

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>

The Avalonian Ediacaran: Lost, but Not Erased

“To see what is so challenging about the science of Broadfoot,” Prof. Ichthy says, emphasizing “challenging” with a wave of their flipper, “we can compare it to the science of the ancestors of more-familiar fauna.”

With another transparency switch, Prof. Ichthy’s grad student reveals a set of reptilian beings well-known to the audience. As with the mammals before them, these are linked both to each other and to the deep past via lines of ancestral descent. Where the lines meet to form ancestral splits, there are speculative illustrations of the precursors of modern forms.

“If scientists want to make sense of the fossils of, say, Ichthyosaurs, well, we have modern examples swimming all around us.” Prof. Ichthy indicates themselves and their audience.

“We can look at these modern examples to test and generate ideas; we can use our understanding of evolutionary history to infer from these modern forms to those ancient forms. We needn’t speculate, as evolution has done our speculation for us. In these circumstances, there is a mutually beneficial, mutually informing relationship between our knowledge of contemporary ichthyosaurs and our knowledge of past ichthyosaurs. Under these conditions, because we know so much, a single incomplete specimen—a single tooth, say—is enough to reconstruct a millions-year-extinct ancestor with confidence. Sadly, this is not so for Broadfoot . . .”




ILLUSTRATION 2.1 "Then Might Those Genera of Animals Return"

Ichthy pauses, glancing at their student, who quickly returns to the first transparency, again projecting the specimen and subsequent reconstruction, and Ichthy continues.

“Broadfoot diverges so radically from living mammals that *now* is, on the face of it, little use in understanding *then*. One of Broadfoot’s teeth would tell us it is a mammal of sorts, and maybe tell us something of its size. Our understanding of fossilization processes, more generally, can get us some of the way. However, without the mooring of Broadfoot-like organisms around us, we are cast adrift . . .”

Ichthy stops for a moment, as if considering whether the “mooring” metaphor fits with the allusion to being “case adrift,” before nodding to themselves and continuing. “In this instance, we are deeply lucky to have such a richly preserved specimen.”

Ichthy pats the skull’s case with a fond flipper.

“Broadfoot, we’ve discovered, has a plethora of unique features: from bill to venom to electrosense. When Broadfoot went extinct, so too disappeared those unique traits.”

Prof. Ichthy repositions their spectacles further down their snout, glaring over them, out at the audience, leaning over the podium with a dramatic air.

“Broadfoot is lost. Lost, but not erased.”

In the first chapter, I provided two motivating thoughts, one historical and one philosophical. The former highlighted the historical importance of establishing a non-human past that differed from the present, while the latter highlighted a philosophical focus on traces in analyses of palaeoscientific knowledge. In this chapter, I’ll provide a conceptual framework that makes sense of what recent philosophy has missed, and of what 19th-century earth scientists were trying to establish. I’ll develop a notion of “loss,” that is, circumstances where some past entities (understood very broadly) no longer exist in the present. For example, from Prof. Ichthy’s perspective, Broadfoot is lost to extinction. And I’ll distinguish “loss” from “erasure,” that is, processes by which traces degrade and disappear. For example, Broadfoot survived erasure by virtue of its fossil’s preservation and discovery. In short, philosophers have missed the importance of loss, focusing instead on challenges from erasure, and early palaeoscientists, faced with traces of a surprising past, wanted to establish that loss had occurred.

Above, Prof. Ichthy compared the investigation of Broadfoot to past ichthyosaurs. Let's consider an analogous question: What is epistemically different about investigating *dinosaurs* and investigating *mammoths*? A natural answer considers these animals' traces. Non-avian dinosaurs are known, shall we say, mostly "geologically": their remains primarily take the form of funny-shaped rocks, mostly capturing previous dinosaur anatomy and morphology, as well as the occasional eggshell or trackway. The mammoth record, by contrast, includes much more than funny-shaped rocks. For mammoths, we have truly biological remains, stored and preserved in nature's arctic freezer. We have their hair, their tusks (sometimes even carved by the hands of our ancestors), and even their DNA. The comparatively short temporal distance between ourselves and mammoths affords a much richer record than that of dinosaurs. All of that is right, but it misses—I want to argue—something in the vicinity of half the story.

A major difference between studying dinosaurs and studying mammals is . . . (drum roll) *elephants*. We have currently existing examples of things extremely similar to (and closely related to) mammoths, and these epistemically and causally link the past to the present in rich ways. As Shoshani have put it:

Unlike researchers who study dinosaurs and have to rely on distant living relatives, such as crocodiles and birds for indirect comparison, students of Proboscidea are fortunate to have living models with which to compare and contrast the extinct taxa directly. This advantage should not be underestimated, because knowledge of the living species helps one to interpret observations of extinct forms and vice versa. (Shoshani 1998, 487)

In the parlance I'll soon unpack, it isn't simply that dinosaurs are more *erased* than mammoths; dinosaurs are also more *lost* than mammoths. In this instance, this is largely due to the ancestrally retained similarities between extinct and extant pachyderms. Mammoth and elephant are both members of the clade *Elephantidae*, and, in this case, ancestral relations have preserved many similarities between them.¹ Although the clade *Dinosauria* is not lost insofar as birds persist, in this case, much less similarity has been retained.

1 Especially so for the Asian elephant: Molecular evidence strongly suggests that African species form an outgroup from mammoth and Asian elephants (e.g., Baleka et al. 2022).

So, many properties possessed by, say, *Diplodocus* are lost even though many others, for instance, their “pneumatized” hollow bones, have not.

I want to articulate the distinction between loss and erasure by discussing two cases from the Avalonian Ediacaran. So, providing relevant information about these is my next task. I’ll then turn to the concepts themselves and set us up for the remainder of the book. This will involve applying the framework to the two motivating thoughts from the introduction, and finally, characterizing the epistemic challenge generated by loss.

1. The First Metazoans

It is commonly thought that metazoan life, that is, complex multicellular life, really got going in the Cambrian, a little under 540 million years ago (Briggs 2015). From that time, animals are largely as we know them today: critters with sensory organs, mouths for consuming, anuses for expelling, methods of locomotion, occupying roles including predators, prey, parasites, scavengers, herbivores, carnivores, browsers, grazers, and so on. But metazoan life didn’t start in the Cambrian. Fossils record a pre-Cambrian history stretching backwards into the Ediacaran. The earliest examples (which we’ll concern ourselves with here) are known from something like 600 million years ago, and that biota² has a striking set of differences from their Cambrian replacements (Narbonne 2005; Narbonne et al. 2012; Droser and Gehling 2015; Dresow 2025). These differences are most stark in the earliest deposits: the Avalonian Ediacaran (Liu, Kenchington, and Mitchell 2015). I’ll provide some general information about the Avalonia before turning to two initial examples of loss.

1.1 Avalonia


The “Avalonia” is used to both pick out a past geographical location—now dispersed across England, Poland, and the northeast of the United States and Canada—as well as identify particular fossil assemblages hailing from the Ediacaran. The assemblage and timespan I’ll refer to as “the Avalonia” are known from deposits in Charnwood Forest, Leicestershire, and in Canadian Newfoundland. A crucial spot is Canada’s Mistaken Point Ecological Reserve

2 There is no happy term for these Ediacaran organisms: terms aligned with particular taxa (“animals”, “metazoans”) imply phylogenetic assumptions, “biota” implies something more unified than some are happy with. Perhaps the least controversial would be “progenitors of soft-bodied macrofossils” but that’s such a mouthful that I’ll just stick to terms like “biota” and hope the reader is sufficiently charitable.



