



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>

Possible Worlds

The applause closing Prof. Ichthy's lecture ranges from polite (various academics) to appreciative (closer colleagues and grad students) to almost inappropriately rambunctious (Ichthy's family, who whoop and whistle). Ichthy beckons their students onto the stage to share the applause, which they do somewhat bashfully. As the applause dies down, a second expectant hush descends, and Ichthy indicates they'd be happy to take questions.

Audience members glance about through the hush, waiting for someone to break the silence. Finally, a younger plesiosaur breaks the awkwardness by raising a flipper and, at Ichthy's nod, bends their serpentine neck upward to speak.

"Uh, so, you've said a lot about the things we know about Broadfoot, that it was, like, a mammal and stuff. And you said some stuff about these big things that happened in prehistory, like, how the age of the mammals ended. But have we also learned anything about how things could be? Like, did these events have to happen? Could the mammals have survived? If they did survive, what would that be like? Would mammals have developed reptile-like intelligence—could they have had professors and universities and sports, and the stuff that, like, we have? I know it sounds kind of silly . . ."

The questioner trails off in some embarrassment before Ichthy comes to a magnanimous rescue: "Could all this knowledge of the actual past also be knowledge of what is possible? That is, about how things could be? An excellent question!"



ILLUSTRATION 6.1 Prof. Ichthy's Audience

Historical scientists generate knowledge despite erasure and loss. How should we characterize that knowledge? Should we call it knowledge, or understanding, or know-how, or know-that? Further, what is this knowledge *of*? Do they know facts about states of affairs in the past, or *merely* what was possible, or only some consistent set of things going on in their minds? What kind of epistemic good or goods are we talking about, and what kinds of things do these goods pertain to?

Heady, deep, philosophical questions indeed.

My immediate answer, you'll likely not be surprised to hear, is pluralistic: given the diversity of evidence historical scientists have at their disposal, the complexity of the world they are trying to understand, and the strategically perspectival, iterative tactics they employ, it would be surprising if there was a monistic epistemic story to be had. Their knowledge comes in many forms and is of many things. While this pluralism is plausible, it doesn't necessarily make for an enlightening discussion. Instead, I want to focus on a particular aspect of the knowledge historical scientists generate: the *modal* element. Above, I made telltale use of the term "merely" in relation to knowledge of what is "possible," even italicizing it for emphasis. In this chapter, I'll argue we should dispense with the diminutive "merely." A major—sometimes the main—dividend of historical science is knowledge of possibility, and it is my aim in this chapter to defend this claim and the legitimacy of modal knowledge. It might be useful to have a slogan:

The science of lost worlds is a science of objective possibility.

What does this amount to? Although sciences of the deep past tell us about the actual deep past—no doubt—they also tell us about possibility in rich, complex ways. Moreover, this modal knowledge is not merely of "epistemic" possibility, that is, what could have happened in the past given what we know. It is also of "objective" possibility. Which is to say, when the science succeeds, it provides true modal information, that is, beyond what actually happened, and concerns what could have, or might have, or must have, or would have happened had something else occurred (Wirling and Grüne-Yanoff 2024). And, as we'll see, this modal knowledge is complex and structured. It concerns not so much what is possible *tout court*, but rather concerns modal spaces consisting of sets of complexly interacting conditionals.

I'll further argue that the best way to understand the justification of this modal knowledge is not, as others have suggested, through the workings of

scientific imagination, but instead through an “artifactualist” account of perspectival tools. In brief, historical scientists construct and examine various objects—perspectival tools—and demonstrate their having various dispositions, that is, sets of modal properties. These dispositions are projected onto past targets as well as onto more general tendencies, patterns, and regularities. As we’ll see in the next chapter, the scientific imagination is expanded through the investigation of lost worlds. But I don’t think imagination, as understood by most philosophers, plays any special role in underwriting or justifying modal knowledge. There’s no special pleading required to understand how scientific representational strategies generate modal, as opposed to non-modal, knowledge.

My strategy is to begin with the old philosophical chestnut of “how possibly” explanations. I’ll carry out some disambiguation and identify a set of “how-possibles” that fit the phenomena at hand. I’ll then argue for a minimally realist account of this objective modal knowledge. What I don’t mean by this is that scientists get it right all the time; rather, I have in mind the humbler idea that the reference of their claims is often modal structures, not actuality. This will involve a commitment to artifactualism.

With that in place, I’ll turn to a case study: this time, astrobiology. Astrobiology is as much a science of the deep past as palaeontology. I’ll argue that construing astrobiology as being in the business of identifying extra-terrestrial life is, surprisingly, missing a large part of the story. As we’ll see, astrobiology (or at least the corner I’ll explore) provides dividends in the form of understanding the relationships and dependencies—the rich modal structures—holding between various biotic and abiotic processes and potential atmospheres. That is, astrobiologists map out how various planetary compositions provide the conditions of possibility of differing atmospheric compositions, given various other biotic and abiotic factors.

Much of the philosophical discussions I’ll be drawing on herein concern modelling, especially the (surprisingly small) literature on the modality of models. I’m not sure whether anything much hangs on any of the practices I’ll cover being called “modelling” or not. The last twenty years or so have seen an enormous amount of work in the philosophy of modelling, and I don’t really see much unity in how philosophers of science conceive of models. Joe Wilson (2024) has recently pointed to various similarities between how philosophers have characterized model-based and historical science, using this to put pressure on attempts to characterize sciences in terms of “prototypical” methods. In this context, attempting to delineate between what’s a model and

what's not a model seems at best a little boring and at worst obscurantist. Presumably, a table (of contents, or of data) isn't a model—though perhaps it sometimes counts as a “data model”—these are nonetheless perspectival tools that scientists make use of.

So, suffice to say, I'll be talking about “perspectival tools” in the sense I introduced in chapter 3: these isolate and instantiate various properties to be explored by providing a space for comparison (as in, say, phylogenetics) or direct intervention (as in experimental taphonomy), or via inference (as in spatial point analysis). We've seen a wide variety of perspectival tools, from Woodruff and Varricchio's (2011) scale model of a dinosaur burrow to various geometric representations of the growth of ancient organisms to phylogenetic and morphological trees, tables, and morphospaces. It is tempting to think of these practices as situating a phenomenon or a set of phenomena within a “logical space” (Elgin 2017, 155); that is, the perspectival tools delineate a set of constraints within which the target is situated. This is about right, but as we'll see, I'll deny we should necessarily think of these spaces in “logical,” or abstract terms: perspectival tools are artifacts, that is, physical objects with their own dispositions. These dispositions themselves set a space of possibility.

The science of lost worlds is a science of objective possibility, then, in two senses. First, it provides modal knowledge of more than epistemic possibility; second, this knowledge is “objectual,” that is, grounded in the dispositions of particular objects.

A quick note before diving in. A discussion of modality in palaeoscientific contexts might immediately call to mind Stephen Jay Gould's famous “life's tape” thought experiment (Gould 1989). To illustrate his view on life's contingency, Gould asks us to imagine “rewinding” life to some earlier, presumably crucial, stage. He asks, were the tape then allowed to run forward, whether we'd see the same history unfold: for many events in life's history, Gould bets that we wouldn't. The thought experiment has been used to frame everything from philosophical analysis of contingency to experimental laboratory work. Here, however, I'll not be discussing the modal nature of Gould's thought experiments. This would distract us from the significantly more grounded modal knowledge produced by the cases we've been examining. Regardless of what we might think about replaying metaphorical tapes (and, indeed, we might think imagination plays a role in these thought experiments), the bona fide modal knowledge apparent throughout the more workaday investigations we've been concerned with, from Avalonian ecology to dinosaur burrows and

pain behaviour, to the phylogeny and development of *Dickinsonia*, is deserving of its own philosophical accounting.

