

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>

Grafting

Prof. Ichthy returns to the slide with the speculative history of mammal life.

“Of course, we shouldn’t exaggerate the differences between Broadfoot and extant mammals. They’re still *mammals*, after all, and thus joined together via ancestry. In recognizing Broadfoot as part of a set of historical trajectories—narratives if you want—we open powerful inroads to understanding Broadfoot, those narratives, and today’s mammals.” Ichthy uses their pointer to trace the black lines from mammals past to mammals present.

“These stories are carried across genetic and developmental pathways, from parents to offspring, across millions of years. Thus, currently living mammals that share in those stories all help us understand Broadfoot, just as Broadfoot helps us understand other mammals.”

Ichthy turns from the projector. “But why restrict ourselves to mammals? We and Broadfoot are cousins, after all: vertebrates. As such, we also share a lineage, another story, tracing back into the deep past. And further, Broadfoot shares a deeper ancestry: that linking all life that we know. As we tell Broadfoot’s story, then, we also tell the story of mammals, as well as our own story, as well as the story of life. Life is a complex, messy thing, and, as such, there are many different lines of ancestry, many different stories we might try to tell. This complexity is a boon for understanding Broadfoot.”

Again, Ichthy begins counting with their flippers. “There are two lessons here, I think. The first lesson is that even though Broadfoot is a weird mammal, its very mammalhood means that we have vital connections to it through its cousins. The second is that our knowledge of the past isn’t always restricted to little fragments or slivers of time: by understanding the connections

between things in the past, we can weave stories which, or so I hope, are at least sometimes a little bit true.”

In the last chapter, I introduced *strategic perspectivalism*. In response to the challenge of loss, the scientist adopts a perspective—characterizes their target—such that certain properties of that target may be isolated and brought into dialogue with something that can be examined.

We’ve seen this strategy play out in *artificing*, where those properties were instantiated and probed in a bespoke model or experiment. In this chapter, we’ll see the strategy take a quite different form, which I’ll call *grafting*. Under this strategy, scientists focus on historical individuals that have been retained into the present, or at least that we have better knowledge of. When you think of a mammoth as a member of the still-existing Elephantidae, or when you examine birds to learn about dinosaurs, you’re grafting.

A “graft” involves linking previously disconnected tissues so they can grow together. In medical contexts, a skin graft attaches healthy skin to damaged (for instance, burned) skin to encourage repair. In horticulture, grafting involves linking the properties of two plants. For instance, combining the fruit qualities of one tree with the more resilient roots of another. Real grafting, then, involves actively linking previously unlinked properties—combining what was once disconnected into one individual. My use of the term “grafting” is more metaphorical than the last chapter’s use of “artificing.” When I use it, I identify a set of practices whereby historical scientists, faced with a putatively lost world, posit existing properties that might be historically continuous with some of the lost world’s denizens, thus “grafting” them together. A downside of the metaphor may be that it implies a more constructivist take on the practice than we might like. But given this chapter’s focus on phylogeny, the horticultural reference is too good to pass up.¹

We can understand grafting as a general strategy that involves understanding a target as part of a causally connected token trajectory. Often (but

1 Somewhat to my chagrin, Hans-Jörg Rheinberger has recently used the term “grafting” in his discussion of experimental practice in the life sciences (2023). His usage is different but, I think, usefully so. For Rheinberger, “grafting” involves incorporating some existing experimental apparatus or procedure into some other apparatus or procedure, thus extending its capacity into new territory, but in a somewhat modular way wherein the grafted objects maintain their identity. This is much like how, say, a variety of apple grafted onto another variety of apple tree does not, as it were, merge with the tree, but keeps its independent identity. This is “a delicate operation, often requiring a considerable degree of adjustment and refinement” (Rheinberger 2023, 66).

not always!), grafting is rooted in conceiving of your target as a historical individual, that is, some token entity or process extended through space and time. But it is crucial to see that past worlds are often—perhaps always—characterizable as many different kinds of historical individuals. Rangeomorphs aren't only construed in ecological, geometric, or morphological terms. They're also characterized as species, genera, orders: as taxa. That is, they are *lineages*.

As a first pass, we can define a lineage as a kind of biological individual, one that's constituted by further individuals connected through relationships of ancestry and descent. For instance, culture generates a plethora of lines of descent. Philosophers and other academics often delight in tracing their lineages: my PhD supervisor (Kim Sterelny) was supervised by Bill Bonney, making him my academic grandfather, I suppose; but I was also a postdoc with Marc Ereshefsky, so I might trace another line of ancestry from him. And this makes Celso Neto, whom Marc supervised as a graduate student, something like a step-sibling, I guess. So, cultural lines of descent might be drawn in multiple ways, although it is certainly not that anything goes. Neto (2019) has argued that biological lines of descent are similarly complex and plural. First, biological categories are often hierarchical, and at different levels we can identify different descendants and ancestors. Parent and offspring genes, homologous traits, individual organisms, microbiomes, as well as groups (consider termite mounds, prides, and so on) might, under the right conditions, be considered lineages, that is, individuals forming lines of descent. Second, lineages sometimes form within us. For instance, through development, our cells form lineages. Third, inheritance is not always vertical, but sometimes horizontal (most obviously in horizontal gene transfer). Horizontal “ancestry” means that biological history is not always best represented by tree-like networks, but sometimes (often?) requires reticulate networks.

For Neto, the upshot is that “There is, then, no single answer to the question, ‘what is a lineage?’” (Neto 2019, 1109), and this pluralism makes room for a plethora of options regarding how we might graft.² When faced with a potentially lost world, we ask, well, what historical individuals or lineages might *not* be lost? Although there are no platypus descendants for Prof. Ichthy to make use of, there are plenty of cousins.

2 Alison McConwell (2019) has made a similar argument regarding the contingency of evolution (see also chapter 8) and the plurality of the biological individuals that might be under selection, and I think similar lessons can be drawn here.

Look out for two crucial lessons in this chapter. First—echoing Prof. Ichthy—if lost worlds can be intertwined into broader historical narratives, they can be made sense of in those terms; second, although perspectives on a particular target might differ, they also often overlap in various ways. This enables them to be compared, contrasted, and used to test one another. This second point underwrites “iterativity,” which I’ll introduce in the next chapter.

As with the last chapter, we’ll start with an example from dinosaur biology, before returning to the Ediacaran. I’ll then consider disagreements and overlapping perspectives, and finish up with a discussion of living fossils. The examples in this chapter will all draw on a particular subset of grafting strategies, namely, those involving phylogenetic reconstruction. It is important to note that not all grafting involves phylogeny: in the next chapter, we’ll see grafting applied to reconstructing the relationship between complex metazoan evolution and the emergence of various global oceanic processes. It’s also worth noting that trace-based reasoning is itself a form of grafting, as it involves linking a token trace or set of traces with their causal ancestors. I’ll focus on phylogeny here because it is a particularly salient example in palaeobiology.

1. Do Dinosaurs Feel Pain?

“So, what was Broadfoot like?” Ichthy muses, “Some things we can be very confident of: we know they were vertebrates, like ourselves. So, they would have hearts and lungs, the usual battery of organs common to our lineage.” Ichthy indicates relevant anatomical parts on themselves. “They would breathe air, metabolize similarly to us, seek out prey, and escape predators. There’s a wide range of traits we can more or less safely assume Broadfoot would have had.”

Ichthy again gestures to the mammalian history behind them. “Taking them as mammals, as I said, we can go further, but there are challenges. For instance, one might assume Broadfoot gave birth to live young, just like all living mammals do. However . . .”

Ichthy waves at their grad student, who quickly switches slides. The next slide shows some fossilized shards, immediately recognized by the initiated as the remains of eggshells. “You’d be wrong in thinking that.”

Did dinosaurs feel pain? At first blush, this is a silly question. On the one hand, we might say: of course they do, insofar as any complex animal does (Browning and Birch 2022; Gibbons et al. 2022). On the other hand, we might consider this unknowable: feeling pain is a subjective, internal state that we have deeply limited access to. Put in these terms, it is hard enough to confirm pain in another human or our pets (due to not having access to those states), let alone long-extinct lineages (Harnad 2016).