FIGURE 2.1 Mistaken Point (Barrett & MacKay Photography, © Mistaken Point Ambassadors Inc., BYND)

FIGURE 2.2 Rangeomorph Fossil
(Detail from Fig. 2, Liu et al. 2015).
Creative Commons CC BY



UNESCO World Heritage Site, a series of dramatic cliffs around two hours' drive from St John's. There you'll find a bunch of fossil beds (on a guided tour, you can't just rock up on your own), which, despite the beautiful location, to the uninitiated might not exactly inspire fossil-fuelled awe (see Figure 2.1).

What are we looking at in Figure 2.1? Most familiar fossils are preserved “hard parts”—shells, bones, and teeth—which have either been preserved or “replaced”: approximately, bones are held in place, retaining morphological information while their carbon matrices are, via water's slow work, replaced with other minerals (Watkins 2024a). By contrast, Mistaken Point's assemblages consist of the imprints of soft-bodied critters, which have been preserved through volcanic silt settling atop layers of mud. Something like 575 million years ago, this was a *benthic*, that is, seafloor, community, far below the reach of light. The denizens of this community left imprints in the sea floor's mud. These imprints were preserved by being coated in repeated layers of sediment rich in volcanic ash.

What kind of world do sites like Mistaken Point reveal? An odd one. The most common fossils are rangeomorphs, “bizarre frond-, spindle-, bush-, or comb-shaped” (Narbonne 2005, 426) creatures basically consisting of a central organic skeleton that supports a set of fractally growing “fronds” (see Figure 2.2). Some fossil forms appear to have been stuck to the seafloor, standing “upright”, while others appear to lie flat.

The seafloor that rangeomorphs attach to is coated by a microbial mat, several centimetres thick (Kolesnikov et al. 2017). This relatively rigid structure both sealed the seafloor to erosion and gave the Ediacaran denizens something to stick to (Seilacher 1999). The ocean is significantly less oxygenated and less homogeneously oxygenated than our own. Charlotte Kenchington has produced a haunting image of the fauna of Avalon, the rendition lit by an impossible light source (Liu, Kenchington, and Mitchell 2015, Figure 2.3). The image uses Emily Mitchell's maps (see below) to represent the organisms' size and spatial position based on data from the fossils *in situ*.

A surprising and, as we'll see, ecologically fundamental feature of Avalonian metazoans is their lack of mobility: the rangeomorphs are interpreted as firmly anchored to the seafloor. In later Avalonian deposits, there are features fairly interpreted as indicating a shuffling movement by some unknown taxa (Menon et al. 2013), and there is evidence of something similar at one contemporaneous site (Liu, McIlroy, and Brasier 2010). While there are always some mysterious geological features acting as centres of interpretative dispute, Liu et al. are firm on Avalonian sessility: “Rangeomorphs,



FIGURE 2.3 Charlotte Kenchington's Avalonian Scene, Depicting Upright and Flat Rangeomorphs. Note the Clever Use of a Positioned Light Source, Illuminating a Scene That Would Have Never Seen Light (Liu et al. 2015). Creative Commons CC BY

arboreomorph/frondomorphs, and the majority of the miscellaneous taxa abundant on these bedding planes are still interpreted as sessile and immobile” (Liu et al. 2015, 1370). The complex multicellular life of Avalonia was incapable of shifting from their spots, very different from the majority of metazoans today.

We’ll be returning to the Avalonia, and the Ediacaran more generally, repeatedly throughout this book, so I’ll save fuller discussion of what we know and don’t know about these critters, and of various hypotheses regarding their feeding strategies, development, phylogenetic alignment, environment, and so forth, for when this is demanded. For now, and to motivate our conceptual engineering in the next section, I’ll recount two stories, both of which highlight apparently unique features of the Avalonian world.

1.2 Pizza Disks

A long-standing puzzle from Avalonia concerns not rangeomorphs but another set of fossils: ivesheadiomorphs or “pizza disks” (Kenchington and Wilby 2014). Ivesheadiomorphs are mostly oval in shape (although they can be more distended), with very little internal detail other than raised lobes (Figure 2.4). They are distributed fairly randomly across Avalonian beds and are, as fossils, frustrating. Part of the frustration is that ivesheadiomorphs are insufficiently alike: there is a distinct lack of regularity across their structure compared to more well-behaved fossil taxa. They are often asymmetrical, and their lack of internal structure makes reconstructing their morphology and phylogeny a baffling endeavour. As with many of Avalonia’s denizens, they’ve been interpreted in a bewildering variety of ways. In 1979, Boynton and Ford interpreted one of the pizza disks, *Pseudovendia charnwoodensis*, as a kind of arthropod, grouping them with other animals sporting chiton exoskeletons, while other interpretations took them to be cnidarians. More recently, Laflamme et al. (2012) argued they were the remains of microbial colonies, while Wilby et al. (2011) suggested their irregularity was due to the differential effects of organisms collapsing in the bed above. In this discussion, I’ll focus on Alex Liu and his collaborators’ explanation, which seems to be generally accepted.

Liu et al. (2010) propose an explanation that ties ivesheadiomorph fossil forms to the putatively unique (indeed—perhaps *lost*) conditions of the late Ediacaran. They draw our attention to differences in decomposition between then and now.

Nowadays, an organism’s death creates a scavenger smorgasbord. Were I to expire, the pattern of my decomposition would be influenced by the activities of various multicellular organisms: a retinue of bugs and creepy-crawlies would find me, seeking a food source (as potentially might my cat), various humans would likely move my decaying body about, perhaps much to their consternation (or at least complaining about the smell), and so on. But then-a-days, in the multicellular, sessile Avalonia, death and decay are different. Death in the Avalonia involves no such smorgasbord for metazoans, as they cannot get to you: Decay is left to single-celled organisms. As such, decomposition would take longer, potentially leading to differing patterns of decay. Add to this the possibility of the microbial mat acting to hold together partially decayed critters, and the less oxidized ocean, and we generate a decomposition pattern potentially unique to the Avalonian: “In such a scenario,

FIGURE 2.4 “Pizza Disk”

Ivesheadiomorph (Fig. 1 from Liu et al. 2011). Reprinted with Permission from *Paleontology*



decay processes may have taken considerably longer than is observed in modern marine environments, implying that organic matter in the form of dead organisms could have been a common phenomenon at the Ediacaran sediment-water interface, and that impressions of these deceased organisms would have been preservable” (Liu et al. 2015, 1364).

The thought, then, is that ivesheadiomorphs are not true taxa after all, but instead a set of different lineages (presumably rangeomorphs and their allies) that have expired, decayed, and been preserved under Avalonian circumstances. Thus, “Ivesheadiomorphs and similar forms represent a morphologically disparate grouping of fossils, linked by their having experienced similar taphonomic processes rather than sharing true biological relationships” (Liu et al. 2011, 610). That is, it isn’t that they are ancestrally linked, but instead that they decayed and were fossilized by a similar process. This matters not only for understanding organismic decay, but also the ecosystems

themselves. As Liu put it in his doctoral thesis: “The late Ediacaran Period, at the time the Avalonian strata were deposited, is considered to be ecologically non-uniformitarian, in that there is currently no evidence for the pervasive bioturbation, macro-predation, or scavenging processes that rapidly break down and remove organic material in Phanerozoic ecosystems” (Liu 2011, 117).

If Liu is right, the processes of decay that produced the pizza disks are lost; that is, they only occurred due to specific, long-past Avalonian conditions. Let’s turn now to a different case.