1. How Possibly: At Least Three Ways

Let's frame this chapter's discussion by distinguishing between different kinds of how-possible explanations. The term "how-possibly explanation" is traceable to William Dray's work in the philosophy of history (1957), where he argued that historical explanation often proceeds according to a distinctive form (one differing from scientific explanation, as understood by Hempel). Dray's notion of "how-possibly" explanation was picked up in various areas of philosophy, particularly in the philosophy of biology's discussions of adaptationism (Brandon 1990; Forber 2010), and along the way, it shifted somewhat from its original meaning. Alisa Bokulich (2014) has provided a nice nutshell of this history, and, along with Thomas Reydon (2023), provides much-needed clarity that I'll draw on. I'm interested in disambiguating forms of "how-possibly" to zero in on various ways that modality might feature in epistemic practices like explanation. As we'll see, many discussions of how-possibly take modal knowledge to be *in service of the actual*, but I'll argue this strategy misses some of what is crucial about understanding lost worlds. Instead, we should (at least sometimes) take the "how-possibly" of historical science to refer to objective possibility, and to be targeted at that objective possibility.

I'll call our first form of "how-possible" *mere-how-possibly*. These are, I think, what philosophers of science often have in mind when they appeal to "how-possibly explanations." To show something is possible is to demonstrate that there is at least one way that it could happen. Typically, mere-how-possibly explanations take a set of constraints and demonstrate that some outcome is consistent with those constraints. Consider challenges to neo-Darwinian adaptationism from appeals to "irreducible complexity" (remember those?). It is denied that some complex trait could be generated via an incremental evolutionary process of selection (therefore, through some argumentative magic, said complex trait must have been designed by . . . someone . . .). Typically, the response is to provide some possible evolutionary routes through which said trait could have evolved incrementally, thus defusing the challenge. Here, the adaptationist need not show how the trait in fact evolved, but rather show how it possibly could (Weber and Depew 2004).

In some scientific disputes, demonstrating that a hypothesis is merely possible is crucial. For instance, "Kon-Tiki experiments" aim to show that

some hypothesis is possible via some kind of physical demonstration (Novick et al. 2020). Thor Heyerdahl drift-voyaged from South America to the Pacific to show that a Pacific settlement from that direction was possible with those technologies, and further Kon-Tiki experiments using Polynesian craft and navigation techniques played a decisive role in showing that non-drift voyaging was possible in the other direction. These demonstrations of mere possibility removed major objections to the relevant hypotheses and were crucial to the eventual acceptance of a non-drift, westward model of Polynesian settlement. Further, because mere-how-possibly explanations concretize at least one possibility within a set of constraints, they can play crucial roles in exploring those constraints and various other forms of scaffolding (see Walsh's 2019 discussion of scaffolding). Mere-how-possibly will return in chapter 7's discussion of palaeoart.

Another form of how-possibly captures Dray's original meaning, so let's call it *Dray-how-possibly*. Here, instead of providing some general modal information—whether some outcome is consistent with some set of constraints—how-possibly explanations, as originally envisioned, involve providing *actual* information, which demonstrates that some *actual* outcome is possible. As Dray put it: “The explanation consists in showing that, in spite of all appearances to the contrary, the event was not impossible after all” (Dray 1968, 391). Explanations of magic tricks often take this form. The magician was able to pull a rabbit out of their empty hat—the trick was possible—because of a combination of misdirection and a secret compartment in the hat that gave it the appearance of being empty. We're given some actual information (the secret compartment), which explains how what was actual was possible.

Consider explanations of global biogeographical dispersal prior to the development of tectonic plate theory. How were various animals able to cross the oceans between, for instance, South America and Africa? In a mere-how-possibly context, the detractor might be happy with an appeal to land-bridges, independent of whether they actually occurred. For a Dray-how-possible explanation, we point to the actual spread of tectonic plates to explain how this could have occurred.

Demonstrations of possibility can play both *mere* and *Dray* roles. If one was informed that Polynesia was settled from Southeast Asia, and expressed bafflement as to how this could be achieved, an explanation of the highly sophisticated navigational techniques and technologies that Polynesian ancestors possessed might remove that bafflement, thus forming

a Dray-how-possibly explanation. Instead, one might simply be interested in whether it is possible to get from Southeast Asia to Polynesia via some kind of craft, and would accept *one possible story* (regardless of whether it is the actual story). As Reydon summarizes: “Dray’s how-possibly explanations tell the actual historical story to show how an improbable event was possible, whereas Brandon’s and Resnik’s how-possibly explanations only speculate about history” (Reydon 2023, 166).

That is, because mere-how-possible explanations are in the business of demonstrating that some outcome was possible, as opposed to showing that some actual event made an outcome possible, insofar as they’re interested in the actual, it is in a merely “speculative” fashion. Bokulich (2014) demonstrates the complex, context-dependent relationship between different how-possible and how-actual explanations.

For my purposes, there is a further similarity between these how-possible explanations. They are typically about, I think, epistemic rather than objective possibility. In these cases, *the possible is in service to the actual*. That is, we care about what is possible because we ultimately care about what is actual: we’re interested in further constraining the set of worlds that could be the actual world. This is necessarily the case for Dray-how-possibly, and typically the case for mere-how-possibly. To get the flavour of this, it’s helpful to work through an example from Massimi. “As new pieces of experimental evidence $e_1 \dots e_3$ are brought in, they refine the how-possible inferences by excluding more and more of the live objective possibilities $x_1, \dots x_n$. This is because new evidence gradually and increasingly constrains what is physically conceivable” (Massimi 2019, 877).

Roughly, Massimi is describing how the scope of what is “physically conceivable” (see below) shrinks as a clearer picture of the actual emerges through incoming evidence. Despite her use of the term “objective,” this is epistemic possibility: some proposition *might* be actual insofar as it accords with the other things we know. In principle (and perhaps with some wiggle room), as knowledge approaches completeness, physical conceivability and actuality approach overlap (we’ll return to physical conceivability in the next section). This is what I mean by saying the “possible is in service of the actual.” We care about mere-how-possibly explanations because we want to know whether some actual outcome could have occurred given such and such a surprising circumstance; we care about Dray-how-possibly explanations because we want to know how some apparently improbable event occurred. These claims are all of the form “such and such happenings are all consistent

with what we know.” For Dray-how-possibly, we say “this event was possible after all because of how the past was”; for mere-how-possibly, we say “it could have happened like this, for all we know.” Objective modality claims are more of the form “if such and such were to occur, then the following would happen,” that is, they aim for true claims about possibility itself. This is the kind of thing I’ll have in mind when I discuss the dispositions of artifacts below.¹

We needn’t think of possibility as playing a mere in-service role. As Daniel Swaim has put it, epistemological accounts of modality do “not fit well with a great deal of scientific practice, wherein the point is to lay bare (to the extent possible) what the structure of physical possibilities is really like, and subsequently, why things turned out to be the way they are” (Swaim 2021, 264). Thomas Brandstetter characterizes models in geology as “how-possibly experiments,” aiming to test hypotheses about what actually happened by determining whether they could have, possibility in service of the actual. But he goes further to claim that some models were “miniature theatres for dramatizing the action of forces on matter over time. In this sense, the model experiments were time machines: they made it possible to visualize geological time itself” (Brandstetter 2011, 135). Beyond being in service of the actual, then, they also provide routes to examine geological processes—and potentially objective modal knowledge.

Let’s consider some discussions I take to indicate a third set of how-possibly: explanations that attempt to map out true objective possibilities, sometimes explaining via situating a target within that modal space, thus providing information about how and why that target could, would, and does occur. This involves appeal to objective rather than epistemic possibility, as—when we have a successful explanation—the space of the possible doesn’t shrink as more information is included. As we’ll see, often but not always, these modal spaces are generated by perspectival tools.