These two interpretations of the question about dinosaur pain imply perspectives. The former implies something like the kind of phylogenetic inference we'll see below, the latter a set of philosophical views about the nature of mind and its accessibility. However, neither of these perspectives is very *strategic* from a scientific standpoint: in characterizing a denizen, the smart money is to adopt perspectives that don't make things too easy (dinosaurs feel pain because any complex animal does) nor too hard (to answer this, we need to solve the problem of other minds). The trick is to find a sweet spot where our interpretation allows us to use a set of tools that underwrite a perspective productively. In other words, neither of the above approaches are particularly *productive* for tackling loss or putative loss. I'll think harder about productivity and what makes a perspective a good one in future chapters. For now, let's examine another way of approaching the question of dinosaur pain.

Les Hearn and Amanda Williams (2019) have tackled pain in the Mesozoic. This is a useful case for us here, first, as a way of introducing grafting and, second, as a way of emphasizing how the evidential meaning of traces often requires grafting. Hearn and Williams' strategy is to consider how neontologists (biologists of the "now"; see Currie 2019b) think about pain in animals and consider how they might detect similar properties using the traces available to them. As we'll see, to some extent, the interpretative work they bring to those traces requires grafting, that is, thinking about dinosaur ancestry. To make this salient, I'll bury the lede by introducing their work without considering ancestry, and then I'll bring it in.

Biologists call the detection of painful stimuli *nociception*, and in non-human animals, it is inferred via two routes. First, the possession of particular internal structures, such as a central nervous system and the relevant receptors. We might ask: Does this critter have the organic machinery associated with pain detection? Does it have the infrastructure required for pain? Second, does it exhibit pain-avoidant behaviours, such as keeping away from pain-causing stimuli and sparing damaged parts? We might also ask: Does this critter behave in ways associated with being in pain? To apply this

approach to dinosaurs, then, we need to know how to infer internal structures and how to infer behaviours.

Hearn and Williams focus primarily on pain-avoidant behaviours. They conduct a literature review, hunting for various dinosaur pathologies and other signals of potentially relevant behaviour. That is, they ask what damaged bones and other signs of bodily harm can be found in the fossil record. So, what would fractured dinosaur bones and the like have to do with inferring nociceptive capacities?

Hearn and Williams take their cue from the qualities associated with pain-avoidant behaviours: for instance, they don't simply look for fractures, but *healed* fractures. An injury healing requires the animal to avoid infection and further damage by altering normal behaviours to spare the damaged part. They also look for odd trackways suggesting pathologies, as well as evidence of social behaviour, particularly care for the young. What did they find? A fair bit: "67 [papers] dealt with healed or healing bone pathologies, mainly fractures, serious enough to have hampered normal functioning (all but eight in dinosaurs); eight described dinosaur tracks showing injury or altered gait; 82 covered dinosaur social behaviour potentially relevant to survival from injuries" (Hearn and Williams 2019, 2).

Hearn and Williams were specifically interested in cases where bones had healed from serious injuries: "Injuries (including subsequent infections) to limbs, shoulder and pelvic girdles, ribs, necks, backs, tails, and jaws, unless described as minor, would impair movement for a significant time, including for predation or escape, breathing and feeding" (Hearn and Williams 2019, 2). The idea is that such injuries indicate various forms of pain-related behaviour. If you just died from the injury, you wouldn't have much chance to spare the limb. Overall, they found that recovery from serious injury was fairly common, identified in 3% of dinosaur species (out of over 1100, so not a small number).

In addition to describing several particularly interesting specimens (for instance, Senter and Juengst 2016), Hearn and Williams pay special attention (rightly, I think) to Ishigaki and Lockley's (2010) discussion of a set of trackways in Morocco's High Atlas Mountains. These belong to bipedal theropods and appear to exhibit a distinctive "limping" gait. Potential limping is noticeable due to an irregularity in the distance between the tracks on one side and the other, a property that is typically regular in theropod trackways. Examining the tracks, Hearn and Williams identify two possibilities:

This evidence suggests that the animal could not support the body weight fully with the digital part of the right foot, and thus it was obliged to use the heel part to support the body weight. A similar, alternative interpretation is to infer that the trackmaker was simply favouring one foot, leg or one side of the body while wading or “labouring” through rather soft mud that impeded its normal progression. (Hearn and Williams 2019, 106)

If indeed we have a limping dinosaur recorded in trackways (and, as they point out, this is contentious!), that seems reasonable evidence of sparing a damaged part. They also cite various pieces of evidence of potential protective behaviour among dinosaurs: herd behaviour, nesting and—as we saw in the last chapter—denning in a burrow.

You might start to wonder what any of this really has to do with dinosaurs feeling pain or exhibiting pain behaviour. Some examples, sure, can be linked to pain avoidance in fairly straightforward ways: dragging an injured leg seems to fit the criteria for sparing a damaged part. However, many seem circumstantial. What would denning have to do with pain-avoidant behaviour, for example? More generally, what does a bunch of data about dinosaur nesting sites, pathologies, and healing have to do with establishing pain-related behaviours?

As warned, I’ve buried the lede here. Hearn and Williams do not understand the relationship between pain-related behaviour and fossils independently of ancestry. First, from the get-go, they point to “evolved and highly conserved systems” (Hearn and Williams 2019, 1) deeply embedded in the history of life, that are related to nominally pain-avoidant behaviour, and they underscore physiological infrastructure related to pain. Second, their argument that various pain-avoidant behaviours in dinosaurs should be equated with feeling pain is based on thinking in terms of ancestry. They point out that “similar pain systems and behaviour among birds or responses to analgesia among reptiles, including opioid receptors in crocodiles, in terms indistinguishable from those used with mammals” (Hearn and Williams 2019, 1). The indistinguishability of pain behaviour and infrastructure between birds, crocodiles, and mammals is relevant because birds and crocodiles “bracket” dinosaurs—they are their closest living relatives, and the common ancestor of birds and crocodiles is also an ancestor of dinosaurs (together they form the group “archosaur”). “While systems and behaviours in extinct species

cannot easily be ascertained, possible forms of these can be found in extant archosaurs” (Hearn and Williams 2019, 2). As dinosaurs are archosaurs as well, insofar as extant archosaurs are representative of that class, crocodiles and birds can be a guide to them.

What’s going on here? The data that Hearn and Williams collate doesn’t speak to nociception in isolation, but only does so when situated with information about dinosaur ancestry. They emphasize the phylogenetic continuity among dinosaurs, birds, and crocodiles as a way of inferring those properties. This is what I’ve called “comparative thinking”: a distinctively biological way of doing business that involves situating a target in its evolutionary history (Baum and Offner 2008; Currie 2021b). This strategy informs thinking about dinosaurs as historical individuals: they are a token lineage that shares properties with currently existing token lineages in virtue of common ancestry; they are, from a certain perspective, parts of the same token lineage.

In virtue of these appeals to ancestry, Hearn and Williams *graft*. By carefully attending to inferred similarities between the past world and the current one—those aspects that are likely not lost—they are able to make headway on a putatively lost world. This is a rather non-serious case (Hearn and Williams frame their paper as a kind of “thought experiment”), but it acts as a comparatively simple example of grafting and allows us to somewhat expand our account of strategic perspectivalism from the last chapter.

Let’s define grafting more carefully. I’ll provide two definitions, one being a subset of the other. *General* grafting appeals to token causal links, while *specific* grafting appeals to historical individuals in particular. I introduce the latter because it is more specific to practices like phylogenetics.

Grafting *generally* involves positing and exploring token causal links among denizens, other times, and the present.

We might understand the general strategy in terms of narratives. Consider the metazoan radiation that occurred during the Cambrian (the “Cambrian explosion”). Some scientists have speculated that the explosion is partly explained by earlier events in the Earth’s history, namely, “snowball events” in which the globe partly or wholly froze over. The thought is that the earlier events constrained metazoan evolution, and that the arrival of balmier conditions triggered the radiation. In this instance, we’re not considering the two past worlds—neither the snowball events nor the Cambrian radiation—in isolation. Rather, we’ve linked them in a particular causal trajectory (I discuss

this in Currie 2014; see also Currie and Sterelny 2017). To do this, we *characterize* denizens as parts of a token causal trajectory, as playing a part in a story, for instance, snowballs acting as constraints on metazoan evolution.

In this chapter, I'll highlight a more specific version of grafting, which generates links by identifying various kinds of historical individuals:

Grafting *specifically* involves characterizing a denizen as a historical individual such that it is continuous with some currently existing historical individual(s) or other past individuals we know about.

Where the general version of grafting concerned token trajectories or narratives, this more specific version concerns a subset of these trajectories: those concerning being part of a historical individual. We should understand “continuity” in terms of whatever persistence conditions we think relevant to that kind of individual, and this might well vary across cases.