1.3 The Ecologies of Avalon

Here’s another tale of Avalonia, also likely related to the lack of metazoan mobility. I mentioned that the Avalonian deposits, particularly those from Newfoundland, were formed in volcanoclastic events. That process of preservation is well-understood, and there is good reason to think they preserved the actual spatial relationships of living rangeomorphs. Treated carefully, then, these fossil beds act as spatial maps of Avalonian ecosystems. This is exciting because it enables the application of models also used to understand contemporary ecosystems, thus allowing us to quantitatively compare then with now.

Emily Mitchell and her team took careful scans of the fossil beds using laser-line probing technology, thus providing a digital representation of the fossils and their relative spatial relationships (Mitchell and Butterfield 2018). This enables the application of *Spatial Point Analysis*. Imagine a simple Euclidean two-dimensional space divided into space points. Now imagine randomly generating a scatter of occupied points across those dimensions; say, for each location, on each run, there will be a 50/50 chance of that location being occupied. Running this over and over and summarizing results will generate a kind of neutral model, a profile that represents the occupation patterns we’d expect if the location points were random. We could then compare this to the spatial distribution of a real ecosystem. Insofar as the observed pattern diverges from the null, we can infer that the ecosystem is structured in some non-random way, presumably by patterns of breeding, competition for resources, predation, and so forth.

Modern benthic systems (and other contemporary ecosystems) are not neutral in this sense: the kind of critter you are, and what other critters are about, make a difference to your spatial position. This is unsurprising, I take it, because ecological relationships—competition, mutualism, predation, food availability, etc.—all shape organism distribution across a modern ecosystem.

That is, spatial relationships are differentiated by *niche*. So, we'd expect the same for the Avalonian, right? I imagine you see the answer coming.

No.

Avalonian ecosystems appear to be neutrally structured, which is to say, the patterns Mitchell et al. generated converge on the null generated by spatial point analysis. Thus, Avalonian ecosystems appear to be dominated by neutral processes: What you are doesn't seem to influence much about where you are. It's worth noting how surprising this is from an ecological perspective. Ecologists examining ecosystems from the past 500 million years find niche-differentiated spatial arrangements, and no wonder, given the ubiquity, complexity, and importance of trophic relationships in metazoan life. Discovering a widespread, long-surviving assemblage of fauna that does not differentiate in this way suggests that a very different ecology is at play. Indeed, Mitchell and company consider a perhaps radical possibility emerging from this result:

This study provides strong evidence that neutral processes structured Avalonian Assemblage palaeocommunities . . . These neutral-process-dominated community dynamics contrast with those observed in the modern marine realm, where neutral processes are typically rare across multiple spatial scales . . . 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)

That is, if Avalonian ecosystems are uniquely neutrally dominated, might this mean that the usual rules of ecology don't apply? Presumably, this neutral ecological structure has something to do with the rangeomorph's sessility. We have remarkable records of the spatial relationships of Avalonian ecosystems, and we can infer their neutral domination. In light of this, Mitchell et al. ask whether neutrality is due to run-of-the-mill ecological dynamics acting in unusual circumstances or if Avalonian systems, with their wildly different organisms, environments, and so forth, simply play by different ecological rules (you'll have to wait until chapter 5 to find out Mitchell et al.'s answer . . .).

This dip into the science of the Avalonian Ediacaran yields two interesting cases of what I'll soon call "loss." First, metazoan immobility combined

with the microbial mat and a less-oxygenated ocean generated never-seen-since patterns of decay and decomposition. Second, at least in terms of spatial composition, Avalonian ecosystems exhibit a near-unique neutral ecological structure. In the next section, I'll draw on these examples—and the more toy-like elephant case—to provide a conceptual scheme for thinking about these issues.

2. Loss and Erasure

“When trying to make sense of Broadfoot, we are doubly challenged.” Prof. Ichthy raises one flipper, “First, we’ve only one specimen, a paltry record.” They raise another flipper, “Second, there’s nothing like Broadfoot around for us to study.” Ichthy lowers their limbs while a grin spreads across their maw, “Happily, this kind of scientific challenge is exactly what gets me excited . . .”

As we’ve seen, the difference between studying a mammoth and studying a dinosaur can be partly captured in terms of their differential preservation: for a mammoth, we have examples preserved in ice, samples of hair, and tusks; for dinosaurs, we have fossilized bodies, eggshell, and trackways. But attending to traces alone misses a crucial factor: elephants. A focus on differential preservation doesn’t quite capture our two examples from the Ediacaran, either.³ We have a fair number of ivesheadiomorph remains, and of great variety, but they are tricky to understand because they don’t resemble anything we’ve seen before. Something similar could be said of the ecology of the Avalonian: We have extremely well-preserved spatial relationships, but an unfamiliar ecology. In these cases, the problem isn’t a lack of traces. So, *prima facie*, attending to traces alone won’t capture the challenges these scientists face, suggesting that philosophical focus on traces has been myopic. Further, as we saw in the first chapter, early geologists and palaeontologists emphasized their discoveries’ importance by leaning on the idea that the past was crucially different from the present. It is time to provide some philosophical clothes for these ideas.

3 Recently, Max Dresow has suggested that a major difference in understanding the Cambrian and Ediacaran biotas is that the former has unusual features, but the latter has insufficiently preserved features. In my parlance, the Cambrian is lost while the Ediacaran is erased (and so might be lost, for all we know). This doesn’t undermine my use of the Ediacaran examples; it is an excellent way into considering how loss and erasure interact (Dresow 2025).

These clothes will be largely stipulative. I'll begin by defining a "past world," and draw on that to further define "erasure" and "loss." These latter notions are intended to capture the epistemic challenges highlighted by bad preservation, on the one hand, and unfamiliarity, on the other. There'll be some philosophical finagling, but I hope the payoff throughout the rest of the book will justify it.

So, on to the notion of a "past world." First off, notice that "world talk" is evocative of a few different things. For many philosophers, it likely brings to mind David Lewis' analysis of modality. For others, it might raise a set of looser conceptions, "life worlds," or something similar. To an extent, my use of "worlds" trades on this ambiguity, but given that we'll be discussing possibility at length later in the book, I think it's worthwhile to connect my use of "worlds" to Lewis'. I do think that this machinery could be jettisoned, but I also think it brings some useful precision to the discussion.

Historical scientists are nominally interested in understanding some set of things in the past (although in later chapters I'll argue that they should have something much more modal in mind). Let's call that set of things a *Past World*.

Past Worlds are subsets of the actual past; that is, they are *actual* but not *complete*. They consist of a set of temporally and spatially connected *denizens*: entities, processes, dynamics, relations, and properties, although which denizens the past world consists of is partly a matter of convention (that is, what scientists are interested in).

Let's tackle those italicized terms one by one.

What is it to say a world is *actual* but *incomplete*? We can understand a *world*, according to Lewisian analysis, in terms of spatial-temporal relations (Lewis 1986; Nolan 2015). If there are any temporal, spatial, or causal relationships between any two things, they are part of the same world. My left hand and the first rangeomorph are in the same world. However distant they may be, some spatial and temporal relationship exists between them. For instance, temporally speaking, my hand and the rangeomorph are something in the vicinity of 575 million years apart. If two objects lack such spatio-temporal relationships, then they occupy different worlds. For example, it might be possible that were true taxa after all (perhaps a kind of arthropod). A possible world semantics would let us understand this claim by saying that in at

least that one world, there are arthropods.⁴ There is no spatial or temporal relationship between true phylogenetic taxa and false taxa, so they occupy different worlds.