A view about how models provide insight into possibility is associated with Batterman and Rice (Batterman and Rice 2014; Rice 2018). The basic idea is that various models and targets are part of the same “universality class.” I interpret “universality class” in a particular way: to be in the same

1 The term “disposition” looms large both in analytic metaphysics and scientific metaphysics. Here, I’m using it as a term to pick out the various modal properties, often conditionals, that constitute the knowledge generated from perspectival tools. I don’t take myself to be committed to any particular claim about the nature of dispositions generally, although (1) I do take dispositions to be necessarily modal, and (2) I don’t think all modal claims are dispositional.

universality class is to share dispositional properties. For instance, the Avalonian ecologies that concern Mitchell et al. are, if their analyses are right, in the same universality class as contemporary ecosystems insofar as when they share the same sets of properties—small populations, fast generational turnover, lack of nutrient constraint—they're disposed to exhibit neutral behaviour. The same might be said of the highly abstract models ecologists use to explain these. There are, I think, similarities between the argument I will provide in the next section and aspects of Batterman and Rice's solution. However, they are focused on how minimal "toy" models capture some fairly universal aspects of real-world targets, while I'm interested in a broader array of perspectival tools, as well as their combinatorial features. The important thing about their account, however, is that possibility is not in service of the actual: in addition to learning some dispositional properties about our target (say, the rangeomorph-dominated ecosystems of the Avalonian), we *also* learn modal information about the universality class (say, the properties of neutrally-driven ecosystems). It is in virtue of two systems sharing membership in the same universality class that the same modal claims apply to them.

Focused on the historical sciences, Meghan Page has argued that one benefit of studying the deep past is probing and testing the fragilities exhibited by regularities established under contemporary conditions. As she puts it:

Historical investigations often play an important and ineliminable role in advancing our knowledge of causal structure and stable regularities. When confronted with complex geological processes that occur over large regions of space and time, historical science acts as an experimental context to test the stability of currently observed regularities. As a result, historical investigations enhance our understanding of the complex causal webs that produce natural processes and help predict how such processes might evolve or shift in the future. (Page 2021, 462)

Sometimes palaeontologists are in the business of using contemporarily established regularities to make inferences about the past, most blatantly in trace-based reasoning. But sometimes they are in the business of probing the stability of contemporarily established regularities by applying them to the past. Consider again the question Mitchell et al. put to the surprising

neutrally driven spatial ecologies of the Avalonian. They ask: do these show that our usual theories about ecosystems simply *fail* under those conditions? By anchoring their analyses across a set of other ecosystems and models, Mitchell et al. argued that when properly understood, our ecological models can, in fact, predict neutrality. This represented an explanation of Avalonian neutrality (it was due to short generation time, low dispersal, and frequent disruptions). But further, it deepened our modal understanding of the factors shaping niche and neutrally structured ecosystems. We now know better under what circumstances ecosystems might exhibit those structures. Often, the pluripotency of perspectival tools is fundamentally modal.

Mitchell et al. thus investigate what, in 2018 (chapter 7), I called *fragile systems*. Here, the target is not some actual event, or even a set of events, but rather the conditions under which certain dynamics hold. In the more extreme cases, such as Mitchell et al.'s, we could characterize these as testing for certain kinds of loss, for instance, of the conditions for the applicability of certain kinds of dynamics.

Expanding from this kind of thought, we can characterize much work across the historical sciences as grappling with various fragile systems; this is particularly true of artificing strategies. Woodruff and Varricchio (2011) empirically investigate a set of dependencies between burrow infill and bone position; Laflamme and Narbonne (2008) model dependencies between rangeomorph morphology and optimization in osmotrophy; Cuthill and Conway Morris (2014) examine osmotrophy in light of fractal morphologies. Each of these artifacts instantiates a possibility space, and examining them generates knowledge of that space (Currie 2020, 2022). Further, the results of artificing are drawn on, integrated, and contextualized with historical reconstruction to provide complex trajectories, such as Butterfield's explanation of the long road from Ediacaran-style oxygenation to that which emerged through and beyond the Cambrian.²

Daniel Swaim's account of *narrative possibility* makes a similar point (Swaim 2021). Focusing on narrative explanations (roughly, explanations that take the form of a token sequence of causal events, culminating in the *explanandum*), he develops an "ecological" approach. This involves, first, central subjects—the token historical individual the story is wrapped around—and their capacities; second, these capacities take on various *affordances* given

² In Currie (2021b and 2024a), I develop a notion of "profiles" that further fleshes out this kind of notion.

the environments into which those historical individuals are embedded. In Butterfield's explanation of the diurnal pump, he pointed to the oxygen-structuring capacities of simple eukaryotes and showed how their being embedded in pre-Cambrian environments afforded more concentrated oxygen and thus a scaffold for the emergence of larger and more complex eukaryotes.

In these circumstances, scientists of the deep past are not only concerned with token denizens in the past—particular events and processes—they are concerned with the dynamics, dispositions, and fragilities of various system-types, whether they are token-historical, historical, or ahistorical. In these circumstances, as our knowledge grows, possibility does not shrink to actuality, although it might become more complex and structured.

These increases in modal complexity and structure were what I had in mind earlier when I appealed to “sets of complexly interacting conditionals.” Consider Mitchell et al.'s work. They shift from modelling Ediacaran ecosystems via a dichotomy between neutrally and niche-structured spatial locations, to introducing various conditionals regarding small populations, dispersal distance, and more. These conditionals interact: if an ecosystem has small populations, it will tend toward neutrality; if an ecosystem is often disrupted, it will tend toward neutrality. These conditionals work together in explaining their result. So, in Mitchell et al.'s work, modal complexity is added as various conditionals concerning the factors associated with neutrality are brought online. The product of the research is a modal space that tracks the conditions under which neutrality and niche differentiation are expected, and a how-actually explanation³ when the Avalonian past world is situated within that space.

I think these extremely common circumstances across investigations of past worlds underwrite a *prima facie* motivation for taking this scientific work to be uncovering objective possibilities.

Fragile systems are explored through the use of perspectival tools. Not all perspectival tools aim to capture objective possibility directly. As we saw in chapter 4, cladistics and other grafting techniques make use of perspectival tools, but the aimed outcome (at least sometimes) is actual ancestral relationships. When we have several possible clades, this should be understood in terms of epistemic possibility. However, these are often used as the basis of developing modal relationships. Knowing ancestry, for instance, is necessary

3 The explanation, as I've characterized it, is potentially a Dray-how-possibly explanation.

for identifying patterns of convergence across taxa, thus underwriting claims about the conditions under which various adaptations might arise (Griffiths 1996; Currie 2013). Recall that perspectival tools are typically pluripotent: applying not only to single instances but drawing out similarities between instances and modelling them more generally. This last function is plausibly understood in terms of objective modal properties.

So, no doubt epistemic conceptions of possibility matter to historical investigation—indeed, the exploration of epistemic possibility is central to chapter 7’s discussion of the scientific imagination—but they don’t exhaust what is modally relevant for the investigation of lost worlds. Historical science is concerned not simply with the actual and with epistemic possibility, but with modality and objective possibility, and, following accounts like Swaim’s, we can understand that possibility as “baked into the world’s structure” (Swaim 2021, 265). But which epistemology should follow this ontology? That is, how do historical scientists gain knowledge of objective possibility?

2. Imagination Bad—Artifactualism Good

“But how might we know about these possibilities? You might start by considering our imaginations.” Ichthy taps the top of their head, their spectacles bouncing a little comically. “I can entertain the possibility that the mammals survived, and that they evolved taxa with lives much like ours. I can imagine, surreal as it may be, some sentient mammal giving a lecture much like this one. And maybe my being able to imagine this gives us some reason to think that sentient, intelligent mammals are, indeed, possible.”