Dinosaurs are weird. Many of their properties are lost. But they're not *that* weird: they share many properties with the birds and crocodiles we know and love. Hearn and Williams adopt a perspective, taking dinosaurs as *potentially pain-avoidant* animals. This enables inferred behaviours from fossils, such as a limping gait, to be relevant to their possession of nociception. But they also adopt another perspective, taking dinosaurs as archosaurs, recognizing continuity among themselves, birds, and crocodiles (just as we took mammoths as Elephantidae, and thus continuous with living elephants). These joint perspectives—seeing dinosaurs as *potentially pain-avoidant archosaurs*—make data and interpretations from fossils evidentially relevant to dinosaur pain avoidance.

Note two things. First, consider the relationship among inferred similarities, the continuity in the graft, and the evidential meaning of the data. The graft identifies a historical individual, say, archosaurs. That enables phylogenetic inferences, for instance: we know dinosaurs have opioid receptors because crocodiles do. And further, the graft makes sense of fossil data: fossils bearing healed pathology are evidence of pain avoidance, partly in light of the healed pathology signalling the same in living animals. It is comparisons with living relatives (and much of the rest of life) that make sense of specimens and what we can infer from them. There is, then, an interdependence among the perspective (how the denizen is characterized), how it is grafted

(what it is taken to be conjoined with as a historical individual), and the relevance of data (fossils, in this instance).

Second, notice that the various perspectives—characterizations—intermingle and overlap. Hearn and Williams characterize dinosaurs as *pain-avoidant*, an ahistorical kind, and as an archosaur, a historical individual. But these are not independent characterizations: dinosaurs are *pain-avoidant archosaurs*. Under the right conditions, perspectives overlap, are combinable, and make empirical contact with one another: this is another element of their pluripotency.

Finally, grafting is closely related to a much more discussed feature of historical science: common-cause explanation (e.g., Cleland 2013). A common-cause explanation makes sense of some correlation by, well, positing a cause common to the correlates. No doubt explanations referring to common causes, and inferences following common causes, are important ways of characterizing some of the epistemic properties of grafting, particularly in its specific historical-individual guise. However, my focus here is less on providing an abstract characterization of the justification of such inferences. I'm interested in how the strategy underwrites tackling loss and putative loss.

2. Stems and Failed Experiments

“Understanding Broadfoot as a mammal challenges some common misconceptions about mammals. We biologists sometimes identify traits as ‘primitive,’ but this isn’t supposed to be derogatory. It just means ‘basal’, that is, present in some common ancestor. However . . .”

Ichthy pauses, glaring over their spectacles with the air of a disapproving (not angry, just disappointed) parent.

“Some call mammals ‘primitive’ in a stronger sense, implying they are ‘failed evolutionary experiments,’ dismissed as nature’s mistakes. Some say they were inevitably replaced by us sophisticated, clever, and dynamic reptiles. But discoveries like Broadfoot put these claims in their place. Mammals were diverse and well-adapted. They were, like us, sophisticated, clever, and dynamic creatures. Indeed, I see no reason, considering mammalian diversity, cognition, or form, to think them incapable of any of the same evolutionary achievements as us reptiles.”




ILLUSTRATION 4.1 Prof. Platy Regards the Ichthyosaur

Prof. Ichthy allows themselves a whimsical smile. “I’ve sometimes imagined what a mammal with ichthyosaur-like intelligence might be like. Maybe in some alternative reality, a mammal, perhaps Broadfoot even, is engaged in the scientific study of a fossilized ichthyosaur! A rather unscientific idea, of course. But I don’t think a sentient mammal is an impossible idea. Nothing about mammals says they couldn’t or wouldn’t evolve reptile-like intelligence. My thought is just this: categorizing mammals as a failed experiment is to undersell their diversity and evolutionary potential. That we are here and they are so diminished, it seems to me, is not due to something lacking on their part, but due to the whims of evolution and history.

“So, contextualizing Broadfoot with other fossil and living mammals helps us see how rich that lineage is—and thus push back against preconceptions of a ‘primitive’ evolutionary past.”

I’ve introduced grafting as a strategy applicable to lost, or putatively lost, worlds. To further explore this strategy and the nature of loss, I want to take a little conceptual tangent into the notion of a “failed evolutionary experiment.” As we’ll see, grafting strategies have the capacity to transform both how we conceive of lost worlds and their epistemic value.

In the last chapter, Cuthill and Conway Morris emphasized the phylogenetic continuity of rangeomorphs while simultaneously asserting the richness

and diversity of the Ediacaran fauna more generally. It's worth quoting them again: "The Ediacaran biota was far from homogenous but instead included diverse phylogenetic lineages and body plans" (Cuthill and Conway Morris 2014, 13124). What makes this diversity worth noting? Part of the story, I suspect, is analogous to Prof. Ichthy's pushback against mammals being characterized as "failed experiments." No, Ichthy insists, the mammals were not a dead end, were not a failed experiment. They were a thriving, rich lineage worthy of understanding on their own terms. Ichthy is here mimicking what has been said of Ediacaran fauna: these strange, often sessile communities were an evolutionary mistake. After all, the metazoans hadn't yet got their act together, and whatever these failed experiments amounted to, they were washed away by the sweep of the Cambrian explosion when metazoans really got going. Interestingly, the Cambrian and the Avalonian have been considered analogous insofar as both are "evolutionary failures": "the Avalon explosion represents an independent, failed experiment with an evolutionary pattern similar to that of the Cambrian explosion" (Shen et al. 2008, 84).

"Failed experiment" claims typically involve a phylogenetic element. As Dunn and Liu have put it, "Members of the Ediacaran macrobiota have been . . . deliberately set apart . . . by suggestions that they were either 'failed experiments' in the history of life, or members of long-extinct higher-order clades" (Dunn and Liu 2019, 512). What is it to be a "long-extinct higher-order clade"? Let's dive into some of the basics of cladism, which will allow us to better understand the notion of a "failed experiment."

One way of understanding "clades" is via *bracketing*, a notion we met in the last section. Take two extant lineages (say, birds and crocodiles), draw two lines backwards in time, along parents, grandparents, etc., to their common ancestor (around 240 million years ago; Green et al. 2014), and then move forwards again, including every lineage that also shares that common ancestor. Performing this procedure with birds and crocodiles produces the category "Archosauria": enfolding not only birds and crocodiles, but also many charismatic Mesozoic denizens, such as dinosaurs and pterosaurs (see Figure 4.1).

As we saw in our discussion of dinosaur pain, dinosaurs being bracketed by crocodilians and birds plays an important role in inferring various properties of theirs. Membership of that ancestral grouping—the historical individual Archosauria—allowed various features of dinosaur fossils (such as healed pathologies), and behaviours inferred from fossils (such as denning) to be made sense of, and made evidentially relevant to, pain-avoidant behaviour in dinosaurs. A little more poetically, this grouping also enables dinosaurs





FIGURE 4.1 Bishop et al.'s (2020) Simplified Archosaur Cladogram (Fig. 1 in Bishop et al.). Creative Commons CC BY

to be part of life's story: in virtue of their placement in Archosauria, they are relevant to understanding and constraining the evolutionary and developmental narrative of the archosaurs.

With (very) basic cladism under our belts, let's return to Dunn and Liu's argument that the Ediacaran fauna have been done wrong by scientists. Consider the upshots of (1) taking the Ediacaran fauna to be a single group which (2) belongs to an extinct higher-order clade. We end up with something like Figure 4.2. The fauna of the Ediacaran cluster together as one clade, and, as a "higher-order" clade, they drifted from the main evolutionary narrative very early on. Here, the Ediacaran fauna are more or less excluded from the history of life: it isn't obvious that they can play a role in informing how we understand current life, and current life has very little to tell us about that fauna.

Figure 4.2 is, in a sense, a cladistic representation of a failed experiment: the Ediacaran fauna are taken as a single group that left no descendants. Further, they are not only lost but fatally disconnected from life more generally: *prima facie*, they have nothing to tell us of real interest about the history of life. Disvaluing extinct fauna has a storied history, from reinforcing a progressive view of history by taking dinosaurs as paradigmatically "extinct" monsters (Rieppel 2020), to reinforcing sapiens exceptionalism by representing Neanderthals as primitive lumbering brutes (Stiner 2017; Meneganzin and Killin 2024). As I'll argue throughout, but particularly in chapter 10, taking loss seriously often involves challenging not just how we think the past was, but the ways in which we value it.