Worlds are *complete* in the sense that they provide an answer (that is, satisfy the truth conditions) for any meaningful proposition. Assuming that Liu et al.'s hypothesis is correct, presumably "are false taxa" is made true by the individual fates and decay of organisms in the Avalonian Ediacaran. If it is possible that ivesheadiomorphs might have been arthropods, then there will be a set of non-actual worlds containing arthropod ivesheadiomorphs. All worlds (according to Lewis, anyway) are possible worlds, and the actual is a possible world which is, well, here: *gestures around*.

Past worlds are subsets of the actual world. That is, they may contain some parts of the actual world, but do not contain all of it; hence, they are actual but not complete. They are, if you like, a bit of the world. So, when I say a past world is *actual*, I mean it is part of *our world*, the one we occupy. When I say it is not necessarily *complete*, I mean that it does not include every aspect of our world. It doesn't contain all it needs to ground all those truth conditions, only some of them. A past world is typically picked out or indicated by historical scientists via a set of *denizens* (or the traces of denizens).

I'm using the term "denizen" as a kind of catch-all, intended to include whatever kind of ontology you'd like. Your metaphysical inclinations may lead you to admit only simples and the void, or processes or structures, or entities with special causal powers, or whatever. However you cash it out, your metaphysical view will nonetheless need to (even if pragmatically) make sense of individual things such as particular rangeomorphs; processes such as the decay that produces ivesheadiomorph fossils; dynamics such as the neutral-driven ecology that structures Avalonian fossil-beds; relations such as those between individual rangeomorphs and the microbial mat they are embedded in. I hope I can remain neutral regarding matters concerning our views on fundamental ontology: at base, "denizens" are whatever entities, processes, dynamics, relations, and properties that constitute the past world.

There are potential downsides to the term "denizen," implying as it does a kind of "thing" or entity, as opposed to a process, say, or a set of dynamics. But I doubt any term will avoid bias in one direction or another, and I quite like "denizen." Again, I'm aiming to be as metaphysically relaxed as I can. If you

⁴ I'm not primarily interested in developing anything like a temporal analogue of modal semantics here (e.g., Rini and Creswell 2012).

believe that ultimately there are only relations, not *relata*, that is fine. For you, any entities and the like are merely derivative of those structures or processes, and vice versa for anyone committed to a more “substancy” metaphysics.⁵

When historical scientists tackle the Ediacaran, say, they more or less indicate a set of taxa, or an ecosystem, figuratively saying “I’m interested in *that* set of denizens.” Perhaps more carefully, they indicate a set of specimens—traces—and thus indicate a set of denizens causally upstream of those specimens. So, what counts as being included in the world is at least partly indicative. This has two important consequences. First, there are *a lot* of past worlds: as many as there are denizens and combinations of denizens. Second, which world we are talking about is highly sensitive to description, depending on which world has been “indicated.” The Avalonian Past World, for instance, is populated by rangeomorphs, the bacterial mat, processes of decay, and so on, but what exactly is included will depend in part on scientific practice.

Although some denizens are explicitly included in a past world, others might be included in less direct ways. Here, we can take our cue from the semantics of fiction. Some truths in a work of fiction are explicit, while others are implicit (Lewis 1978). Prof. Ichthy is an Ichthyosaur, has graduate students, is delivering a lecture, and so on. These are true in the fiction due to being explicitly stated in the text. But we might also infer from Ichthy being a professor that there is a university to which they belong, that marine reptiles can be PhD candidates, and so on. These truths are implied by both what is explicit in the fiction and our background expectations. Although past worlds are not fictions (or at least I’m not committed to them being so), there are also implicit denizens that, to some extent, fill the world out. For instance, although less complex multicellular organisms rarely figure in reconstructions of rangeomorph ecosystems, they likely haunted these worlds, and sometimes become explicit denizens when investigation demands it. The flexible boundaries of past worlds, as we’ll see, are a crucial factor in how historical scientists investigate them. I’ll call using this flexibility *strategic perspectivalism*, and it’ll be introduced in the next chapter.

You might worry whether I’m talking about worlds as scientists represent them or worlds in themselves. In the past, some palaeontologists spoke of an Avalonian past that included a true taxon identified as “ivesheadiomorphs”; now they do not. It is helpful to distinguish between *putative worlds* (worlds as

5 And while I’m at it, if you’ve no metaphysical time for this “possible worlds” talk, don’t despair: I’m only referring to the actual after all. Well, don’t despair yet . . .

scientists represent them) and *true worlds* (the worlds themselves). Assuming we've got something right about the Avalonian, the true world doesn't contain an ivesheadiomorph taxon; that was simply a putative world. All worlds as scientists represent and discuss them, of course, are putative worlds, and they are true worlds when those representations are veridical or sufficiently so. Most of the time this distinction won't matter, but I'll wheel it out when necessary.

Okay, we have a notion of a *past world*: a set of denizens, indicated by scientists. Given some past world, what kinds of features make a difference to our epistemic access to it? One set of factors refer to our knowledge, technologies, and scientific skills and capacities. Another set, and what concerns us here, are at least partly features of the world itself. Recall the contrast between dinosaurs and mammoth. The latter has more and richer *traces*, downstream causal remnants. But the latter also has *elephants*: the denizens, insofar as they are members of Elephantidae, still exist. Dinosaurs, in the terms I'll now introduce, are both more "erased" and more "lost" than mammoths.

Some past world is *erased*, relative to our current world, insofar as either (1) there is no downstream causal information of its denizens or (2) that information is not retrievable in the present.

Erasure is a partially epistemic notion. It may be that some past denizens literally leave no causal information (as Havstad suggested in chapter 1). Such denizens are permanently erased. If information persists, we need sufficient know-how and luck to be able to discover and make sense of it. As such, erasure is relative to our current knowledge states. Something might have been erased, but due to some new technological innovation or fossil discovery, become unerased. In this sense, erasure is not necessarily permanent. Dinosaur muscle does not leave direct causally downstream remnants in quite the same way as mammoth muscle sometimes does, although it is not completely erased, as we can infer it from features of preserved bone (compare with my definition of "trace" in Currie 2018a, chapter 3). Erasure, then, turns on whether causal information has been preserved and whether we can interpret it. Let's now turn to loss. A world is *lost* insofar as its denizens do not exist in the present.

Although mammoths and elephants are differing taxa, both are members of Elephantidae, and insofar as we are interested in understanding mammoths, understanding them as members of that still-existing group is an enormous boon. For dinosaurs, we have still-existing birds—the clade Dinosauria still exists; however, that is often not the aspect of the past world that interests scientists (that is, they don’t pick out those denizens). Lineages like the enormous sauropods, the remarkably head-gear-ed ceratopsids, or the larger predatory theropods have left no direct descendants, and birds have diverged from them dramatically. Of course, some aspects of mammoths—members of Elephantidae adapted to Arctic-like conditions, for instance—are lost. So, whether loss has occurred is in part dependent on how we characterize denizens. Let us, then, define loss in terms of denizens:

Some denizen is *lost* if it doesn’t currently exist or doesn’t exist for all we know.

So, loss is a property held by a past world relative to now (in principle, it could be relative to other times as well), which is determined by taking the set of denizens constituting the past world and checking which are now present.⁶ We’d call a world “lost” in case some (or enough) of its denizens are lost, or if sufficiently significant denizens are lost. Similar to how our characterization of denizens is sensitive to context—at minimum, what scientists indicate—whether a world is lost is also sensitive to context.