Ichthy pauses dramatically. In the audience, one graduate student nudges another with a slight smirk, indicating that they’ve heard what’s coming before.

Ichthy puffs their chest out. “But I am a scientist, and just picturing or considering things in the cage of my mind should carry very little weight indeed. Instead,” Ichthy gestures to the toy Broadfoot still held by one of the students, “we build things. And the things we build tell us about what could be.”

I've thus far laid out a few options regarding the kinds of modal knowledge sciences of the deep past might generate and argued that, at least sometimes, objective possibility better captures the dividends. This leaves a set of crucial epistemic questions open, especially what Wilring and Grüne-Yanoff call "the epistemic question" concerning modal models. They ask: "What must models or modelling practices be like in order to provide good reasons for modal claims (*under what conditions*) and why does fulfilling those criteria make models epistemically valuable with respect to modal truth (*in virtue of what?*)? [italics in original]" (Wilring and Grüne-Yanoff 2021, 4).

Multiple questions, then. In this section, I'll make some steps toward answering them. I'll start with a set of views that appeal in some sense to the cognitive, imaginative capacities of scientists, and argue that these don't capture what is going on in the cases we've discussed. I'll then turn to, I think, a better kind of view—artifactualism—and argue that it does.

The argument in this section is narrow in several senses. First, if you'll permit me instrumental use of a much-maligned distinction, I'm stepping into the *context of justification* regarding the role of imagination. That is to say, I'm not here interested in the role imagination plays in the generation and development of ideas (discovery and pursuit) but in claims that point to imagination playing a special role in justifying scientific modal claims. Second, I'm interested in the justification of modal knowledge derived from perspectival tools. As we'll see, some of the prominent views on the table accord imagination a special justificatory role for that modal knowledge, and I deny this. So, imagination might well play justificatory roles throughout science, but it doesn't play a special one in this case. Third, I'll be taking imagination in an "imagistic" way, that is, as a mental faculty that involves sensory-like representations. As we'll see in the next chapter, there are other ways we can understand imagination that are more useful for our purposes. Indeed, in next chapter, I'll argue for an artifactualist account of scientific imagination generally.

I'd advise not getting too caught up in what I mean by "knowledge" here—I'm using it loosely, in a way that could include close by epistemic notions, such as understanding or grasping. I also don't think my usage needs to distinguish between propositional and non-propositional ("know-how" or ability-based) conceptions of knowledge. Although I'll speak of dispositions and express counterfactuals propositionally, I suspect this is a matter of presentation.

2.1 Imagination Bad

The idea that models are valuable because they provide modal knowledge is a move often made by philosophers trying to make sense of highly abstract models that bear little or no resemblance to their target systems—sometimes called “minimal” or “toy” models (Reutlinger et al. 2018; Gelfert 2019). Given that such models surely can’t give us true information about actual targets per se, perhaps they can give us true information about what is possible. Such models idealize in order to isolate a system’s essential dynamics. As Michael Weisberg has put it, a minimal model “contains only those factors that make a difference to the occurrence and essential character of the phenomenon in question” (Weisberg 2007b, 642). The epistemic questions ask, then, what reason do we have to think that such modal information is justified? And, if models provide knowledge of objective possibility, in virtue of what do they do this?

Some answers to these questions appeal to scientific imagination. Following Wilring and Grüne-Yanoff’s characterizations, there are at least two versions of the proposal. One takes models to describe fictional worlds and rests their credibility on the imagination-driven credibility of fictions. Meanwhile, another set is interested in the “conceivability” of worlds, that is, the capacity of relevant experts to imagine them. What unifies these, more or less, is that the justifying circumstances under which the models do their modal work are—so far as I can tell—inside scientists’ heads. For Michaela Massimi, for instance, modal claims derived from scientific tools are justified when the relevant scientists can *physically conceive* of those claims. That is, they can imagine it under the constraints of their background knowledge and the laws of nature they know. Till Grüne-Yanoff, by contrast, takes the fiction described by models to be judged as *credible* in ways similar to how we judge the credibility of fictions: although fictions depart from the actual world, we still seem able to judge them as more or less credible. “A model is credible—and hence a good guide to justified possibility claims—when the model world is internally coherent and a competent user of the model would judge the development in the model world to be intuitively plausible, conditional on the model set up” (Wilring and Grüne-Yanoff 2021, 6).

So, in addition to requiring coherence, such views take a scientist’s *judgement*, a judgement underwritten by some imaginative act, to be necessary for the justification of the modal knowledge generated. Reflecting on the modal claims we’ve seen historical scientists make, I find such appeals implausible.

Let's run imagination-views through a simple case. In chapter 3, one of our examples of artificing involved Woodruff and Varricchio constructing a burrow, filling it with rabbit bones, and probing the relationship between bone deposition, initial bone position, different kinds of burrow inflow, and different compositions of flow. They primarily did this to test hypotheses concerning a token past world: whether the critters found in the fossilized Jurassic burrow should be interpreted as starting within the burrow or having been deposited from the outside. They also intended the model results to have more general applications, applying to fossil burrows across a wide range of actual and non-actual instances. As I argued in chapter 3, their experiment is pluripotent. In addition to testing hypotheses concerning that fossilized burrow, they are reasonably understood as aiming at objective modal knowledge: namely, of the dependencies between those variables in burrows. They thus tell us about the dispositions of skeleton deposition in burrows. And knowledge of these dispositions underwrites counterfactuals and other modal claims. Their argument depended on this being the case: we should interpret the dinosaurs as having been in the burrow at the time of death because if they had not been in the burrow, the bones would have been distributed differently.

Let's focus on the counterfactual: *if the critters had been outside the burrow, then they wouldn't have left that pattern of bones*. By *credibility*, this modal information is justified by the relevant scientists considering the relationship between the fictional world where the critters didn't start in the burrow—imagining it, using their background knowledge—and then, given the model setup and their own expertise, concluding that the counterfactual is “intuitively plausible.” By Massimi's “physical conceivability” view, similarly, we should accept the counterfactual because it is physically conceivable to the scientists in question, that is, given the relevant constraints and relevant “laws,” they can imagine this outcome.

I don't doubt that Woodruff and Varricchio did imagine and think through the results of their studies in the way described by these views, but I don't find this a convincing way of accounting for the epistemic question. Appeals to imagination explain neither the conditions under which their experiments generate modal knowledge nor what justifies it.

Scientists' imaginings are extraneous to the justification of the modal knowledge they generate. By setting up an experimental apparatus and performing various interventions, Woodruff and Varricchio empirically generate and explore a modal space. They are able to map, in the experimental system,

the conditional relationship between initial and final bone position against various differences in inflow, bone density, etc. The justification for this conditional knowledge turns on features of the experiment, their data-recording practices, and statistically relevant properties concerning how representative the experiment is. The addition of the scientists doing some imagining doesn't seem to bring any obvious justificatory weight: after all, there is a physical object that generated the data. No doubt this requires interpretation, but I don't see any particular reason to think this interpretation would require imaginative capacities, especially. If their empirical arguments hold water, they have established a set of dispositions possessed by the experimental apparatus they built.

However, generating modal knowledge pertaining to the experiment isn't the end of the justificatory story. We must also consider whether that set of dispositions can be projected onto both the Jurassic burrow that it is intended to probe and onto more general characteristics of burrow infills (perhaps concerning the "universality class" at hand). I find imagination extraneous here as well. As we saw in chapter 3, Woodruff and Varricchio carefully point out various differences and discrepancies between the object of intervention and the eventual target, and possible idiosyncrasies about their experimental apparatus that could undermine its representativeness. For instance, their experiments involve a single, complete rabbit skeleton, as opposed to the three incomplete skeletons of the target. In their discussion, they note that "[t]he single rabbit skeleton consists of 230 elements compared to the observed 122 in the Blackleaf Formation specimen. The number of total elements likely would impact the depositional process to some extent, but it is not clear a priori, whether the rabbit assemblage should be increased or decreased in number for proper modelling" (Woodruff and Varricchio 2011, 144). Although this appears somewhat inconclusive, I take the point to be that it is unknown what kinds of effects having more or fewer elements would have, but that, regardless, it is unlikely to affect the main thrust of the experiment, which is to determine the relationship between initial skeleton position and position in deposition.

