outsourcing going on, but the deployment of that knowledge is not foundational.⁴

Further, we might imagine Mitchell et al.’s work as being about the “laws” of ecology. That is, there is a set of “rules” that range over all living ecosystems. And in a sense, their work supports this idea: the same model that should lead us to expect niche-driven spatial differentiation in most ecosystems should also lead us to expect neutral-driven differentiation in cases where competition is limited, assemblages are frequently disrupted, and so on. But notice the role of “laws” here. While Mitchell et al. certainly framed the discussion—set up an interesting challenge—it was via an anchoring strategy, the close attention to contextual knowledge about the Avalonian, as well as

4 Khalifa and Goldberg’s account of “foundational” is about the epistemic status of propositions, approximately whether they are questioned within the relevant epistemic practice. I’m here less interested in the epistemic status than I am in the epistemic *function* of foundations. Even if some scientific propositions are unquestioned by historical scientists (so they are “basic” in Khalifa and Goldberg’s sense), they do not—or at least do not typically—play a basic or foundational role in their reasoning.

various analogous cases, that progress was made (see Sánchez-Dorado 2024a for an emphasis on iteration between various “idiographic” and “nomothetic” models in the Earth sciences).

The lesson here is that models of highly abstract, robust, and general regularities (or “laws,” if you must) can play important roles in both identifying and solving challenges from loss. But it is worth noting that their very abstractness can be limiting. They afford a perspective that allows unifying very many targets—as we saw in Mitchell et al.’s case—but if opportunities for iteration between perspectives are limited, especially without a rich trace record and investigation of diverse analogies, this in itself is not particularly productive. This leads me to be somewhat sceptical of the kind of investigative strategy that Maynard-Smith and Szathmari suggest. Although laws and extant animals can teach us much, without an actual anchor into the deep past (provided by, say, fossils), I’m not confident that we’ll see as much substantive progress (Currie 2019d).

So, I want to suggest it isn’t laws per se, but the adoption of strategic perspectives and iteration that allow us to discover much about lost worlds. Our models of highly general regularities, to be employable, must often be made bespoke: carefully tailored to fit the local, idiosyncratic contexts that must be grappled with if we are to understand the deep past.

Prof. Ichthy takes a deep breath. “So,” they begin, “earlier I was worrying, well, at least feigning worry, about whether my work on Broadfoot amounts to mere speculation. I hope now to have convinced you—or at least to have convinced myself!” They give a self-deprecating chuckle. “That the worry is unjustified. Even though Broadfoot is a strange creature, unlike anything we’ve seen before, we are nonetheless able to piece together this ancient lineage. So, how do we do it?” Once again, Prof. Ichthy’s use of ‘we’ is a self-conscious nod to colleagues and students.

“I’ve emphasized how important taking a *perspective* on Broadfoot is. If something looks unique, then, think of ways in which it might not be unique. Broadfoot is a unique species, that’s for sure. But they aren’t the only mammal, they aren’t the only swimmer, they aren’t the only venomous creature, they aren’t the only vertebrate, and so on. These perspectives let us link Broadfoot with an array of other things, be they living analogues, their surviving relatives, the things we build, and contextual factors from the environment Broadfoot lived within. And with each of these new perspectives and links, we get more

opportunities for bringing our knowledge into conversation, iteratively working our way closer and closer to a rich and, I think, plausible picture of this weird, wonderful little animal.” Ichthy pauses, perhaps considering how exactly to finish the lecture.

“I mentioned earlier that one reason to care about Broadfoot and the science of Broadfoot is that thinking about the past can be transformative of how we think about the present and the future. And I can say I’ve learned a lot thinking about Broadfoot and the other mammals they shared that past world with. Although I only see that past world in parts, it provides such a broad perspective—a grand vista—from which to think about our world and how changeable and how fragile and, yes, how valuable it is.”

Ichthy pauses again, their flipper resting fondly on the case containing the skull.

“Were it not for this specimen, I could never have imagined a world with Broadfoot. And I’m so glad that I got to, because the world I live in is so much the richer for it.” Ichthy smiles a little bashfully as they indicate the end of their lecture with a small nod.