FIGURE 4.2 Informal Cladogram
Depicting the Ediacaran Fauna as a
Dead End or Failed Experiment



Dunn and Liu argue that representing the Ediacaran fauna in ways similar to Figure 4.2 is a mistake, and understanding why is illuminating of grafting. Here’s how they state their position: “We advocate abandoning the failed experiment perspective and embracing phylogenetic thinking in order to make progress in determining the phylogenetic positions of these organisms and realizing their evolutionary significance” (Dunn and Liu 2019, 512). But what could this mean? Figure 4.2 is a cladogram (well, of sorts), so it does seem to constitute phylogenetic thinking of some kind. To see what Dunn and Liu have in mind requires a small—but surprisingly important—conceptual addition to how we’re thinking about clades.

A clade consists of some common ancestor and all of its descendants. A common way of identifying clades, as illustrated earlier, involves identifying at least two extant lineages, their common ancestor, and all other descendants sharing that common ancestor. This conception of “clade” in fact picks out what palaeontologists would call a “crown group.” Add to your conceptual repertoire another notion, that of a *stem group*. Recall that all life on Earth comes from a single origination event. This means that for any two organisms, there will be a common ancestor: for instance, there will be some common ancestor of ourselves and the Ediacaran fauna. A clade’s *stem group* includes some ancestor and all of its descendants *minus the crown group* (Briggs and Fortey 2005; Bryse 2008). I find that definition difficult to




FIGURE 4.3
Cladogram
Depicting Two
Possible Stem
Groups for the
Ediacaran Fauna

wrap my head around. Here's one way of getting it. Let's add a third lineage to Figure 4.2, say, fungi. We thus imagine adding the common ancestor of fungi and metazoans, representing a larger clade that includes both fungi and metazoans (the Opisthokonts, in case you're interested, which will include non-metazoan organisms). With this in place, a further question arises: are the Ediacaran fauna on the line leading from the Opisthokont common ancestor to the fungi, or to the metazoans? We can represent that question with something like Figure 4.3.

Notice that how we answer this question makes a big difference to how Ediacaran fauna fit into the history of life: if the denizens of the Ediacaran are stem-group fungi, then they are part of that narrative, and can tell us something both about the base of the Opisthokonts and about how the fungi evolved. But if they are stem-group metazoans (which, as we'll see, palaeontologists now think many of them are), then they could be crucial for understanding the evolution of metazoans. After all, they are part of the story, dividing fungi and metazoa. So, simply by drifting somewhat further back in time, by adding the notion of a stem group to our phylogenetic repertoire, we transform how we characterize the Ediacaran fauna, what might be relevant for understanding them, and how they might help us understand broader-scale evolutionary patterns.

Further, as Dunn and Liu point out, considering the Ediacara as a possible stem group has occurred alongside palaeontologists taking a more fine-grained taxonomic perspective. These more recent studies often suggest, as we saw in Cuthill and Conway Morris' analysis last chapter, that the taxonomic status of these critters is neither as simple nor as monolithic as Figure 4.2 represents. So, they do not form an easily dismissed phylogenetic clump.

In the next section, we'll look at a case in more detail, but I want to wrap this section up with a few lessons moving forward.

First, the strategy I've described—of switching from thinking about a lost lineage as a failed experiment to thinking of it as a potential stem group—is an example of *grafting*. After all—and as we'll see in more detail in the next section—to determine where some lineage might sit in a stem requires emphasizing points of continuity between the putative stem and the relevant crown group. We might ask, what characters do the Ediacaran fauna and the metazoans have in common that sets them apart from fungi (and vice versa)? So, in light of an apparently mysterious lineage, one apparently unlike anything we've seen, we seek to find principled strategies to determine ways those lineages might not be lost, and thus be incorporable into an overall narrative.

Second, notice that we didn't make any directly empirical moves in shifting between Figure 4.2 and Figure 4.3: we added another node to the cladogram, and thus generated another common ancestor. You might say, we simply changed perspective. This is a general feature of phylogenetic thinking: although there are facts of the matter about which lineages are more related to which, whether two lineages are part of the same clade turns on which clade we're talking about. Considering the crown group including ourselves and crocodiles, we're in the same clade as the dinosaurs; considering the crown group including ourselves and cats, we're not in the same clade as dinosaurs. So, when using this kind of thinking to make sense of various lineages, we can—as it were—zoom in and out of various levels of description to get the best grip on the traits we're interested in.

You don't get much more strategically perspectival than that!

Third, and as a perhaps-too-late caveat, my description of phylogenetic, cladistic thinking has been purposefully casual: apologies to any phylogeny nerds out there. I've not said anything about the inferential strategies that go into constructing phylogenies (e.g., Sober 2002), nor about the complex analyses required to understand characters (let alone what counts as a character: DiFrisco, Love, and Wagner 2020; Meneganzin, Ramsey, and DiFrisco 2024). Further, I've only hinted at the rich ways that these underwrite comparative

methods in biology (Currie 2021b). Max Dresow and I are developing an argument that the palaeontological innovation of the stem group transforms how fossils can be treated in phylogenetic contexts, and that involves a much more sophisticated treatment than I can manage here.

Regardless of all that, we can begin to see the grafting strategy as it plays out in evolutionary palaeontology: use phylogenetic thinking to emphasize the continuity between lineages with lost traits and better-known lineages, thus finding lines of ancestry between them, seeing them as historical individuals, and thus providing a perspective to start building understanding. Let's turn to a more specific case to see this in action.

3. *Dickinsonia*

Let's wind the clock forward, past the Avalonian, and later into the Ediacaran. As the clock's hands spin, we wave goodbye to the microbial mat around the same time as burrowing metazoans arise and the oceans start to become more reliably oxygenated. The sessile Eden of the rangeomorphs gives way as worm-like creatures evolve. As they eradicate the microbial mat, we must also bid farewell to pizza disks in the subsequent fossil record. We'll return to the fall of the Ediacaran in the next chapter, but for now, let's focus on perhaps the most well-known of the Ediacaran denizens: *Dickinsonia*.

The Ediacaran period is named after the old rocks of the Ediacara Hills, nestled in the Flinders Ranges of South Australia. It was there in the 1940s that *Dickinsonia* was identified by Reg Sprigg. Its fossils are oval, nearly perfectly bilateral imprints with rather beautiful lines radiating from a central axis (Figure 4.4). *Dickinsonia* varies in length from a few insignificant millimetres to a mighty near meter and a half. Its fossils have also been found in later Ediacaran beds in Russia, Ukraine, and China (Brasier and Antcliffe 2008; Ivantsov and Zakrevskaya 2022).

As is typical for Ediacaran fauna, *Dickinsonia* has long been a phylogenetic mystery (Sprigg originally thought it was a jellyfish, but over the years it has been described as everything from a fungus to a worm to something entirely unknown). Gold et al. (2015) describe the basic problem with trying to place *Dickinsonia* in the history of life: "Part of the difficulty in classifying *Dickinsonia* comes from its morphological simplicity—which would be consistent with its placement as a primitive organism—but also means that it lacks morphological synapomorphies to ally it with any living (crown) group" (Gold et al. 2015, 315).



FIGURE 4.4 *Dickinsonia* Fossil (Detail from Fig. 1, Bobrovskiy et al. 2018). Creative Commons CC BY

In the last section, I informally introduced us to phylogenetic thinking, which, for the purposes of this book, is most interesting as an example of grafting. That is, it allows us to conceive of various critters as historical individuals—clades—and thus find continuity among them, continuity that might underwrite various inferences. But to apply phylogenetic analysis, we need *characters*—that is, traits possessed by the relevant lineages that can be used to empirically inform who is more related to whom. Dinosaurs and birds share various characteristics, such as three-toed feet and pneumatized bones, which is evidence of their shared ancestry. Now, look again at *Dickinsonia* in Figure 4.4. What morphological characters does *Dickinsonia* have? On what basis might we think it similar to or different from any living animal? As Gold et al. reference in the quote above, in modern phylogenetic analysis, we need to identify “synapomorphies.” A synapomorphy is a character *uniquely* possessed by a common ancestor and its descendants—it is thus a trait that can distinguish clade membership from non-membership. For instance, the human chin has long been identified as a distinctive morphological trait among the primates; that is, we have it, but no other primate does. It is thus, potentially at least, a character that might be a synapomorphy, dividing our own taxon from that of our relatives (Meneganzin, Ramsey, and DiFrisco 2024).