A fairly straightforward way to understand loss and its context-sensitivity can be drawn from the philosophical literature on similarity (e.g., Goodman 1972, 437–446).⁷ Let’s take a very broad conception of how we might group denizens: say, two denizens can be grouped so long as they are similar (I’ll expand on kinds below). This proposal runs into an immediate problem. Without any limitations on similarity, every denizen is infinitely similar and dissimilar. For instance, a rangeomorph and my hand are both similar insofar as neither is the moon, neither existed on January 27, 1980, neither can speak Japanese, and so on. Similarity *tout court* is wildly profligate and unhelpful, to say the least. As such, similarity claims should be made in the context of

6 A way of understanding loss that continues the Lewisian theme would be to appeal to “alien properties” (Lewis 1986, 91–92). Whereas alien properties, for Lewis, are properties that do not exist in some world but do in another, we could understand loss as an alien property in a past world relative to ours.

7 Thanks to Tyler Brunet for pointing this out.

some set of properties or measurements. To say that mammoths are not lost because elephants exist is to say that elephants are similar to mammoths in specific ways. In particular, it is to say both that elephants are similar with respect to many pachyderm properties (they both have trunks, for instance) and, perhaps, that they are both Elephantidae. Thus, insofar as loss judgements are judgements of similarity, loss judgements are sensitive to context. That is, a denizen is only lost under some characterization.

Presumably, all past worlds are to some extent lost and to some extent erased. As I've said, denizens are profligate: They could include such insignificant aspects as, say, my spatial position in a room. That property is then lost when I move from that position. Going similarly fine-grained on erasure has similar results: presumably, some information from a few seconds ago is not retrievable.

Following from this, erasure and loss are both plausibly understood as *graded*, potentially along multiple dimensions. A past world will be erased insofar as information about that world is not contained in the present; a past world will be lost insofar as there are denizens of that past world that are not present now. On these definitions, holding fixed a set of denizens at least, one world might be more or less lost than another. For instance, mammoths are less erased than dinosaurs, as we have more, and more detailed, mammoth fossils. However, as we discover more fossils or discover more ways of interpreting them, dinosaurs might become less erased. So, erasure can be understood in graded terms. Loss is somewhat trickier on this front, as matters of existence are not typically considered to be vague: you either exist or you do not. However, mammoths, understood as members of Elephantidae, are not lost; but understood as the species *Mammuthus primigenius*, they are lost. This is sensitivity to description, not vagueness, but the speed at which scientists often switch between characterizations generates an effect very close to true vagueness. Further, under a more epistemic conception of loss, we might be uncertain whether some denizen is lost or not; so, we might have varying credence in loss. Further, considering loss as it applies across past worlds, differences in loss in those worlds' denizens will lead to gradation. There are senses, then, in which loss is graded as well.

On my view, in many circumstances, absences can become lost. That is, in principle, an absence could be a denizen, or at least a relevant property of a past world. Part of what makes the Avalonian Ediacaran mysterious isn't just what populates it, but what doesn't populate it. The lack of non-sessile complex metazoans marks an enormous difference between Avalonian and

later times. So, one of the properties of the lost Avalonian past world is a *lack* of non-sessile complex metazoans. Including absences as properties, I think, is crucial for capturing elements of loss as it plays out in the epistemology of historical science. It also further increases the plenitude of lost worlds to infinity. However, I don't think this undermines the distinction: as with issues concerning similarity, loss being indexed to a particular set of denizens as partly determined by scientists' interests sufficiently constrains the notion in particular applications.

The upshot of this is that we shouldn't really call a world "lost" *tout court*. Rather, when I say a world is lost, I typically mean that a particular set of denizens is lost, relative to now, in some significant way. I don't think I mean anything fancy by "significant" here. Our own interests will do: The loss of the Avalonian Ediacaran is significant in part because it makes understanding that early example of metazoan life so difficult—and so enticing.

Because denizens can be characterized in so many different ways, and loss turns on the existence of denizens in past worlds and now, it will be helpful to say something about how we *kind* denizens. That is, how do we sort them into sets or groups? I can imagine a more restrictive view than mine that requires loss to appeal to natural kinds, say, or classifications as scientists use them. There is, indeed, a rich literature in the philosophy of biology and philosophy of science on classification (e.g., Richards 2016), and classification in the Earth sciences is often complex, with particular schemes playing multiple roles simultaneously (Witteveen 2024). Continuing with my (hopefully not too frustrating) non-committal approach to metaphysics thus far, I won't favour any particular story here about kinds or classification.⁸ Nonetheless, there are three broad ways of characterizing denizens and linking them together, which are helpful for my purposes (I'm adapting these from Khalidi 2021; see also Currie 2019a). (In chapters 6 and 7, I'll draw on Catherine Elgin's notion of "exemplification" to explain artifactualist reasoning. She uses the language of "features" rather than "types," and I think that could work well here also, but the distinction between historical tokens,

8 My metaphysical leanings do tend toward fairly expansive, "promiscuous" ontologies (Dupre 1993), and it might be that some of that is being smuggled in here. I suspect that more restrictive ontologies could nonetheless accommodate much the same account of denizens and their properties as I have, it would simply result in some of these being considered mind-dependent or derivative of more fundamental ontological structures. I take it that all of these "denizens" would still be epistemically relevant, however.

types, and ahistorical types does the work I need it to.) So, let's list and discuss Khalidi's three notions.

- (1) Historical token (i.e., historical individuals: *that* mountain, *that* lineage, *that* organism)

Philosophers of biology often take many biological categories, particularly those linked with taxa, such as species, orders, etc., as *historical individuals* (Hull 1976). Historical individuals have identity conditions similar to the individuality of ordinary objects: they will have particular *persistence conditions*, often something like spatio-temporal continuity, and potentially particular *generation conditions* (see Ereshefsky 2014 for a particularly cogent discussion that emphasizes the “historicity” of individuals). For instance, the Dodo lineage began when it speciated from its pigeon-like ancestors and ended when the species went extinct sometime in the second half of the 17th century. My time as a postdoctoral researcher at the University of Calgary began with my signing a contract in 2014 and ended with my shifting to the UK in 2017. When we consider loss in the context of historical token individuals, we ask whether that denizen has survived into the present. Squanch, the individual mammoth, is lost due to her death. The mammoth species is lost due to extinction. The historical individual Elephantidae (which mammoth belongs to) survived into the present and thus is not lost. So far as we know, no rangemorph taxa survived into the present—they left no ancestors.

Epistemically speaking, historical individuals often act as denizens as well as provide information about denizens; they can thus block both loss and erasure. Elephants, *qua* Elephantidae, and regarding their “pachydermy” properties, ensure that mammoths are not lost in those regards. In virtue of their ancestral relations, elephants also possess information about their common ancestor with the mammoth, and thus are partly traces of them. I don't see any particular problem with entities playing both roles.

A crucial split between ways of characterizing denizens is taking them as token individuals on the one hand and as types, sets, or categories on the other hand. I'll divide that latter division into two.

- (2) Historical type (membership requires the *same process type*, e.g., igneous rock)

Some classes are such in virtue of having undergone the same kind of generation process or, potentially, maintenance process (Millikan 1999). Igneous

rock is a typical example: any rock formed of cooling volcanic lava is igneous, and will share various properties in virtue of this. That is all it takes to be igneous. A biological trait is an *adaptation* if it has been selected for in the past. So, we can group adaptations in terms of what they are an adaptation to. Wings are adaptations for flight for birds, bats, and pterosaurs. Consider particular adaptive radiations, such as that which spread our species out of Africa and across the globe, or that which took the Cambrian's metazoans into evolutionary prominence. We can consider both of these as individuals: the Cambrian radiation completed many millions of years ago, so it is now lost in that sense. But we can also understand them as historical types: an adaptive radiation is an event whereby some set of taxa, in virtue of a novel adaptation or set of adaptations, occupied new territory. Insofar as adaptive radiations still occur and are still occurring, that historical type is not lost. Forms of speciation are also historical types: cladogenic speciation, for instance, occurs when two populations split, that is, a branching event as one lineage becomes two. Insofar as this still sometimes happens, cladogenesis is not lost.