It is this careful consideration of possible differences, similarities, and confounders that underwrites the justification of projecting the modal conditionals, not the capacities to conceive of the conditionals holding in those other cases, nor judging them as credible. The experiment involved isolating a set of properties thought to be possessed by the denizen, determining the

objective modal features of those properties, and then arguing that we should expect similar properties with similar modal features in the lost denizen.

Philosophers familiar with the literature on projectability and experimentation would, I hope, have found this argument somewhat obvious: of *course*, the projectability of experimental results turns on their internal and external validity (Weber 2004; Currie and Levy 2019). I agree, and this is precisely my point. I don't see any special difference between projecting modal information about an experimental setup versus projecting other kinds of information. That is, where imagination-based views make a special appeal to the role of imagination in justifying modal knowledge, I think the same basic structure holds for both modal and non-modal knowledge.

Perhaps a thought experiment will help (perhaps . . .). Let's imagine that Woodruff and Varricchio simply *cannot* conceive of the dinosaur bones ending up in the positions they were found in if they started within the burrow. A failure to conceive, in this instance, involves not finding the result "intuitively plausible." We might envision this as the scientists running a kind of mental simulation, and the patterns of bones just not behaving as they ought. Or the scientists being unable to figure out, through their own reasoning, what follows from their results. I can imagine this being very frustrating for the scientists, and I can imagine this motivating them to do further exploration. And indeed, I can imagine them complaining that they *don't understand* their results. But, crucially, would this undermine the counterfactual claim: "if the critters had been outside the burrow, then they wouldn't have left that pattern of bones"? I don't think so; after all, why should a scientist's capacity to run mental simulations make any kind of difference to *empirically demonstrated* modal knowledge? It is that empirical demonstration, I want to say, that justifies the counterfactual, not the capacity of some scientist or other to imagine it.

Following Grüne-Yanoff's definition of credibility (see above), you might respond that the scientists in the thought experiment are irrelevant because either the model isn't credible or the user is not competent. But I don't think this kind of response works. It seems to require claiming that credibility or competence turns on whether the scientist understands the system they are engaging with. Another set of scientists, reading Woodruff and Varricchio's results, would primarily judge their credibility just as we usually do: Is their evidence good? Was the experiment set up well? Has external validity been sufficiently established? Is the lab capable of this kind of work? When a scientist asks after another's credibility, they ask these kinds of questions; they

don't ask after scientists' capacity to imagine. Again, I don't mean to claim here that imagination plays no role in science, or even in justification; rather, I'm arguing that it doesn't seem to play any kind of special role in establishing the modal knowledge generated by studies like the one at hand.

You might complain that I'm taking "imagination" or "conceivability" too literally—to say that something is "conceivable" isn't to say that some actual agent needs in fact conceive it, but rather that some sufficiently idealized agent would do so (what has been called a "non-epistemic" account of conceivability; Massimi 2022, 148). It is sometimes a little ambiguous how psychologized these terms should be taken to be. By Massimi's account at least, "imagining" appears to be a practice both attributed to a particular epistemic subject but also an epistemic community, suggesting that conceivability is at least to some extent de-psychologized (depending on how happy you are with communities having minds). However, I think ultimately we should read Massimi's and other appeals to "imagination" in this context to place psychologistic capacities in a crucial position for justifying the modal deliveries of models. The constraints of background knowledge, laws, etc., work for Massimi to hold imagination to some kind of empirical account: "Scientific imagining has to respond to the tribunal of experience" (Massimi 2022, 150).

You might further complain about my choice of case: it is an experiment, not the kind of toy model that imagination-folks are concerned with! Fair enough: let's consider a much more "modely" case: Emily Mitchell and company's attempt to understand the rangeomorph-dominated ecosystems of the Avalonian Ediacaran. As we saw, ecological models of the spatial and taxonomic relationships of those fossil beds turned up a unique, neutral-dominated pattern. In response, Mitchell et al. adopted a strategy I've called *anchoring*, navigating among various natural systems toward informing a model of those systemic dynamics. They considered various properties of those Avalonian systems—their relative immaturity due to frequent disruption, small sizes, lack of resource constraint—and then compared these with contemporary systems that also exhibited these properties, demonstrating that they too exhibit the telltale signs of neutrally structured ecosystems. This involved modal claims, such as *if the Avalonian ecosystems had more resource competition, then they would be more niche-structured*, and *if the Avalonian ecosystems were more resilient, they would be more niche-structured*. The justification of these counterfactuals doesn't at all seem to me to involve the scientists imagining or conceiving of them, regardless of their expertise and how much constraint the "tribunal of experience" provides. It comes from

the specific arguments they provide and the demonstrations they perform. A thought experiment à la the one above can be rerun: even if Mitchell or her colleagues were unable themselves to imagine the counterfactuals, if they failed to physically conceive of them, or failed to judge them as credible, the justificatory work would nonetheless be carried by the dispositional properties they establish in their models, and the links they empirically draw between models and real-world targets.

A final problem: I'm not sure whether imagination-based answers to the epistemic question of modality are well-placed to capture objective modal knowledge. Recall that epistemic possibility is, basically, what is left open given the constraints of some individual or epistemic community's knowledge. As Knuuttila has put it, epistemic possibility "depends on scientific knowledge of the actual state of the world. Any scientific claim that is not ruled out by scientific knowledge is epistemically possible" (Knuuttila 2024, 65). By contrast, *objective* possibility is not simply a claim about what is consistent with our knowledge of the actual, but involves modal claims beyond actuality. The counterfactual statement *If there were more competition for resources in Avalonian systems, then they would be less neutrally structured* is epistemically possible if the statement is consistent with what scientists know. It is objectively possible if, regardless of scientific knowledge, the counterfactual is true (or sufficiently veridical, or justified, or whatever). As we saw, one way of differentiating between these is whether the space of possibility shrinks as more information comes in. The epistemically possible is geared toward the actual, and so—when we are lucky—these will converge. Not so of the objectively possible. It might turn out that Avalonia's systems are not neutrally structured after all, but the above counterfactual conditional is not thereby ruled out. My argument is that sciences of the deep past provide us with knowledge of objective possibility, not merely epistemic possibility. Now, consider the imagination-based accounts on the table.

Both families of imagination accounts—those relying on "conceivability" and those relying on "credibility"—involve a criterion of consistency and a psychologistic criterion: roughly, the relevant representations must be consistent with what else the agent knows, and the agent must grasp them in the right way. This is intended as an answer to the *justification* of the modal knowledge. The consistency aspect appears to simply track epistemic possibility. To be consistent with a set of scientific knowledge *just is* to be epistemically possible vis-à-vis that knowledge. It no doubt, then, confirms epistemic possibility. But does it give us any particular guide to objective possibility? It

might do: if we think our scientific knowledge is really well-grounded, if we think it tracks the structure of the world, then perhaps it does. But in that instance, the justification isn't to do with imagination, but concerns whether we have good scientific knowledge. How about the psychologistic criteria? Does an expert being able to grasp, or run a mental simulation, or draw the right inferences, or in some sense "get" the result, add epistemic weight to a conclusion being about objective possibility? In some cases, perhaps, as a matter of testimony, the intuitive judgements of experts under the right conditions might be a good reason for us to take a hypothesis seriously. But given how epistemic agents are situated within their knowledge, we must ask: what are they running mental simulations of? Well, it would be their understanding of the relevant phenomena under the constraints provided by their background knowledge. In other words, epistemic possibility. So, only when epistemic and objective possibility track one another should we expect imagination to track objective possibility. And, as we've just seen, that appears to only be a matter of the consistency of our background beliefs.