What potential synapomorphies does *Dickinsonia* have? Well, to answer that question, we need to first ask another: What *characters* does it have? With so few features preserved, identifying meaningful morphological characters is extremely challenging. Gold et al.'s strategy is to find characters by appealing not to morphology, but to ontogeny: how *Dickinsonia* grows. That is, instead of looking for a morphological character, they seek out a developmental character (perhaps it should be called a "synapontogeny" . . .).³

The first step is to recognize that *Dickinsonia* grows via serially repeated structures. This enables us to characterize *Dickinsonia* as a critter that grows via serial repetition. This perspective underwrites a set of comparisons with various eukaryotes that also grow that way. As we'll see in the next chapter, this is a form of *anchoring*, so I'll leave a thorough description of that aspect until then. Suffice to say, Gold et al. argue that *Dickinsonia* grows in a very particular way: "A form of terminal addition, a developmental mechanism currently exclusive to bilaterian animals" (Gold et al. 2015, 316). If so, this suggests that *Dickinsonia* should be phylogenetically aligned with those animals, that is, "terminal addition"—a developmental character—is a synapomorphy among bilaterian metazoans.

Gold et al. generated a set of phylogenetic trees and ancestral state reconstructions, the former capturing patterns of relatedness and the latter identifying which traits belong to which common ancestors (let's not worry about the technical details here). In accordance with many other studies, the crown group of bilateral animals bottoms out sometime in the Cambrian; that is, the common ancestor of all currently living complex animals lived sometime after the Ediacaran and *Dickinsonia* had passed. Crucially, "All analyses agree that the probability of terminal addition existing in the last common ancestors of animals is less than 5%" (Gold et al. 2015, 319). So, the common ancestor of all animals didn't grow that way, which is to say, that mode of growth is not a "primitive" (or "basal") trait inherited by animals (and lost in some), but is rather "derived." However, their results do "suggest there is strong support for terminal addition being the ancestral condition of the bilaterian animals, and that there is little support for this developmental trait existing at deeper nodes of the eukaryote tree" (Gold et al. 2015, 320). In other words, basal bilaterian animals grew by terminal growth, and this emerged within that clade, not earlier.

3 Kirsten Walsh takes the blame for this neologism . . .

What does all this tell us? Growth via terminal addition in animals evolved in the bilateran animals, not before. *Dickinsonia* is at the very least allied with animals, and it grows via terminal addition. So, unless it evolved that character independently, it should be placed in the stem of bilateral animals. It is part of the story (even if going on something of an evolutionary tangent) of how we, bilaterally symmetric animals, got our start. Terminal-growth addition is a developmental synapomorphy in bilateral animals and *Dickinsonia*.

What is important about getting Dickinsonia's phylogenetic placement right? For Gold et al., "Our inability to place Ediacaran fossils onto the eukaryotic tree has impeded our attempts to constrain molecular clocks, assess the intensity of Earth's earliest extinction events, and interpret the significance of the Cambrian 'explosion' of animal life" (Gold et al. 2015, 315). Molecular clocks (roughly, figuring out the time between extant lineages' ancestral splits from one another using estimates of the pace of mutations in their respective genomes) need to be calibrated against timing derived from fossils. If we don't know where to place fossil taxa in phylogenetic trees, they can't be used for calibration. Further, if we don't know where Ediacaran fauna place in the history of life, we can't investigate the effects of major extinctions and other macroevolutionary events across various branches. For instance, we might wonder whether the Cambrian explosion represents the descendants of a few lucky survivors from the Ediacaran, or if there was a smoother transition between the two. Gold et al. take their view to support something like this latter view: "Large bilaterans were an important component of the Ediacaran biota, thus connecting this fauna with the following radiation of animal life in the Cambrian" (Gold et al. 2015, 322). Thus, this specific attempt at grafting also underwrites a more general instance of the strategy: folding *Dickinsonia* into the narrative of metazoan life.

I've simplified Gold et al.'s work for the purpose of highlighting grafting. In the face of a strange, unknown—putatively lost—lineage, Gold et al. hunt for particular traits possessed by the lineage that could be meaningfully brought into dialogue with taxa we are more familiar with. Morphological traits were unpromising, partly because of *Dickinsonia*'s apparently simple, lost body plan and because of the erasure of what useful morphology it may have had. But we do have information about how the creature grew (or at least Gold et al. argue that we have such information), and so we can hunt for ontogenetic synapomorphies instead of morphological ones. This echoes the strategy we saw for artificing in the last chapter: developmental features

are isolated and investigated before being recontextualized, in this case using various phylogenetic methods.

So, when faced with a lost world, try to find aspects of that world that are not lost and weave them into a historical narrative. This brings denizens into productive dialogue with time periods we have better information about. That is the essence of grafting strategies.

With the ideas underlying grafting in hand, I want to turn to two related discussions. The first, in the next subsection, expands on strategic perspectivalism with an initial pass at disagreement. The second, and this chapter's conclusion, will turn to what might be a rather extreme example of grafting in biology (the companion to chapter 3's discussion of de-extinction): living fossils.

4. Conflicting Perspectives

I've been emphasizing the perspectival nature of palaeoscientific knowledge: what counts as a past world is at least partially up to scientists' interests; the distinction between loss and erasure turns in part on which denizens scientists indicate; how denizens are characterized turns on the perspectival tools at play. You might also accuse me of infusing strategic perspectivalism with a fair bit of epistemic optimism (Currie 2019c): I want to say that the clever use of perspectives and the tools that generate perspectives explains how progress can be made in light of loss and putative loss.

But, you might worry, there are two closely related epistemic problems with a perspectival account, even a merely strategic one. If the results generated by perspectives are tied to those perspectives, then (1) how can that knowledge travel beyond the contexts of those perspectival tools, and (2) how can apparent disagreements be resolved? This worry should be familiar to philosophers of science, echoing as it does the more extreme worries about incommensurability associated with Thomas Kuhn.⁴ Less extremely, these worries often arise for pluralistic accounts of representation, as Massimi summarizes: "Whenever there are situations requiring a plurality of models for a given phenomenon, each model might accurately represent some aspects of the phenomenon at the cost of blurring, idealizing, or distorting other aspects. The outcome is that different models may deliver different,

⁴ Although incommensurability needn't be understood in such stringent ways (see Bocchi 2022, for instance).

sometimes incompatible, or even inconsistent images for the same phenomenon” (Massimi 2022, 53).

Notions of pluralism, perspectivalism, and what to do with inconsistent models are long in the tooth in the philosophy of science (Wimsatt 1987; Weisberg 2007a; Potochnik 2010; Massimi 2021). While I’ll return to these themes, in this context I’ll focus on how apparently conflicting perspectives might relate to uncovering lost worlds.

The possibility of inconsistency between models (or perspectives) and the related apparent difficulties of transporting knowledge from a perspective or resolving disputes and disagreements arising from differing perspectives share a common root. I’ll call this root *perspectival isolation*. In some circumstances, particularly those favoured by more traditional perspectivalists and pluralists, perspectival isolation is not a problem, but a solution, due to the limitations of any particular representation or perspective. Here’s Elgin’s version of this strategy: “Different models make manifest different features of the target. Diverging models afford different perspectives on the same reality. And it is no surprise that different perspectives reveal different aspects of that reality. There is no optimal model for the same reason that there is no optimal perspective. Every perspective, in revealing some things, inevitably occludes others” (Elgin 2017, 180).

So, one—and, I think, the received—way of handling apparent disagreement is to follow this kind of divide-and-conquer strategy. Models do not truly disagree, but rather imperfectly describe different elements of the same reality. The model’s representative capacities are thus, in a sense, isolated from one another. I’m not sure if this strategy will cut it in many of the circumstances I’m interested in. Scientists faced with lost worlds no doubt flexibly adopt various complex perspectives as they attempt to get a grip on lost denizens, but these perspectives are far too much in dialogue, and indeed sometimes in true disagreement, for this brand of ecumenicism to have wide applicability. So, toward the next chapter’s discussion of iterativity, understanding how different perspectives might speak to one another and, yes, disagree with one another is crucial.

Insofar as perspectives are isolated, that is, self-contained and without overlaps with other perspectives, it is hard to see how disagreement might be resolved (or, for that matter, how perspectives might speak to one another in the first place). However, we’ve already seen in this chapter that perspectives are often combinable: dinosaurs are not only characterized as archosaurs and pain-avoidant critters in isolation, but as *pain-avoidant archosaurs*. In the

next chapter, I'll argue that forms of epistemic iteration overcome perspectival isolation in practice. As a warm-up to that argument, let's look at a putative disagreement about the phylogeny of *Dickinsonia*.