As a final example, consider Liu et al.'s explanation of iveshediomorphs. Where previously the similarities were understood in terms of historical individuality—the fossils represented ancestrally linked taxa—Liu et al. argue that they are in fact similar due to undergoing a similar process, namely, decaying under Avalonian conditions. Potentially, then, “ivesheadiomorph” is a historical type: to be one is to have decayed and been fossilized under those circumstances. And further, it seems to be a lost historical type.

(3) Ahistorical type (membership turns on shared properties/
dynamics: e.g., possessing traits, ecosystem/ecotype)

Where historical types demand certain generation procedures for membership, there are various ahistorical types that rely simply on shared properties or characteristics. Typical examples of these include essentialist conceptions of objects: “gold” is anything with the relevant chemical composition, for instance. But consider as well the ecological dynamics that interest Mitchell et al. To have those dynamics would be to approximately instantiate a neutral-driven profile in an ecosystem. If their work is right, then there are no (or very few) systems like that around now. But further, recall their worry that Avalonian ecology “plays by different rules.” Here, the worry might be not so much that there are no instances of that kind of dynamic around anymore,

but that the models we use to understand contemporary ecological dynamics don't gain purchase in the Avalonian: ecology itself works differently.

So, historical scientists indicate putative past worlds and try to understand them, both via engagement with traces from those past worlds and by examining those historical individuals, historical kinds, and ahistorical kinds that have survived. The accessibility of a past world turns on the loss and erasure of that world's denizens relative to now. Denizens can be characterized as token historical individuals, and historical or ahistorical types at various levels of grain. Individual mammoths are lost, as is that taxon understood as a species, and they are lost when characterized as members of Elephantidae adapted to Arctic-like conditions. But understood as Elephantidae, or as cold-adapted mammals, or as far-roaming herd animals, they are not lost.

Before applying the framework in the next two sections—and throughout the remainder of the book—there's a central objection to tackle: Can we really distinguish between loss and erasure? If we consider loss and erasure's relationship, we find them intimately connected: Every case of erasure is a case of loss, and likely any real-world case of loss will include some erasure. For instance, a mammoth fossil being destroyed via subduction is a case of erasure, but it is also a case of loss, specifically, loss of that fossil. A mammoth's dying is a case of loss (of that individual), but it will also involve the erasure of various traces of living processes. The two are asymmetric in this regard. As traces can count as denizens, erasure can always be understood in terms of loss. However, in principle, a perfect recording machine of a mammoth would ensure against erasure despite loss. This is merely in principle, however: in the actual and near-actual circumstances that concern us throughout most of this book, loss will bring erasure. Do these features suggest that the distinction I've drawn collapses? If we expect these categories to be applied independently of scientists' interests and practices, then perhaps. However, appealing to the context-sensitivity of erasure and loss saves both the distinction and its function.

Once we fix a set of denizens, we can ask whether examples of those denizens have survived into the present, to what extent loss has occurred, and whether downstream information about those denizens in the past has been retained, that is, to what extent erasure has occurred. Elephants are, in a sense, both the same kind as mammoths (as members of the same phylogenetic family) and contain traces of mammoths insofar as they contain traces of

their joint ancestry.⁹ In Liu et al.'s explanation of ivesheadiomorphs as false taxa, that is, a phenomenon generated by the unique decay conditions of the Avalonian, one of the denizens included in the past world is the decay conditions themselves. The traces of the decay conditions have not been erased, but the decay conditions themselves have been lost. What makes this the case is partly what is included in the past world and what is not. To see this, let's examine Liu et al.'s argument in a little more detail.

Liu et al. rely on a careful interpretation of patterns of ivesheadiomorph morphology, linked to a particular picture of decay and preservation in the Avalonian. In short (see Liu et al. 2011 for details), they argue that ivesheadiomorph morphology reveals a pattern of more or less linear, effaced decay: from a rangeomorph-like beginning, you can arrange them into a pattern of increasingly decayed and distorted signals. From this pattern, a mechanism involving slow microbial decay coupled with preservation potentially occurring anywhere throughout the process is suggested. And this mechanism, they argue, would generate pizza-disk forms from once-living frondoid forms. Thus, it is concluded that pizza-disks are dead rangeomorphs and their allies (see Figure 2.5).

Liu et al. are not, in this instance, trying to understand a world occupied by ivesheadiomorph fossils, but one populated with the conditions leading to those fossils. That is, the past world includes rangeomorphs, a deoxygenated ocean, the microbial mat, and so forth. When scientists try to understand erasure processes themselves, they often do investigate worlds where things like fossils are taken as denizens—if we're trying to understand how subduction processes might have distorted or destroyed fossils, for instance—but we can identify which is which by attending to the interests of the scientists at hand: which denizens they indicate. Liu et al. aim to reconstruct rangeomorph decay processes. As such, the past world they indicate includes denizens relevant to understanding the decay of pizza disks.

Although interest-sensitivity saves the distinction from collapse, loss and erasure are still connected. For example, to the extent that loss has not occurred, then some erasure has not occurred. Currently living elephants contain information about past living elephants: they are (potentially, see footnote 9) traces of them. An upshot of all this is that if you really want to

9 There is an open question here regarding whether we take common-cause cases, such as the ancestral continuity between mammoth and elephants through their common ancestor, to count as traces. I don't think much turns on this for present circumstances.



FIGURE 2.5 Liu et al.'s (2011) Explanation of Ivesheadiomorph Decay, from Living Forms (Upper Left) to Highly Effaced and Pizza-Disk-Like (Bottom Right) (Fig. 13, Liu et al. 2011). Reprinted with Permission from *Paleontology*

think about all of this in terms of our accessibility to information, you could likely pull it off, but you'd need to recover a distinction along the lines of loss/erasure to capture the aspects of historical knowledge I'm interested in.

To summarize, if we take a past world in a sufficiently maximal way, loss and erasure nearly collapse; however, if we hold a set of denizens fixed, then we can identify loss and erasure in that context. Thus, the distinction is meaningful across the (many, many) contexts that are relevant to understanding historical knowledge.

So, we have the basic conceptual machinery on the table. Palaeontologists aim to reconstruct a past world and, in so doing, indicate a set of denizens. Their investigations are thus challenged by both erasure, that is, the loss of information causally downstream from those denizens, and by loss, that is, those denizens' cessation.

3. Tracecentrism and History Redux

In the last chapter, I claimed that the philosophy of the historical sciences has been too narrow, focused too much on traces. I also discussed how the early sciences of palaeontology and geology justified their importance by claiming that the past differed significantly from the present. The distinction between erasure and loss allows us to capture these ideas more carefully.

Tracecentric accounts of the epistemology of the historical sciences claim (or at least imply) that the central challenges of uncovering the past turn on the erasure of traces, and characterize the method of historical science as relying heavily on inferences from traces. There's no doubt that the past has had a causal effect on the present, and thus these effects provide an inroad to the past, which is a foundational insight that shouldn't be disregarded. However, to uncover the past we need, yes, the traces, but also an understanding of both how those traces form and the kinds of denizens that could leave those traces. Ichthy can tell us about Broadfoot not simply because they have the specimen, but because they understand fossilization processes, which they gained in part from study of similarly acting processes in the present, and because they live in a world with mammals, albeit mammals quite different from Broadfoot. If no such mammals existed, if they too were lost, then Ichthy's epistemic situation would be significantly more difficult.