You might worry that my arguing that imagination enjoys no special role in underwriting modal knowledge doesn't undermine the role of imagination if it underwrites knowledge, or at least knowledge generated from models, more broadly, as is often the case for some views, especially those sometimes called "fictionalist" (see Toon 2016; Levy and Godfrey-Smith 2020; Salis 2020). The fictionalist, I think, can happily agree with the negative claim of this section. Typically, appeals to imagination by fictionalists are not appeals intended to directly justify claims derived from models. Rather, imagination matters for fixing a model's *content*. For a fictionalist, claims about models (that, say, Cuthill and Conway Morris' geometric programs represent rangeomorph development via some fictionalist semantics) should be understood as invitations to imagine: "We can regard a model's equations (or even a verbal description of the model) as a prop that, given suitable (scientific) principles of generation, imply what the model's content is" (Levy 2024, 102). Fixing a model's content is crucial for identifying claims derived from the model, but this doesn't say anything in particular about what justifies that content or propositions derived from it. As such, one might entirely go along with my arguments that imagination plays no special role in answering the epistemic question about modal models, and still adopt a fictionalist view (in chapter 7, however, I will disagree with fictionalism about the scientific imagination more generally).

So, at this stage at least, my argument isn't that imagination isn't crucial for knowledge generally, but simply that it plays no special role in justifying claims to objective modal knowledge.

So, answers to the epistemic question about modality that appeal to imagination won't work because (1) imagination is extraneous to justification, (2) they put too much weight on scientist's grasping compared to the work put into designing, tweaking, and making arguments concerning perspectival tools, and (3) they appear to give us epistemic, not objective, possibility.

Happily, there's another view within the modelling tradition that provides, I think, an answer much more aligned with the modal knowledge we've discussed: *artifactualism*.

2.2 Artifactualism Good

Emphasis on imagination—mental simulations, “grasping,” or other psychologicistic accounts of the modal power of perspectival tools—doesn't sit well with the *bespoke* nature of many perspectival tools. Palaeontologists do not simply grab well-known or well-used models and imagine them in new scenarios. Palaeontologists build, adapt, and tweak model design, construction, and deployment toward particular idiosyncratic features of the systems they're trying to understand. As I'll argue in chapter 7, imagination is deeply coupled with perspectival tools. As such, I think an artifactualist account has a better chance of answering questions regarding the justification of modal knowledge.

Artifactualism is a set of views concerning models, elements of which are “implicit, in many practice-oriented approaches to modelling” (Knuuttila 2024, 122). Artifactualists emphasize models' materiality in the context of their use and, in various senses, emphasize the non-representational aspects of models. The view has been largely developed by Tarja Knuuttila (2005, 2011, 2021) with the central contrast being representational, particularly fictional, accounts of models (see also my attempt at a similar view, Currie 2017b, and de Oliveira's more radical form, 2022).

Most accounts of models are *representation-first*. That is, they think of models as primarily standing in some relationship to some actual target system, and ask in virtue of what that relationship is representational. On such views, Cuthill and Conway Morris' model is fundamentally a model of rangeomorph growth. Such accounts start with what Knuuttila has called “external representation”: “the relationship of a model to a real-world target

system” (Knuuttila 2024, 65). Artfactualist accounts begin instead with the model itself. A model is not fundamentally a representation of another system, but a system *itself*—a set of interdependencies, expressed through some (to use Knuuttila’s (2011) language) *mode* and implemented in some *media*. The former is the symbolic devices at play—the language, if you want—and the latter is the material substrate at hand (in 2017, I made approximately the same distinction by distinguishing “vehicle” and “content,” terms that are too representationalist in retrospect). To use a model, scientists must *interact* with it in some way, as Guilherme Sanches de Oliveira puts it:

In order to be used by scientists, even [mathematical] models must be implemented in some way that enables interaction and manipulation, including, for instance, as markings and inscriptions with pen on paper or chalk on a blackboard, or as a programming code typed up and displayed on a computer screen. It might be tempting to think of “the model” as transcending, or being independent from, any particular physical implementation—still it’s precisely as some physical implementation or other that the model enables scientists to intervene in some way (e.g., changing parameters, settings or variable values) and to visualize or otherwise measure the effects of those interventions. (de Oliveira 2022, 4)

Let’s briefly work through an example. A mathematical model will follow some mode. Cuthill and Conway Morris’ (2014) model of rangeomorph development consists of a procedure for generating geometric patterns based on the serial repetition of self-same fractals. This procedure is the *mode*. However, this mode must be realized in some *media*. Cuthill and Conway Morris used a computer, but in principle they could have used pen-and-paper, or straws, or imagined it in their heads. In a certain sense, the various media are different ways of “instantiating” or “realizing” the mode, but each medium has very different properties that are often highly relevant to its epistemic possibilities. Computational media are constrained by hardware and software and are often able to generate a large number of runs on comparatively short time-scales. Pen-and-paper media leave much more open to human calculation error and take much more time—however, in computational cases, human error in the form of coding mistakes is often much, much harder to discover.

When it comes to modal knowledge, artifactualist accounts—whether “hybrid” versions that make concessions to representation (as favoured by myself and Knuuttila) or the more “radical” forms defended by de Oliveira—have a crucial advantage over views that appeal to imagination. The various artifacts scientists build are not considered representations in the first instance, but are instead considered “systems-in-themselves.”⁴ As systems-in-themselves that scientists have constructed and examined, their use is primarily intended to establish the dispositions of those objects. That is, they generate knowledge of objective possibility concerning those objects. The trick, then, is to provide good reason to think the relevantly same dispositions will occur in the objects they ultimately want to understand, or in the “fragile systems” they’re targeting.

The artifactualist perspective also provides some clues as to under what conditions modal knowledge might be justified. Knuuttila does appeal to imagination-based accounts in her artifactualist justification of modal models: “The question of conceivability is crucial for modal inferences” (Knuuttila 2024, 9). However, the way this plays out in her discussion is less about mental simulations and more about serious study of “whether the possible systems studied by modelling could be actualized” (Knuuttila 2024, 11). As she points out, often synthetic biologists attempt to determine whether theoretical models can be built using organic components, explaining how the various media through which the structures are realized enable different kinds of knowledge. I have made a similar point about how ecologists navigate between highly abstract models, field data, and simple “bottle” experimental systems:

A coupled differential equation is one way of representing a set of assumptions about the causal interactions between trophic levels. Positing such an equation sets a particular modal domain, and allows us to explore it. That is, it underwrites our learning how to navigate the domain specified by the equations. We see the same in bottle experiments. But—crucially—the mediums are different, the way we specify our assumptions are different, and both the modal domains, and the way we

4 This term—“systems-in-themselves”—might be somewhat misleading, as it might imply that they do not have functions. As artifacts, they *do* have functions. However, these are often open-ended, in flux, and, crucially, they are not restricted to *representational* function (Currie 2017b; Knuuttila 2024).

explore and navigate these domains are different: thus, I argue, different understandings are generated. (Currie 2020, 924)

For artifactualists, modelling practices often involve using different media to realize “the same” mode. This is done to establish the more robust dispositions across those systems of interest. In light of the book thus far, we already have the means of describing this: *anchoring*. When we’re trying to understand lost worlds (and I bet in many other scientific circumstances), toy models and other scientific artifacts are other anchors across which we iterate.