As we've seen, Gold et al.'s analysis of *Dickinsonia* turned on its growing via terminal addition, which was inferred mostly via phylogenetic bracketing. Hoekzema et al. (2017) dispute this interpretation, arguing that the lineage grew in a two-step manner, first via insertion pre-terminally with some inflation, and then a second step dominated by inflation. That is, the early stage of *Dickinsonia*'s growth mostly involves adding components, while the later stage largely focuses on increasing the size of those components. Crucially, their model denies that the lineage grew via terminal addition.

As opposed to phylogenetic bracketing, Hoekzema et al.'s evidence is based on examining a sequence of *Dickinsonia* fossils, interpreted as varying across ontogenetic stages. They use these to reconstruct and analyse the life history of *Dickinsonia*, thus testing different models of its development.

On the face of it, we have here the use of two different sets of perspectival tools. Where Gold et al. develop a novel developmental character and employ it via phylogenetic bracketing, Hoekzema et al. use an ontogenetic interpretation of fossils to generate an apparently conflicting story. If the arguments are made based on differing tools—and thus differing perspectives—can this avoid perspectival isolation? That is, is there a substantive empirical dispute between them? I want to insist yes, and to see why, we'll need to fill out Hoekzema et al.'s argument in more detail.

As we saw, Gold et al. argued for a bilaterian affinity for *Dickinsonia*, and Hoekzema et al. are explicitly in the business of disagreeing with their interpretation: "Recent studies into growth in *Dickinsonia costata*, and arguments for a bilaterian affinity based on ancestral state recognition, rely on assumptions regarding growth in this taxon that we here argue are incorrect" (Hoekzema et al. 2017, 2). How does this argument go? Everyone agrees that *Dickinsonia* grows via the successive insertion of units around a central axis, where each unit is, in effect, one of the lines crossing the centre of the distinctive morphology visible in Figure 4.4. In addition to the units, we can also identify two "ends" of *Dickinsonia*: the "deltoidal" region (or "D"), where the lines curve outward to form what is often interpreted as a head, and the "anti-deltoidal region" (or "AD"), which is typically interpreted as the end where insertion occurs (see Figure 4.5).

Hoekzema et al. measure three things across their specimens: (1) unit length, (2) unit count—the total number of units—and (3) unit "number."




FIGURE 4.5 Measuring *Dickinsonia* Growth (Detail from Fig. 1, Hoekzema et al. 2017). Reprinted with Permission from Royal Society B

This last measurement needs some unpacking. They describe it thusly: unit number is the “progressive number of appearance of each unit, considering the possibility of generative zones at either the D-end or the AD-end of the organism” (Hoekzema et al. 2017, 3). “Unit number” is the order in which that unit would have been added if the addition were at the anti-deltoidal end or the deltoidal region. The former, taking the growth to occur anti-deltoidally, is the assumption behind terminal-growth, and is supported by the fact that the smallest units are located there. But Hoekzema et al.’s measurements allow them to compare the relationship between unit length and count against two versions of unit number, one starting from the D-end and another from the AD-end, which can be understood through an overall ontogenetic model. At base, they can ask, which insertion point makes the most likely developmental sequence?

Hoekzema et al. compare the developmental sequence taking the (usually assumed) AD-end insertion route, that is, terminal addition, and the D-end insertion. The major difference between the two is the regularity of unit

length and unit counts against overall size. We would expect the youngest unit, that is, the first one inserted, to be a similar size across specimens, but on AD-end insertion, we see nothing of the sort. By contrast, if insertion occurs at the deltoid, “New units neatly and consistently exhibit increasingly greater lengths as they are added” (Hoekzema et al. 2017, 6). On this basis, then, Hoekzema et al. argue that there is good reason to take *Dickinsonia* to grow not by terminal addition, but by deltoidal addition.

Hoekzema et al. abstract from their empirical measurements to generate a simplified growth model that they claim can generate a variety of *Dickinsonia* morphologies, a result that they further argue has phylogenetic consequences: “Although different *Dickinsonia* species have disparate morphologies, they can be rationalized by a common morphometric model, substantiating their coherence as a natural group” (Hoekzema et al. 2017, 6). We saw this kind of inference earlier in Cuthill and Conway Morris’ unified developmental model for rangeomorphs in chapter 3. Producing a simplified, unified developmental model is taken as evidence for phylogenetic continuity. This is an example of how grafting and artificing can come together.

So, Hoekzema et al. and Gold et al., in a sense, have differing characterizations—perspectives—on *Dickinsonia*. Each incorporates different perspectival tools. Where Gold et al. incorporate (1) adult specimen morphology with (2) a presumed growth structure and (3) ancestral-state reconstruction, Hoekzema et al. use (1) measurements of *Dickinsonia* across ontogeny, which (2) allow a comparison of two growth models and (3) the generation of a theoretical geometric model. It might be tempting to think that these two perspectives—both made available in the horizon of alternatives thanks to those tools—are simply ones that scientists can select; and so, the disagreement is not truly empirical. After all, isn’t it simply which perspective we choose to adopt that determines our endpoint? Massimi characterizes this as “perspectival disagreement”: “Perspectival disagreement envisages the possibility that two or more epistemic communities from different scientific perspectives might agree about either the reliable experimental procedures or the justificatory principles, while also disagreeing on which scientific knowledge claims these in turn support” (Massimi 2021, 561–562).

I take it that perspectival disagreement occurs when scientists adopt the same characterization of a denizen—use the same perspectival tools—but disagree about, say, what a satisfying explanation looks like, or what the right kind of question to ask is. As mentioned in the last chapter, much discussion of perspectivalism targets this kind of problem. And I do think, if you’re

interested in resolving that question, it does present a philosophical problem. However, I don't see much—or any—perspectival disagreement in the cases I've been examining.

For instance, compare the differing characterizations of historical individuals we've examined in phylogenetics. Here, although there is a fact of the matter regarding the cladistic relationships between taxa once they are characterized, there is nothing at all stopping us from combining these. They are different perspectives insofar as they involve different characterizations of the lineage, but in this instance, a single representation can capture these differences, namely, a phylogenetic tree.⁵ By contrast, Gold et al. and Hoekzema et al.'s differences diverge from this: they legitimately disagree. Their two perspectives are not simply talking past one another. The presumed growth structure in Gold et al. is more or less one of the growth models that Hoekzema et al. compare in their study. An initial point to realize is that although the perspectives differ, they also overlap.

If the two perspectives involve differing characterizations of the denizens, and the results they generate clash, we might ask whether the clash is due to more than differences in those characterizations. That is, can the disagreement be considered truly empirical? Yes, it can. Hoekzema et al. have provided arguments that undermine Gold et al.'s interpretations, ones that they or others could develop responses to. For instance, one might argue that Hoekzema et al.'s approach more directly tests the developmental sequence of *Dickinsonia*, while Gold et al. depend on more circumstantial tests, and we might then take it as stronger evidence. Further, as I'll explain in the next chapter, that disagreement itself can act as a launch pad for further studies.

Seeing why differences in perspectives needn't undermine productive empirical disagreement provides a vital clue to our story about how challenges from loss are tackled: we learn new things from adopting strategic perspectives, but that new knowledge is not limited to those perspectives and, as such, the knowledge is not perspective-bound, and the perspectives are not immune to empirical constraint or testing.

Recall Massimi's perspectival₂: perspectives enable inferences that reach beyond what is explicitly represented in that perspective. I take it that the

5 In instances of *discordance*, that is, when apparently nested lineages, such as species and parts of their genomes, have differing phylogenetic histories (Haber 2019), a single phylogenetic representation may be unable to easily capture the state of play. I don't think this is an instance of perspectival disagreement, however, just a case of limited representational capacity.

features of recontextualization and pluripotency that I articulated in the last chapter, and we see present in the case at hand, make further sense of how this might work. It isn't only that a perspective affords inferences; it is also that perspectives can be brought into dialogue: constraining, testing, and integrating with one another.

It is worth returning to the thought that under some conditions, perspectives may be combined. Again, Hearn and Williams didn't only conceive of dinosaurs as potentially pain-avoidant, or as archosaurs, in a sequential or isolated fashion: rather, they conceived of them as pain-avoidant archosaurs. Such mixed perspectives are common and crucial for understanding how lost worlds might be reconstructed. In visual cases (if we're to believe early 20th-century psychologists), a gestalt switch occurs as we swap between differing perspectives—call to mind the philosopher's favourite chimera, the duck-rabbit. But here we have a disanalogy with scientific perspectives. Our (or at least my!) visual perceptual tools disallow seeing the duck-rabbit as a duck and a rabbit simultaneously. But phylogeny allows us to represent both ancestral information about historical individuals—being an archosaur—as well as ahistorical kinds, such as having pain-avoidant behaviour. Thus, some perspectival tools not only enable perspectives to be put in empirical tension due to overlaps, but they can also sometimes act to integrate perspectives.