Such loss-induced difficulties are pressing for unusual extinct taxa, such as some dinosaurs. Although we have a fair number of traces of, say, *Diplodocus* and other sauropods, without existing analogues, our capacity to make sense of them is diminished. Considered as historical individuals, sauropods are not completely lost because the Dinosauria clade they belong to survives in modern Aves. But birds depart so far from these massive, quadrupedal animals that they constrain our understanding insufficiently. Understood as an ahistorical type, say, as a large quadrupedal animal, we have some mammals to consider (elephants, for example), but it is often unclear how similar we should take elephants and sauropods to be. Or if sauropods lived in herds, we have plenty of animals in herds (an ahistorical ecological and behavioural kind), but it is difficult to tell how similar herd behaviour in sauropods is to herd behaviour in contemporary animals. The significant loss between sauropods and anything that exists now makes it hard to know which aspects of today's possible sauropod analogues can be projected back to sauropods.

In short, it seems as if successful investigation of past denizens requires (1) their traces and (2) those denizens not being completely lost. That is,

something must still exist that can be linked to those denizens as historical individuals, historical kinds, or ahistorical kinds. Existing examples enable us to understand how traces form, and the kinds of entities that might form them. I'll further articulate this idea in the next section, but the point here is just that to the extent that an account of the method and challenges of historical science is focused on traces and decay only, it misses the critical importance of loss.

As Rudwick and other historians of science have shown, the discovery and articulation of a pre-human history revolutionized our understanding of the world and further provided a motivation for studying the deep past. Although other cosmologies had histories, they lacked an open, discoverable and surprising narrative about that history.¹⁰ In a sense, these geologists discovered loss: the denizens of the past differ from the denizens of the present. In the last chapter, I contrasted Professor Ichthyosaur, a caricature of Lyell, with Prof. Ichthy. Professor Ichthyosaur sees the geological world according to a steady-state model. Although particular historical individuals might come and go, the general dynamics of geology, the environment, and biological responses to that environment are not lost but continue acting today. Although Lyell might not want to say that the marine reptiles that would reappear in Jurassic conditions would literally be ichthyosaurs, he nonetheless seems to be committed to them being ichthyosaurs in many of the ways that matter. By contrast, his detractors like De la Beche (and presumably Prof. Ichthy, if only they were to know of Lyell's views) see the world as being in significantly more flux, much more contingent and more path dependent. It thus potentially contains more loss. They deny that the processes and dynamics that apply today apply in the humanless worlds of the past.

Loss's sensitivity to description allows us to capture different positions regarding the relationship between the past and present. A steady-state perspective sees much loss, no doubt, but that loss is not epistemically significant: by examining today's entities and processes, we can recover all we need to understand the past. Individual losses can be reconstructed via an understanding of currently ongoing processes, and, indeed, the losses are likely insignificant due to the circularity of the Earth's processes. This thought underwrites Lyell's uniformitarianism. As our metaphysical picture shifts from a steady-state model to an increasingly contingent one, we see more significant

10 As Chis Manias points out to me, Christian cosmologies certainly include notions of loss!

loss, making the past more mysterious—uniformitarian reasoning becomes ineffective.

Recall Mitchell's question about the ecological dynamics of the Avalonian Ediacaran. It isn't simply that Avalonian ecosystems are lost; she also worries that the very ecological dynamics that they instantiated might be lost too: they raise the possibility that the Avalonian played by different ecological rules. If we want to understand the Avalonian, potentially we don't simply need to consider quite different denizens, but these denizens might break the general models we use to understand how ecosystems form and function.

Thus, the conception of loss I've characterized provides a challenge to philosophers interested in understanding historical science, and it has the resources to capture various positions on the nature of history. On one extreme, loss is never significant—the world moves in stable patterns, which we can uncover by watching those patterns unfurl. On the other extreme, significant loss is rampant, and history is wildly contingent and path-dependent. The challenge, I take it, is to understand how to navigate between these two extremes. I'll take up this challenge in chapter 8.

4. *The Challenge of Loss*

"Broadfoot is unique. And this uniqueness is baffling. Us scientists," Prof. Ichthy gestures magnanimously, including various graduate students, "deal in numbers: statistics, data, measurement. We take nature's chaos and find order with a ruler. And, of course, we can measure our specimen." Prof. Ichthy places their lecturing rod alongside the encased skull.

"My pointer here is about thirty centimetres long, and the skull is about one third of the length of the pointer. So, I can surmise a skull of approximately one third of that length, that is, in the vicinity of ten centimetres. But what are we to make of this measurement without anything to compare it to? Is ten centimetres normal for members of Broadfoot's species? Is this specimen an outlier or an infant? Should we treat it like a typical mammal in terms of head length, when we know it's not a typical mammal in so many other ways? And if we must ask these questions of something as simple as skull size, what are we to make of its apparent aquatic adaptations—the bill and webbed feet—to say nothing of its reptile-like features, or its plethora of other mysterious, unique properties?"

Ichthy pauses for effect, leaning their lecturing rod against the lectern.

I've stated that Broadfoot was venomous, that it hunted using electrosense, that it was the only egg-laying mammal we know of. But how could we possibly find that out? We interpret fossils against the backdrop of the wide variety of things we know, and that knowledge requires numbers. Again: statistics, data, and measurement. If some fossil feature is unique, it seems, we haven't any way of gaining the grounding required to make scientific sense of it.

Ichthy takes a step from the podium, lays a flipper on the encased skull, and says in a stage whisper that carries across the audience: "Might all my research merely amount to non-scientific speculation?"

Ichthy provides a conundrum: if scientific knowledge depends on examining multiple examples, how can we have a science of the unique? As we'll see, although loss doesn't necessarily generate uniqueness *per se*, it can generate similar problems.¹¹ The next three chapters could be read as responses to Ichthy's conundrum. To identify the challenge that arises from loss, I'm going to take a tangent into a related discussion: *N=1* problems as they play out in macroevolution and astrobiology.

An *N=1* problem, traditionally understood, holds when there is some entity or process we want to understand, but our sample set is one, that is, we have only one specimen. The term has currency in two intimately connected debates:¹² discussion of life's contingency at a macroevolutionary scale, and investigation of life beyond our planet's boundaries (Powell and Mariscal 2015; Powell 2020; McMahon 2021). I'll argue that a single specimen does not an *N=1* problem make. Instead, *N=1* problems are problems where we have insufficient specimens or data to establish the background theory required for legitimate interpretation. With *N=1* covered, I'll argue that the challenge of loss is similar in character.

Questions about the robustness or fragility of evolutionary outcomes at a macro-scale have much energized palaeontologists, theoretical biologists, and philosophers of biology. For instance, as we'll explore downstream, the Avalonian world ended at some point, and the Ediacaran fauna were replaced by the funky, complex, and strange metazoans of the Cambrian. Was that

11 For discussion of epistemic problems related to uniqueness, see Tucker 1998; Buskell and Currie 2021.

12 As Rose Trappes points out, there are also *N=1* cases in medicine. I'm not sure if these play out in the same way.

outcome inevitable or contingent? Did it have to happen, and, if so, what necessitated it? And, further, what were the factors that shaped that modal profile? Were rangeomorphs and their ilk out-competed by novel species? Were they fragile in the face of environmental collapse? Or did they simply get unlucky? Overall, we want to know what the choice-points in evolution's history were: where its path could have gone one way or another. Discussion of contingency, convergence, and what upshots it might have for the nature of the past has garnered a fair bit of attention from philosophers (for instance: Beatty 2006; McConwell and Currie 2017; Ereshefsky and Turner 2020), and we'll return to it in chapter 8.