A useful way of making sense of artifactualist epistemology is Catherine Elgin’s notion of “exemplification” (Elgin 2017). Something *exemplifies* a property when it possesses that property. In virtue of having that property, examining that something can be a way of demonstrating features and dispositions associated with that property. I might represent a *wing* in tap-dancing verbally: the dancer scrapes their foot out to one side, then quickly hits the floor before stepping down. Alternatively, Kirsten might exemplify a wing by, well, putting on her tap shoes and performing one. Similarly, Cuthill and Conway Morris’ geometric model exemplifies some properties of rangeomorph development in virtue of in fact possessing those properties. Exemplifying and examinations of exemplifications can act as demonstrations of, for instance, modal properties of the relevant objects.

So, here is the basic strategy for generating modal knowledge according to an artifactualist account. In standard, say, experimental settings, we ask after the projectability of results, the representativeness of the specimen, and so on. The basic answer to these questions involves finding some warrant, a set of reasons, to think that the specimen and results are sufficiently similar to our ultimate target. Steel (2008), for instance, relies on what he calls *comparative process tracing*. This involves identifying the mechanisms responsible for experimental behaviour and tracking whether the same mechanisms are present in the target. The artificed dinosaur burrow and careful comparison between it and reconstructions of the Jurassic burrow are an example. Now, having identified a disposition in the system we are examining—mapping out a modal structure for it—we then, in just the same way, ask if that modal structure is representative (and *how* representative it is) of the non-experimental structures we’re interested in. If we’re thinking of exemplification, we’d here ask whether the exemplified properties are in fact the ones we’re interested in, in the system of interest.

You might wonder what account of justification I have in mind. I don't take myself to have committed to any particular position between, say, internalist accounts that take justification to lie in the beliefs of agents or externalist accounts that incorporate aspects of the world into justification. I think that either internalists or externalists can make sense of what I've said, but I suspect my externalist leanings are probably quite obvious throughout the discussion.

One way of reconciling appeal to epistemic concepts like imagination and artifactualism is to tie imagination closely to "understanding," but rather than appeal to "theoretical" conceptions of understanding (e.g., De Regt 2017), appeal to more embodied, situated, or "know-how" conceptions of understanding. On these views, what it is for a scientist to understand (and, let's claim, "imagine") is to have certain abilities or capacities (see Leonelli 2009; Le Behain 2016; Currie 2020). As Leonelli puts it: "Understanding can only be qualified as 'scientific' when obtained through the skilful and consistent use of tools, instruments, methods, theories, and/or models: these are the means through which researchers can effectively understand a phenomenon as well as communicate their understanding to others" (Leonelli 2009, 190).

Embodied forms of understanding marry naturally with artifactualism. Scientific understanding acts through the objects scientists construct and interact with. It is the capacity to do things with these objects: to draw inferences, make predictions about their behaviour, manipulate them in predictable ways, and so on . . . that constitutes understanding. If imagination is similarly embodied, then it too folds into an artifactualist account of modal knowledge. In chapter 7, we'll return to these ideas, but for now I'm left to draw some general lessons and apply them to a final case.

3. The Science of Lost Worlds Is a Science of Objective Possibility

Scientific justification, at least of the kind of modal knowledge that concerns us here, doesn't happen inside scientists' skulls. Justification lies in intimate connection with artifacts ranging from material specimens to theories to experiments and other kinds of apparatus. In light of loss or the possibility of loss, historical scientists construct perspectival tools: traces, epistemic artifacts, material objects—media—that realize a mode. Demonstrations of these artifacts' dispositions ground scientists' knowledge of the dispositions of the systems that interest them. These artifacts have genuine modal properties,

and in virtue of deep understanding—when we have it—of the relationships between those artifacts and the systems they model, those modal properties can be projected via inference onto the past. Thus, when the conditions are right, historical scientists don't simply tell us what actually happened, but can tell us what might, could, or would have happened under different circumstances. Further, those artifacts often instantiate complex, multi-layered modal spaces, which not only enable testing various hypotheses about the actual past, but in themselves represent genuine epistemic achievements.

The argument from the last section, at base, claimed that imagination-based accounts of the justification of modal models—even if they work for the toy models they target—can't make sense of the practices of scientists interested in uncovering lost worlds. Modal claims are justified not based on some scientific capacity to conceive of possibilities, but via demonstration in an artifact that exemplifies the properties of interest. The results of said demonstration are then transferred to the target system or system type.

On this view, no special answer or pleading, nor an appeal to imagination, is required to justify modal knowledge.

I've argued that to understand the sciences of the deep past, you should understand them as tackling two interwoven challenges. First, erasure: information from the past is often fragile and incomplete, typically requiring interpretation through complex chains of reasoning and rich background theory. Second, loss: past worlds are sometimes deeply unfamiliar due to containing unseen, mysterious denizens. Historical scientists always navigate challenges from at least the possibility of loss. Both loss and erasure contribute to the heavily modal nature of scientists' work. Tackling erasure requires understanding the dependencies between traces and the past denizens they are sourced from, as well as various forces and sources of bias and information loss; at base, it requires understanding the dynamics of information loss and decay. Tackling loss (and its possibility) involves understanding the fragility of systems. The well-mixed oxygenation of the oceans is crucial for supporting life as we know it, but that turns on the biological pump generated by the daily migration of all of those aquatic creatures, and that pump itself has a history.

Faced with loss and erasure, then, scientists of the deep past construct artifacts that—they hope—exemplify the kinds of properties they need to understand to make sense of past worlds. These artifacts are (1) applications of strategic perspectivalism: they isolate particular properties; (2) bespoke: they are coopted and geared toward local and often highly particular questions;

(3) pluripotent: while informing us about various instances, they also explore and probe fragile systems, dynamics, and regularities (that is, possibility); and (4) iterative: they are in dialogue with other artifacts, with traces, and often with observational data from our contemporary world. These iterative explorations grant modal knowledge about the conditions under which phenomena occur, their persistence, and so on. This modal knowledge is not merely epistemic, but objective.

This set of claims underwrites another regarding the nature of success in the historical sciences. As Derek Turner emphasizes, we will likely get many details concerning past worlds wrong (Turner 2004, 2005, 2007). The dinosaurs may have been located outside of the burrow, rangeomorphs may not have been obligate osmotrophs, Avalonian ecosystems may not have been neutrally structured, and the emergence of metazoa may not have turned on the emergence of key innovations that structured their oxygen environments. But this is just to say we didn't get the *actual* world right.

We may still, nonetheless, have gained crucial insights into modality. We may come away knowing more about the relationship between burrow infill and bone placement, the properties and limits of obligate osmotrophy, the dynamics and conditions under which neutral ecosystems occur, and how organisms can influence oxygen flow.

So, the sciences of past worlds are sciences of possible worlds, and that possibility is objective. It is objective both in the sense of going beyond epistemic possibility and in the sense of being objectual, that is, grounded in the dispositions of objects.

4. Astrobiology and the Actual in Service of the Possible

In chapter 2, I introduced problems generated from loss via an analogy to astrobiology. Many astrobiologists examine the atmospheric composition of exoplanets (inferred from traces in the form of wavelengths detected by telescopes) in an attempt to identify “biosignatures” (that is, atmospheric compositions that require feedback from a biosphere). There is a sense in which they have a single N of life: that found on Earth. This raises a question: whether or not life on our planet is representative of life generally. If we get that wrong, our characterization of biosignatures could be too narrow (hence generating a bunch of false negatives) or too broad (generating false positives). As I argued, only having 1 N is not so much a problem concerning the amount

of specimens or data, but a problem of being able to generate sufficiently rich background knowledge required to make sense of data. If astrobiologists had a bunch of living planets to examine, they would be much better positioned to develop the background theory required to detect unambiguous biosignatures (if such a thing can be had).