To summarize in somewhat of a rush: some facts are generated/graspable/confirmable from a perspective, but are not perspective-bound, as they may, for instance, be taken into (or even themselves form) new perspectives, as we often see in the “characterize-examine-recontextualize” process. This means that by shifting between perspectives, we can cobble together information that isn't available otherwise. In other words, even though the fossils are very distant from the past world, and the past world is very different from now (is lost), under the right conditions, we can nonetheless recover a lot of knowledge about it.

5. Coda: Living Fossils

The grad student switches transparencies again, this time revealing an organism familiar to the audience. The animal looks, at first glance at least, like a rather unremarkable lizard. The slide's appearance initiates a small cheer from a group of long, sleek, eel-like pleurosaurs in attendance. Prof. Ichthy gives them a small wave before continuing.

"I've been explaining that although mammals and Broadfoot share many properties, nothing too much like Broadfoot has survived. It's a unique extinct lineage. Which is to say, there's no analogy for what we marine reptiles have in our familiar old friend Peaks-on-Back."

Ichthy again nods at the pleurosaurs before turning to consider the image behind them.

"Peaks-on-back is a strange, rather wonderful animal. While clearly related to pleurosaurs, it nonetheless retains a number of highly primitive, conserved traits. Many of its traits are otherwise only known from fossils. Indeed, back in the day, the discovery of Peaks-on-Back certainly caused quite a scientific stir!" Ichthy chuckles to himself.

"If you'll permit me some rather unscientific speculation, were there scientists in the Age of Mammals, they may indeed have wondered after Broadfoot as we do after Peaks-on-Back. The mammals we know and love, those still surviving today, are what we call "placental mammals." But Broadfoot seems to be something quite different: only one of two genera we have so far found which represent the "monotreme" group. I can imagine a mammalian scientist, upon discovering a living Broadfoot, thinking that a stranger from an earlier time had visited them. Imagine, something known only from fossils just wandering out of the bushes. A lost world appearing before you. A fossil, but alive."

Throughout this chapter, a picture has begun to emerge concerning how grafting—that is, emphasizing historical continuities between past and current worlds—can make progress in cases of loss. In the face of a putative lost world—that is to say, a denizen with certain unique, unknown, alien properties—we seek out continuities in token historical individuals. This involves characterizing our target in ways that enable grafting: dinosaurs are bracketed by birds and crocodiles, rangeomorphs are stem lineages—but of fungi or animals?—*Dickinsonia* shares a growth strategy with animals. These practices of characterization and recontextualization enable points of comparison that further underwrite situating the denizen in some way: within the evolutionary narrative implied by a cladogram, within models of developmental processes, and so on. These situating practices extend the horizon of alternatives, that is, the perspectives available to scientists. But the components of these horizons are not independent: they partially overlap in ways that enable empirical dialogue between them. Thus, by flexibly switching between cases,

the lost property can slowly emerge. In the next chapter, we'll concretize this further by developing a notion of "iteration," but before getting there, I want to briefly turn to "living fossils," as they might be thought of as an extreme example of grafting.

Although Prof. Ichthy could not possibly know, his speculations are right: the (European) discovery of the platypus at the turn of the 19th century caused various kinds of stirs. Darwin infamously took the platypus to be a living fossil (Gruber 1991; Moyal and Marks 2019):

In fresh water we find some of the most anomalous forms now known in the world, as the *Ornithorhynchus* and *Lepidosiren*, which, like fossils, connect to a certain extent orders now widely separated in the natural scale. These anomalous forms may almost be called living fossils; they have endured to the present day, from having inhabited a confined area, and from having thus been exposed to less severe competition. (Darwin 1964, 107)

Darwin pointed to the platypus (*Ornithorhynchus*) and lung-fishes (*Lepidosiren*) as examples of forms that are anomalous insofar as they are very different from other living animals—widely "separated in the natural scale"—such that they most closely resemble taxa only known from fossils. Darwin's commitment to a linear, incremental, and progressive conception of evolution led him to explain these "living fossils" by appeal to "refugia": protected places where a lack of competition has enabled their less-adapted forms to survive. Further, Darwin's commitment to there being transitional forms made these "living fossils" potentially important for bolstering that aspect of his theory (Turner 2019b has a nice explanation of the importance of living fossils for Darwin, as well as the term's origin and pre-Darwinian use; see also Werth and Shear 2014).

What is a "living fossil"? It is more or less accepted that the uses of the concept across biology are unified at best by family resemblance. That is, the term picks out a set of non-equivalent notions—being rare, being morphologically similar to fossils, possessing evolutionary stasis, being at biodiversity risk due to smaller geographical range or populations, possessing basal characters, etc.—of which lineages called "living fossils" possess some but not all. There is disagreement as to the usefulness of the term, how the concept should be understood and deployed, and what the concept is for. Two quotes

from Werth and Shear, I think, capture the problem and the promise. First, the problem: “We must not expect that all aspects of coelacanth or horseshoe crab biology have existed unchanged for hundreds of millions of years. Alas, this dashes Darwin’s hopes that living fossil taxa might provide unsullied knowledge about the world as it existed long ago” (Werth and Shear 2014, 3).

That is, living fossils cannot be understood as fortuitous time machines. Darwin is making a mistake in thinking living fossils are “transitional forms.” This is because they themselves are products of millions of years of evolution in their own right. A platypus, for instance, is not an ancient mammal, but a lineage that has followed its own evolutionary path throughout the Cenozoic. As Olivier Rieppel has recently put it, “As living fossils are characterized by the retention of some strikingly ‘primitive’ or ancient features, these are nevertheless mixed with specializations of their own” (Rieppel 2023, 6).⁶ Nonetheless, returning to Werth and Shear, the promise: “No doubt living fossils offer, as Darwin anticipated, an unprecedented opportunity to study—with fascinating, attention-grabbing narratives—major questions concerning extinction, competition, and rates of evolutionary change, both morphological and genetic” (Werth and Shear 2014, 4).

So, as Lidgard and Love argue (2018), various perspectives associated with living fossils—apparent morphological stasis (or at least slow change), survival in refugia, and so on—are all questions deserving evolutionary focus. Love and Lidgard suggest that the purpose of living fossil concepts is to underwrite research programs targeting these questions. But how might living fossils be relevant for understanding lost worlds? Although they don’t provide “unsullied” knowledge, can they help at all? Let’s examine a case.

Tuatara (*Sphenodon* on formal occasions) is a lineage of lizard-like organisms that can be found on various islands on offshore Aotearoa (New Zealand). To my endless delight as an undergraduate student, they could also be found in a glassed enclosure near a walkover at Victoria University of Wellington.⁷ They are named after their distinctive spines (“tuatara” approximately translates from Māori into “peaks on the back”). The lizardiness of tuatara is superficial. First, the basic lizard form appears to be basal for amniotes: the earliest amniote fossil basically looks like a lizard, so it is

6 Jack Ashby argues against calling Australia’s menagerie of mammals living fossils, connecting this practice with colonial attitudes and with conservation issues (Ashby 2021, 2022).

7 I see that in 2017, they were moved into the University’s science building. I hope they are happy there.

no surprise that some amniotes have lizard-like morphology. Second, there are major differences between tuatara and true lizards. Tuatara prefer cold weather, are nocturnal, lack external ears, possess a third eye (which is visible when they're young, but covered over once adults), and have extraordinarily extended lives compared to their lizards distant cousins, averaging sixty years but sometimes hitting a century (Crook 1975). They have further unique traits: "*Sphenodon* is unique among modern tetrapods in having thoracic ribs split into jointed tubular proximal and flaring distal parts. Additionally, it is displaying gastral ribs, a feature shared in living tetrapods with crocodylians only" (Klein and Scheyer 2017, 3).

Tuatara are the last survivors of a long-lived and diverse taxonomic group, the Lepidosauria. Lepidosauria is the sister clade to the Squamates, an extremely populous group that includes lizards and snakes (over 9,000 species). In the deep past, the Lepidosauria included many tuatara-like critters as well as larger land-based carnivores and *Pleurosauros*, 1.5-meter-long aquatic predators hailing from the Late Jurassic (Dupret 2004); hence their inclusion in Prof. Ichthy's audience.