Astrobiology aims to understand the diversity of life in the universe: if there is other life out there, how similar will it be to life on Earth? How might we detect extraterrestrial life? Can we take current life as representative of all other life or not? This matters for how we might go about detecting extraterrestrial life: how will we know it when we see it (Mariscal 2015; Cleland 2019)? Some connect macroevolution and astrobiology rather intimately: Rachel Powell (2020), for instance, argues that various convergences on Earth, and regularities drawn from them, can underwrite fairly robust predictions about what alien life will be like.

Both discussions of macroevolution and astrobiology run into a similar problem: on the face of it, we have only one specimen of life. Our N , then, is 1. Life on Earth has a common origin; so, its modal profile could have been shaped by quirks inherited from that common origin. In answering questions about contingency, we ideally want independent events that generate data to inform our hypotheses, but common origins mean the data is not independent. As Kim Sterelny has put it, even in instances where distantly-related lineages have evolved down similar pathways, "The envelope of evolutionary possibility for each lineage will depend in part on homologous developmental resources" (Sterelny 2005, 588). Similarly for astrobiology: without relevant, independent contrasts, it doesn't seem obvious that we can determine whether life on Earth is as it is because that is what life is like everywhere, or if there are particular quirks about life on Earth and Earth's particular conditions that determine its nature. Some argue at least, that common origins, and commonality in many of Earth's conditions across time, mean that $N=1$.

Although $N=1$ problems are often cashed out as due to a lack of specimens, that can't be the whole story, and this matters for understanding loss. There are many cases in palaeobiology and other historical sciences where N seems to equal 1, but these don't appear to have the same epistemic consequences

as reported in macroevolution or astrobiology. Professor Ichthy's Broadfoot specimen is the only one available: it is the only direct trace of the animal's existence. But, as the Professor points out, we know enough about fossilization and mammals to be able to make solid claims about the past based on that single specimen. In real-life vertebrate palaeontology this is common: new species are identified through single teeth and other partial, singular specimens, fairly regularly (see Currie 2018a, chapter 1). This is made possible by the rich epistemic background palaeontologists draw on. Because we know so much about mammal teeth, a single tooth— $N=1$ —is sometimes sufficient to underwrite well-evidenced claims about the past.

So, $N=1$ problems are problems not *simply* because we have insufficient specimens, but because we lack the epistemic background required to understand the relevant causal and modal structures that the "1" is embedded in. As Ben Jeffares has pointed out regarding our knowledge of fossils: "Our understanding of the palaeontological record is not built *de novo* from the observations of fossils, but constructed on the foundations of our knowledge of the contemporary natural world and the deployment of a general regularity" (Jeffares 2008, 474).

So, although for some fossil taxa we have only one specimen, because fossilization is so well understood, having a single specimen doesn't necessarily block knowledge.

This generalizes: traces are only traces, that is, evidence of past worlds, because we have some grip on regularities concerning how those traces form. We are only able to build on that knowledge because we have some grip on the kinds of dynamics and relationships that hold between these kinds of systems. To get that knowledge—at least it seems—we need to observe and study examples of those denizens. This is what makes $N=1$ problems tricky. With only a single example, we seem to lack the resources to develop the background knowledge required to make sense of counterfactuals concerning life's path or surprising extraterrestrial signals.

The apparent inability to build robust epistemic backgrounds is what makes loss an epistemic challenge. The denizens of the Avalonian Ediacaran, the strange fossilization processes due to the microbial mat and lack of metazoan mobility, the surprising neutrally-driven ecological relationships between rangeomorphs, as well as their uncertain phylogenetic relationships, uncertain metabolic and feeding strategies, uncertain developmental pathways, and so on, are lost. There is (for all we know) nothing like them around today, meaning that (putatively) there is nothing we can look to in building

our background theories. As such, on the face of it, we lack *any* specimens with which to constrain our theorizing about them. This isn't a $N=1$ problem. In some sense, we're faced with an $N=0$ problem.

Of course, it is not literally $N=0$. We *do* have traces—the world of the Avalonian might be lost, but it isn't erased—so we have some partial access. Those interested in life's contingency and in astrobiology have extremely rich access to a sample of life—it is all around us—while those studying lost worlds often have extremely limited access.

To see the idea, let's briefly compare the astrobiologist and the palaeobiologist of the Ediacaran (we'll return to astrobiology in chapter 6). The former is often in the business of trying to detect biosignatures from exoplanets. The chemical composition of an exoplanet's atmosphere is retrievable from its wavelength, and some atmospheric compositions could signal the presence of life (heaps of oxygen, say). So, the light from exoplanets counts as traces, and the theory of biosignatures works analogously to a theory of fossilization. From where was that theory developed? Well, observations of our own Earth (also studies of early Earth and comparisons with early Mars, but let's not worry about that): to simplify, we know the kinds of biosignatures to look for—we have the basics of how to make sense of the relationship between atmosphere and biosphere—because we are intimately familiar with a particular atmosphere and biosphere. An $N=1$ problem is generated because we don't know how representative our planet is: it could be that Earth-like atmospheres can be generated by abiotic processes, or it could be that there are biota that have profoundly different biochemical signatures than ours. But still, that 1, the Earth, is a crucial lynchpin. But also note that *N isn't* in fact 1, at least regarding planetary atmosphere: there are many traces from those.

Now compare this to studying something lost, like the Ediacaran pizza disks.

The pizza disks are the result of conditions of decay and preservation unique to the Avalonian Ediacaran. There is no analogue of Earth's life for astrobiology here. Studying something lost from the deep past is a bit like trying to understand exoplanet atmospheres without being able to look at the Earth's atmosphere. There are—on the face of it—only the traces to go on. This is the challenge: how can we build the kind of robust background theory required to interpret traces when their sources are lost? Without sufficient denizens surviving into the present, such theories seem beyond our reach. Thus, we cannot make sense of lost denizens. With such limited information

about those lost denizens themselves, analysis of their traces can do little to improve the situation.

As with Prof. Ichthy's rather staged admission of doubt, I'm surely laying it on a bit thick here. Clearly, loss doesn't seem to block scientific knowledge. We've already seen a few hints as to why. A lack of erasure can mitigate loss, and a lack of loss can mitigate erasure. Prof. Ichthy's impoverished, but still present, mammalian fauna underwrites their ability to make sense of the strange fossil specimen they discovered and hence infer many of Broadfoot's surprising properties. With a sufficiently rich set of traces, we can puzzle our way to understanding loss even without contemporary examples, as we saw with Mitchell's uncovering of the surprising ecological dynamics of Avalonian systems and Liu et al.'s arguments for ivesheadiomorphs' ichno-taxic status. But as loss becomes increasingly rampant, the challenge becomes pressing. What can we do then?

That question drives the next three chapters, whose aim is to abstractly capture three (non-mutually exclusive) strategies historical scientists adopt when faced with loss. One set of strategies involves emphasizing what in the target world has *not* been lost: they *graft*, that is, seek aspects that are historically continuous with the past world and seek to understand those. Another set of strategies is, in a sense, less direct. Scientists *artifice*, that is, construct a system with similar properties as the lost world. They resurrect it. And finally, they might *anchor* by navigating between a set of differing systems to construct a model of their target. These strategies are often used simultaneously, but characterizing them separately, I hope, helps us understand their underlying logic. As we continue, a characterization of an overall strategy arises: as we'll see, loss is tackled through a combination of *strategic perspectivalism* and *iterativity*.