Astrobiology is often characterized as a multidisciplinary (or at least a multiple-lines-of-evidence) effort to discover extraterrestrial life. This is typically tied to understanding the origins of life on our own planet. NASA emphasizes this aspect, for instance, characterizing astrobiology as “the broad-based study of the origins of life here and the search for life beyond Earth” (Kaufman 2022). It is, then, at least doubly historical: by understanding various processes that operated in Earth’s *past* astrobiologists hope to develop ways of detecting life from exoplanet *traces*; they’re further interested in the *pasts* of neighbouring planets and the possibility of their having housed life or having once had conditions relevant to it (hence the importance of the “warm & wet” early Mars hypothesis).

So, we can more or less divide astrobiology into two potentially quite different but mutually informing tasks. First, what kinds of processes lead to life? That is, under what conditions might life arise? Second, what signals might there be of extraterrestrial life from exoplanets, that is, can we develop an understanding of biosignatures? I’ll focus on this latter question.

As I’ve indicated, much philosophical discussion of astrobiology, and epistemic worries from astrobiologists themselves, turns on the extent to which we can use life on Earth as something like a representative sample when understanding life in the universe. The only specimen we understand in any detail is Earth and Earth’s history. So, how can we know whether biosignatures developed based on our understanding of Earth have the right scope? If they’re too narrow, we’ll potentially miss a lot of life; if they’re too broad, we’ll potentially misidentify it. Worse, atmospheric composition underdetermines the processes that maintain or generate them. Earth’s oxygen-richness is due in part to biotic processes, but on other planets, oxygen-richness could be due to abiotic processes (Meadows et al. 2018). Further, as various sinks could mask oxygen production, a *lack* of oxygen is also no indicator that earth-like life is not present (indeed, it’s only been in the last 400 million years or so, a tenth of its life, that Earth would have produced a stable oxygenated signal).

The basic shape of this problem should be familiar to taphonomists. Many mineralogical forms could be produced by biological or purely geological origins. Thus, telling whether a weirdly-shaped rock’s properties are indicative

of biological or geological signals turns on having a rich understanding of the various ways that various processes could generate those properties. In most circumstances, trace-based reasoning doesn't occur independently of analysis of specific details of context. The same holds for biosignatures: biosignature interpretation relies on "understanding a planetary environment to identify characteristics that may potentially reduce our ability to see a biosignature, even if a strong biological source exists" (Meadows et al. 2018, 632). But this shouldn't be a surprise: it is another way in which recontextualization is crucial for successful recovery of lost or putatively lost worlds.

Notice that I have presented biosignature research as involving modality, but primarily in the service of the actual: we're interested in the dependency relationships between exoplanet atmospheric composition and possible biospheres because we want to know whether that planet has life. But we needn't view astrobiology in these terms. We might be interested in how different atmospheric compositions might come about, in mapping kinds of possible atmospheres, and the kinds of biotic and abiotic processes that could generate and maintain them. To illustrate this and see why it matters, let's turn to a recent example.

Nicholson et al. (2022) are interested in understanding biosignatures in the context of "nutrient-limited" biospheres. Their study focuses on conditions of the early Earth during the Archean, before the "Great Oxygenation Event," when our planetary biota was largely dominated by methanogens. "Given that Earth spent roughly a third of its lifetime in the Archean, it is natural to begin our study of the vast possibilities for biosignatures with this long-lived and comparatively simple biosphere [references omitted]" (Nicholson et al. 2022, 222). The aim of their modelling is to provide environmentally relevant biosignatures for planets in that kind of state. The approach is to use a relatively simple model of methane exchange, but let the biological aspects act as free parameters, evolving and interacting with the atmosphere. As such, they hope to map out a modal space consisting of possible arrangements of atmospheric composition and biological states.

In a nice example of artificing (and very characteristic of astrobiology), Nicholson et al.'s model consists of an atmosphere and an ocean, with three chemicals circulating between them (hydrogen, carbon dioxide, and methane). This simplicity is important for their goal; they characterize their target abstractly to highlight—exemplify—the relevant factors: "As we are interested in the overall behaviour of the simple life-environment coupled system and are not trying to recreate the climate of a real planet, we use these simplifications

to keep the abiotic environment simple” (Nicholson et al. 2022, 225). It’s also another example of how even abstract models are rendered highly bespoke in strategically perspectival contexts, to facilitate the particular kinds of isolation and eventual recontextualization crucial for the investigation’s aims.

Microbes are modelled as a single species that metabolizes carbon dioxide and hydrogen into methane and water. This generates adenosine triphosphate (ATP), which is used to generate biomass. Microbe population dynamics are determined by ATP availability: at each timestep, some proportion starve due to a lack of ATP, and some proportion reproduce due to sufficient ATP (and some proportion do neither). The biological and abiological chemical composition is coupled in various ways: most obviously by the availability of ATP, but also by, for example, the “burial rate,” wherein microbe deaths act as a hydrogen sink as they are buried under the ocean floor. The basic experimental setup is to run the abiotic system until it reaches an equilibrium, and then introduce the microbes, varying factors like death rate, and how ATP-greedy the microbe is.

The major result of the various studies is that atmospheric methane remains largely insensitive to biological action, but is highly sensitive to oceanic hydrogen. Thus, variation in methane in such a planetary system can be chalked up to hydrogen variation regardless of the processes of microbes. The upshot of this is that population dynamics are less important than nutrient-availability. On planets like the early Earth, it is oxygen-availability that dampens the system: “When considering a nutrient-limited biosphere it is more important to accurately model the processes that regulate the availability of the limiting nutrient, and determine the limit to which life can exploit this nutrient, than it is to model any specific population dynamics for that biosphere” (Nicholson et al. 2022, 237). For the direct purposes of developing an atmospheric biosignature for this type of planet, this isn’t great news: detecting oceanic oxygen and factors affecting it (volcanism, for instance) on exoplanets seems like a big ask, to say the least.

How should we characterize the epistemic dividends of work like Nicholson et al.? Their aim is certainly to help develop biosignatures: that makes sense of the various decisions made in constructing their simulation. But what we learn doesn’t only pertain to detecting biosignatures. We’re also learning, at least potentially, about the kinds of dynamics that might rule atmospheric composition and various other global systems. We’re mapping out possible planets. And this mapping isn’t merely in service of the actual:

astrobiologists make sense of complex dynamics and dependencies between global systems. It is, at least in principle, a science of objective possibility.

There is no doubt, then, that astrobiology can be characterized in part as the search for extraterrestrial life. But to leave our description there is to do it an epistemic disservice. Astrobiology doesn't simply ask whether we can detect life on other planets. At its best, astrobiology is interested in the space of possible planets: what kinds of atmospheres could there be, and how could those atmospheres be generated and maintained? But this exploration isn't empirically unrooted—it is not a mere modelling exercise—but is anchored by empirical examples, from the signatures of distant exoplanets to the current and past Earth.

There's a more general lesson here concerning the dividends of historical science. Typically, the explicit aim is to answer questions about the actual past—for example, did this dinosaur burrow? Were rangeomorphs governed by ecological dynamics different from those of contemporary ecosystems? Was there life on distant planets? But the dividends of the study are modal spaces: dependencies between bone position and burrow infill, ecosystem properties that determine neutral spatial distribution, possible constellations of atmospheric composition and underlying biotic and abiotic processes. Perhaps ironically, the focus on the actual provides anchors for the exploration of possibility.

You might complain here: surely the scientists take themselves to be trying to uncover biosignatures, not trying to explore conditionals concerning atmospheres and possible biospheres? And indeed, this being an aim of research like Nicholson et al.'s needs to be taken seriously to understand why they undertake the research in the way that they do, why various aspects are salient, and so on. I agree with this. But I distinguish between the *aim* of an investigation and the *dividends* of an investigation. Dividends can be by-products, unintended upshots, or unexpected windfalls. Scientists needn't be aiming for dividends that they nonetheless achieve.

In a certain sense, then, we might think of these cases as not only involving the possible in service of the actual, but the actual in service of the possible.