On the face of it, the tuatara represents an incredible story of evolutionary survival: a phylogenetic, taxonomic, and morphological outlier that somehow outlasted all of its relatives. If it were to disappear, then many denizens would become lost. As such, an account of lost worlds might give us some way of telling us under what conditions we should value living fossils, and what their epistemic use might be.

Recently, philosophers have presented defences of the notion of a living fossil. As we saw above, Love and Lidgard present a defence appealing to pursuit-worthiness (Love and Lidgard, 2018). In broad agreement with Werth and Sheer, they argue that "living fossils" concepts structure various investigations: into rates of molecular evolution, morphological stasis, and so on. So, the term is useful insofar as it underwrites and partially links those investigations.

Derek Turner (2019b) has provided a somewhat more realist conception and defence, based on what he calls the *phylogenetic living fossil* concept. This consists of three criteria, which he argues "really do sometimes co-occur in ways that are theoretically and practically interesting": (1) morphological stability, (2) rarity, and (3) high phylogenetic biodiversity value (Turner 2019b, 23). He is particularly interested in emphasizing the conservation value of living fossils; hence, the emphasis on biodiversity value (Turner and Han 2023). Turner defends the tuatara as a paradigm phylogenetic living fossil:

“The group is species poor, and exhibits prehistorically deep morphological stability. The group is also evolutionarily distinct.” (Turner 2019b, 23).

I’m not interested here in negotiating between Love and Lidgard and Turner (Lidgard and Love 2021). Nor am I really interested in defending the idea of living fossils against the various objections scientists have pushed, although I agree with Watkins (2021) that declaring a putative living fossil not so based on it lacking some particular feature, such as having a fast rate of molecular evolution, is to miss the point.

Instead, I want to ask whether phylogenetic living fossils like tuatara might provide some special insights into our knowledge of the past, something which Turner briefly mentions: “Horseshoe crabs and other taxa that exhibit morphological stability over the deep past might also have epistemic value to scientists interested in reconstructing the deep past” (Turner 2019b, 23). Turner gives the example of horseshoe crabs’ unusual blood, grounding a basis for inferring similar blood in their ancestors. In other words, classic grafting: inferring a past property based on shared ancestry. Aja Watkins (2021) focuses more specifically on the kinds of inferences living fossils can ground, particularly about the relationship between phenotypic similarity and evolutionary stasis, arguing in particular that it isn’t molecular rates of evolution that living fossils primarily evidence, but phenotypic and developmental change and stability.

I think there is more to explore concerning the epistemic value of living fossils. To see this, let’s turn our attention to one of the tuatara’s extinct relatives.

Palaeopleurosaurus could be understood as a transitional form between the Late Jurassic marine *Pleurosaurus* and its terrestrial ancestors from the Triassic. Like *Pleurosaurus*, it is a long, thin—somewhat snake-like—and potentially marine reptile. There is dispute as to whether they were obligate aquatic animals or partially lived on land. Klein and Scheyer (2017) examine bone histology, that is, bone microstructure, to figure out the critter’s habitat. The underlying idea is that in most secondarily aquatic vertebrates, particular properties of bones are adapted to aquatic conditions, for instance, “Increase in bone compactness and/or deposition of secondary lamellar tissue within bone cavities during intense remodelling” (Klein and Scheyer 2017, 4). As such, by examining the histology of various secondarily-aquatic vertebrates, properties associated with aquatic adaptation can be identified. So, they characterize their target as a secondarily aquatic vertebrate, thus enabling comparison with, for instance, cetaceans (Gutarra and Rahman 2022). Because

vertebrates that have returned to the ocean did so from similar morphological and histological starting points, it is thought likely that they might evolve parallel traits in light of those changes.

In addition to characterizing their target as a secondarily aquatic vertebrate, Klein and Scheyer's analysis is rooted in a further comparison: between the bone histology of *Palaeopleurosaurus* and that of its one remaining relative, the terrestrial tuatara. It is worth quoting and unpacking their conclusion:

Aquatic *Palaeopleurosaurus* shows osteosclerosis (i.e. bone mass increase) in its gastralia, some osteosclerosis in its rib, and none in its femur. However, bone compactness in the (proximal portion of the) rib of the terrestrial *Sphenodon* is only slightly lower. There is no difference in femoral micro-anatomy between *Palaeopleurosaurus* and *Sphenodon* (Castanet et al. 1988), so that the former resembles the terrestrial condition of the latter. Thus, the microstructure of the studied bones of *Palaeopleurosaurus* does not reflect a clear aquatic lifestyle. (Klein and Scheyer 2017, 7)

In the above quote, Klein and Scheyer note similarities between the terrestrial tuatara and putatively aquatic *Palaeopleurosaurus*. Although in the latter there is some increase in bone mass and compactness, it is not so much more than seen in tuatara. And the femurs are basically the same. The idea is that if *Palaeopleurosaurus* was an obligate aquatic animal, its bones would be far more suited to that habitat, that is, they would diverge from tuatara in those regards. They go on: "Finally, we hypothesise that the life history of the aquatic *Palaeopleurosaurus* differs from that of terrestrial *Sphenodon* due to an irregular spacing of growth marks, the presence of non-annual rest lines and a possibly reduced lifespan in the extinct *Palaeopleurosaurus*, although this hypothesis needs further testing with more individuals" (Klein and Scheyer 2017, 7).

You can read an organism's growth, speed, and age by looking at a cross-section of bone, similarly to how you might read tree rings. It looks like *Palaeopleurosaurus* did not share tuatara's long lifespan. Klein and Scheyer further note that a proposed function for tuatara's strange ribs, that "the ventral portion of the thoracic rib with its partially overlapping, spade-shaped anterior and posterior expansions was supposed to form, together with the

gastralia, a protective ‘armor’ for the viscera” (Klein and Scheyer 2017, 6), is undermined by the comparison with *Palaeopleurosaurus*.

So, although Klein and Scheyer deny that we should think of tuatara as living fossils (they call the term “erroneous”), their being a phylogenetic living fossil in Turner’s sense is crucial for their analysis. Their study drew on three perspectives. First, taking *Palaeopleurosaurus* as a secondarily aquatic vertebrate, which enabled some general ideas about what aquatically adapted bones should be like. Second, careful examinations of bone histology in *Palaeopleurosaurus* fossils. Third, taking their target as Rhynchocephalia, that is, a member of the same ancient clade that tuatara hail from, enabled them to use tuatara as a relevant point of comparison. Tuatara’s being terrestrial is here a boon, as it provides a point of contrast for exploring just how aquatically adapted *Palaeopleurosaurus* is. Note how this goes beyond Turner’s suggestion of projecting tuatara-properties into the past. Instead, the tuatara acts as a crucial comparator for studying *Palaeopleurosaurus*: the two are grafted into the same phylogenetic story and thus become mutually relevant for understanding one another.

Also, note the coupled nature of the relationship between fossil analysis and comparative work with tuatara: the two are mutually informing. And indeed, we even learn something more about the tuatara themselves. The investigation demonstrates how palaeontologists bring together now and then in mutually-informing ways. It further demonstrates how flexibly adopting multiple perspectives and navigating between them can allow progress on long-dead, largely lost animals.

So, where does this leave the notion of “living fossil”? Scientists are right that we shouldn’t literally treat such lineages as still-living fossil denizens. “Living fossils” retain some properties, yes, but evolution ensures they differ in many ways from their ancestors. But the way they differ from their forebearers can provide crucial information for uncovering those forebearers. Things usually identified as living fossils are certainly not special or unique in this regard: sharks are in no way rare enough to count as living fossils, but nonetheless, their morphological stability certainly underwrites much work in understanding extinct sharks.

Consideration of living fossils like tuatara demonstrates at least four things about the science of lost worlds. First, and I hope you’re getting bored of me making points like this, whether or not some lineage is a living fossil depends on how you characterize it: what traits are you interested in? Second, living fossils bring the importance of loss into stark relief. If tuatara didn’t

exist, we wouldn't know some of the things we know about *Palaeopleurosaurus*. Third, it extends Turner, Watkins, and Love and Lidgard's points about the epistemic value of living fossils: they stand as examples of survival against loss with the potential to tell us a lot about the past. And fourth, their epistemic value lies partly in their capacity to act as comparators—members of the same lineage—as targets in the past.

